Centrum für Europäische Politik
FREIBURG | BERLIN

cepInput

No 12 | 2025 15 July 2025

The Algorithmic Hand

How Large Language Models Disrupt Competition and Democracy

Anselm Küsters

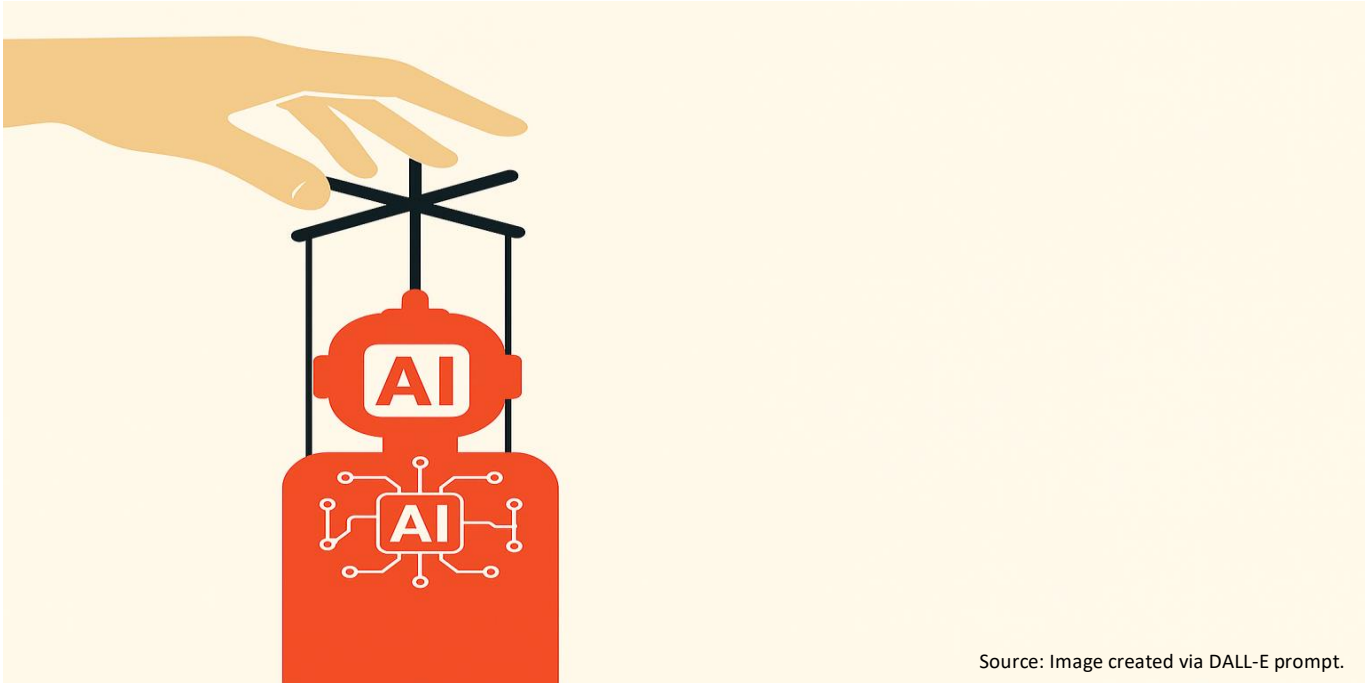

Source: Image created via DALL-E prompt.

Large language models may soon evolve from helpful search companions into information gatekeepers, quietly influencing what people buy, believe, and vote for. As this platform shift unfolds, Europe's ordoliberal promise of dispersed market power and democratic debate is confronted by an AI stack that learns and then amplifies our collective biases. This cepInput asserts that unless the competition rules are redesigned and the EU acts swiftly, the "invisible hand" could become a private entity that serves only a few Big Tech firms.

- Three LLM bias vectors – namely skewed training data, corporate safety-layer filtering, and persuasive framing in the model's tone – reinforce each other to influence economic choices and political judgements. Empirical evidence from several technical LLM experiments and examples related to misinformation campaigns demonstrate how easily these vectors can undermine consumer and citizen sovereignty.
- Consequently, this study formulates a new antitrust "theory of harm": LLMs combine market dominance, informational asymmetry, and the ability to set norms, creating "algorithmic capture" that blurs the line between economic and political power. While antitrust tools remain necessary, they must be supplemented by new metrics that register discursive plurality and new measures that tackle attention bottlenecks.
- On this basis, the study suggests a two-pronged approach: robust enforcement of merger control, DMA, DSA, and AI Act to maintain market competitiveness, and "utility engineering" combined with citizen assemblies to align model value functions with democratically chosen norms. Investments in domestic open-source models form part of a "EuroStack" roadmap that helps restoring the benefits of Smith's invisible hand.

Content

1 Introduction: from search engines to algorithmic hands

To google or not to google? In a relatively short amount of time, the familiar way of using the web, i.e. typing a query into Google (or any other search engine) and reading through the ranked hyperlinks, has begun to disappear. Conversational AI chatbots powered by large language models (LLMs), such as ChatGPT or Perplexity, respond with complete sentences, follow-up questions, and targeted recommendations, meaning the user never needs to leave the chat window. A representative survey in the US in December 2024 found that 27% of respondents already "use AI tools instead of search engines",[1] a figure that would have been unthinkable even two years earlier. While adoption is more prevalent among younger age groups, European polling suggests a similar trend.[2] The consumer research and consumer survey that the British Competition and Markets Authority recently commissioned indicate that "for some users and some use cases, AI assistants such as Gemini are beginning to be consumed in a way akin to general search".[3] The research also found that most AI assistant users anticipate increasing their use of AI assistants for "search" tasks in the future.[4] A so-called "platform shift" – from web search to LLM-based assistance – may thus be only a matter of time.[5] But what will happen when the informational infrastructure on which markets and, arguably, democracies rely is no longer a polycentric network of websites, but rather a small group of probabilistic text generators?

From an ordoliberal perspective based on free markets, rule of law, and a "competition democracy nexus" that asserts a positive link between unbiased competition and democratic policymaking,[6] this technological inflection point carries constitutional weight. Adam Smith's famous "invisible hand" metaphor is based, despite many misunderstandings,[7] on the key idea that no single entity controls price signals. These signals then lead to allocative efficiency when utilised by market participants. On this basis, ordoliberalism maintains that competition in goods and ideas is beneficial when power is dispersed. However, the Freiburg tradition has always recognised that technical change can upset this delicate balance.[8] Scholars such as Franz Böhm and Hans Großmann-Doerth warned that markets collapse into political privilege whenever private intermediaries combine economic and norm-setting authority.[9] As the following sections show, such a fusion is evident today, with large language models

---

1 See: https://www.techradar.com/tech/people-are-increasingly-swapping-google-for-the-likes-of-chatgpt-according-to-a-major-survey-heres-why.

2 See: https://www.eu-startups.com/2025/05/how-europe-views-ai-insights-from-our-polls-and-expert-reactions/.

3 See: CMA, Proposed decision (24.6.25), p. 33.

4 Thinks Insight & Strategy qualitative consumer research report, Exploring consumers' search behaviour, para. 6.6.

5 See also: Salomé Viljoen, Jake Goldenfein, and Lee McGuigan, 'Design Choices: Mechanism Design and Platform Capitalism', *Big Data & Society* 8, no. 2 (July 2021): 1–13, https://doi.org/10.1177/20539517211034312; Andrei Hagiu and Julian Wright, 'Multi-Sided Platforms', *International Journal of Industrial Organization* 43 (November 2015): 162–74, https://doi.org/10.1016/j.ijindorg.2015.03.003.

6 Elias Deutscher and Stavros Makris, 'Exploring the Ordoliberal Paradigm: The Competition Democracy Nexus', *Competition Law Review* 11, no. 2 (2016): 181–214. For ordoliberalism in general, see: Josef Hien and Christian Joerges, eds., *Ordoliberalism: Law and the Rule of Economics* (Oxford: Hart Publishing, 2017); Malte Dold and Tim Krieger, eds., *Ordoliberalism and European Economic Policy: Between Realpolitik and Economic Utopia*, Routledge Studies in the History of Economics 225 (New York: Routledge, 2019).

7 Bas Van Bavel, *The Invisible Hand* (Oxford: Oxford University Press, 2016).

8 Walter Eucken, 'Technik, Konzentration Und Ordnung Der Wirtschaft', *ORDO: Jahrbuch Für Die Ordnung von Wirtschaft Und Gesellschaft* 3 (1950): 3–17.

9 Franz Böhm, 'Das Problem der privaten Macht. Ein Beitrag zur Monopolfrage', in *Grundtexte zur Freiburger Tradition der Ordnungsökonomik*, ed. Nils Goldschmidt and Michael Wohlgemuth, Untersuchungen zur Ordnungstheorie und Ordnungspolitik 50 (1928; repr., Tübingen: Mohr Siebeck, 2008), 49–70; Hans Großmann-Doerth, 'Selbstgeschaffenes Recht der Wirtschaft und staatliches Recht', in *Grundtexte zur Freiburger Tradition der Ordnungsökonomik*, ed. Nils Goldschmidt and Michael Wohlgemuth, Untersuchungen zur Ordnungstheorie und Ordnungspolitik 50 (1933; repr., Tübingen: Mohr Siebeck, 2008), 77–90.

integrating retrieval, ranking, and framing of information signals within a single pipeline. With the large-scale integration of LLMs into our economic and societal life, this paper argues, the invisible hand is no longer unbiased. The ordoliberal project of safeguarding both consumer and citizen sovereignty in an "economic constitution" must therefore confront a novel intermediary: the algorithmic hand.

The economic mechanics of this shift are already evident. While the traditional "attention economy" in Web 2.0 extracted value by keeping users looking at screens, LLM interfaces aim to capture the user's next action, i.e. to turn attention into intention.[10] As generative AI models learn detailed preference maps from ever larger sets of training data as well as conversational history and autonomous actions in the web, they can predict and influence purchase or voting decisions with far greater accuracy than display ads or sponsored links.[11] What appears to be a neutral answer engine is actually a device that shapes preferences and is situated at the heart of tomorrow's digital architecture. According to ordoliberal theory, such bottlenecks are potential sources of restrictions on the competitive process and thus call for robust *Ordnungspolitik*. It is precisely the intermingling of economic and societal power in this type of technology that is the key problem from an ordoliberal perspective, as it undermines the separation of powers on which the invisible hand depends. LLMs will combine market dominance (control over web traffic), informational asymmetry (knowledge of each user's cognitive profile), and norm-setting capacity (the ability to privilege certain values). By doing so, they threaten free price formation and open democratic deliberation – twin pillars of the social market economy.

The remainder of this paper therefore treats LLMs as constitutional actors rather than as neutral tools. Section 2 sketches how these models work and maps the three systemic bias vectors — inherent training data imbalances, corporate filtering, and framing effects — that cause the algorithmic hand to be partial. Later sections will demonstrate how these biases result in market foreclosure, attention capture, and electoral distortion (sections 3-5). Finally, this paper will outline potential remedies, ranging from enhanced merger scrutiny to the emerging science of "utility engineering", culminating in an ordoliberal policy blueprint for the EU (sections 6-7). Before exploring these remedies, it is crucial to understand precisely how bias enters the models that will soon mediate many spheres of modern life.

2 Mapping systemic bias vectors in large language models

LLMs do not enter the marketplace as neutral devices. Before a response appears on the user's screen, the model has already passed through several decision gates that influence its outputs. To understand why these gates exist, we need a clearer idea of how LLMs actually work. These models process text using a small set of well-established techniques.[12] First, a text stream is segmented by tokenisation, breaking words into shorter units that the system can handle efficiently. Each token is then mapped to a numerical vector, or "embedding", whose position is refined during training so that tokens occurring in similar contexts cluster in the same region of a high-dimensional space.[13] Inside the so-called transformer architecture, which forms the basis of ChatGPT and Co, stacks of "attention layers" assign

[10] Yaqub Chaudhary and Jonnie Penn, 'Beware the Intention Economy: Collection and Commodification of Intent via Large Language Models', *Harvard Data Science Review*, no. Special Issue 5 (30 December 2024), https://doi.org/10.1162/99608f92.21e6bbaa.

[11] See the subsequent sections and the examples and literature cited therein for evidence.

[12] Torsten Hiltmann, 'Hermeneutik in Zeiten der KI. Large Language Models als hermeneutische Instrumente in den Geschichtswissenschaften', in *KI:Text. Diskurse über KI-Textgeneratoren*, ed. Gerhard Schreiber and Lukas Ohly (Berlin, Boston: de Gruyter, 2024), 203–34, https://doi.org/10.1515/9783111351490.

[13] Chongyang Tao et al., 'LLMs Are Also Effective Embedding Models: An In-Depth Overview' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2412.12591.

greater importance to the tokens that are most relevant for predicting the next one.[14] This allows the model to capture context, i.e. both syntax and long-range dependencies in a text. Training the neural networks underlying these models involves adjusting millions of internal parameters so that, given any partial sequence of tokens, the model assigns a high probability to an appropriate continuation. During the deployment of the resulting LLM, this prediction step is repeated continuously. This enables the LLM to respond to user questions by appending one word (token) at a time until the reply is complete.

At least three key decision gates lie between the model architecture and what the user ultimately sees – and these gates can introduce bias. To begin with, the first bias vector originates in the training data used for the machine learning: the corpus. In the context of today's gigantic models, this corpus basically consists of the whole content of the internet. However, as many critics have pointed out, public web scraping over-represents English-language sources and Western discourse.[15] This means that token frequencies already exhibit a political bias before the first parameter is trained. Systematic reviews confirm the downstream effect: larger frontier models reliably provide answers to US policy questions from a centre-left perspective and disproportionately focus on American affairs, while newer multilingual capacity introduces further asymmetries.[16] Even if developer would attempt to remove these biases by expanding the corpus, the way word embedding works still tends to reinforce stereotypes; for instance, female names are clustered (in the high-dimensional space) nearer to caretaking verbs.[17] In ordoliberal terms, one could say that using any web-scraped corpus as training data acts like a distorted price signal, misrepresenting the true distribution of social preferences and thus distorting processes, such as competition, that depends on accurate information. Crucially, increasing model size does not dilute these biases.[18] In short, the invisible hand learns from a distorted reflection of society.

A second bias vector emerges after training but before the tokens reach the user in the form of input and output filtering. For instance, take the case of the popular reasoning model "R1" developed by Chinese AI firm DeepSeek: When the same model weights are run locally on a laptop, the answer to "Is Taiwan a nation state?" is nuanced and recognises de facto sovereignty.[19] When the query is routed through DeepSeek's official app, the reply becomes vaguer. When the prompt passes via a third-party wrapper that combines proprietary retrieval with additional guardrails (such as Perplexity), the answer is so evasive that it loses all political substance. It is not necessarily the neural network that varies, but the amount and design of filters, such as blocked web domains and "system prompts" that are injected between the user and the model.[20] As these filters are privately governed and dynamically updated,

[14] Kevin Clark et al., 'What Does BERT Look At? An Analysis of BERT's Attention' (arXiv, 2019), https://doi.org/10.48550/ARXIV.1906.04341; Ashish Vaswani et al., 'Attention Is All You Need' (arXiv, 2017), https://doi.org/10.48550/ARXIV.1706.03762.

[15] This section is based primarily on: Shangbin Feng et al., 'From Pretraining Data to Language Models to Downstream Tasks: Tracking the Trails of Political Biases Leading to Unfair NLP Models' (arXiv, 2023), https://doi.org/10.48550/ARXIV.2305.08283; Roberto Navigli, Simone Conia, and Björn Ross, 'Biases in Large Language Models: Origins, Inventory, and Discussion', *Journal of Data and Information Quality* 15, no. 2 (30 June 2023): 1–21, https://doi.org/10.1145/3597307.

[16] Yejin Bang et al., 'Measuring Political Bias in Large Language Models: What Is Said and How It Is Said' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2403.18932.

[17] Tolga Bolukbasi et al., 'Quantifying and Reducing Stereotypes in Word Embeddings' (arXiv, 2016), https://doi.org/10.48550/ARXIV.1606.06121.

[18] Ze Wang et al., 'Bias Amplification: Large Language Models as Increasingly Biased Media' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2410.15234.

[19] This is based on an experiment reported by Azeem Azhar and Nathan Warren on January 29, 2025 (blog title: „DeepSeek's shock: 9 critical things you need to know"), and own interpretation.

[20] Anselm Küsters and Manuel Wörsdörfer, 'Exploring Laws of Robotics: A Synthesis of Constitutional AI and Constitutional Economics', *Digital Society* 4, no. 46 (4 June 2025): 1–33, https://doi.org/10.1007/s44206-025-00204-8; Pengfei Liu et al.,

they create a sort of "informational tollgate" that is hard to scrutinize from the outside. Furthermore, these so-called "safety layers" often merely reflect the geopolitical interests of the host company's jurisdiction: References by the R1 model to Tiananmen Square, for example, disappear in some regions (China) but not in others. This has the effect of imposing a soft law of private censorship that sits outside democratic accountability, shaping both consumer choices and public debate. Power thus migrates from the decentralised internet to a handful of LLM stack owners who effectively decide which queries and which answers are acceptable.

Thirdly, even when the factual content remains unchanged after input and output filtering, the way in which an LLM *phrases* that content can influence perception. A study by Buyl et al. demonstrates that identical prompts in English and Chinese elicit different moral evaluations of individuals such as Edward Snowden, revealing that stylistic assumptions, not just factual beliefs, vary across language contexts.[21] Similarly, German researchers found that larger models such as Llama-3-70B, when prompted in German, tended to align more closely with the platforms of left-of-centre parties, whereas the same models queried in English tended to remain neutral.[22] Complementary experiments show that "Western-centric AI models homogenize writing toward Western norms", based on a cross-cultural controlled experiment with 118 participants from India and the US who completed specific writing tasks with and without AI suggestions.[23] From an ordoliberal perspective, this framing bias is especially harmful because it acts like a hidden subsidy in a market: preferences appear to be the consumer's or citizen's own, yet they have actually been influenced by the intermediary's stylistic defaults. Thus, the marketplace of ideas, which in theory should be self-correcting, risks being quietly captured by LLM rhetoric.

Note that these bias vectors compound. When a contaminated corpus meets filtering rules and persuasive framing, different types of falsehood are presented with compelling fluency. Most users lack the time and expertise to reflect on or even audit LLM probabilities, so these bias goes largely uncontested. Indeed, as the literature notes, scale often amplifies rather than mitigates distortion because larger models pick up fine correlations, including spurious ones, from their vast input space.[24] Once a bias takes hold, it reinforces user beliefs, which then feed back into the data stream through clicked links and re-queried prompts. Taken together, these bias vectors of LLMs threaten the ordoliberal aspiration to keep economic and political power separate, as argued next.

3 Economic repercussions: from attention to the intention economy

Digital capitalism has often been described as a surveillance system in which firms monetise the finite time and cognitive capacity that users devote to screens.[25] Technology analyst Dan Hayes refers to attention as "the defining resource of our age", suggesting that, in highly competitive contexts,

'Pre-Train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing', *ACM Computing Surveys* 55, no. 9 (30 September 2023): 1–35, https://doi.org/10.1145/3560815.

[21] Maarten Buyl et al., 'Large Language Models Reflect the Ideology of Their Creators' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2410.18417.

[22] Luca Rettenberger, Markus Reischl, and Mark Schutera, 'Assessing Political Bias in Large Language Models' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2405.13041.

[23] Dhruv Agarwal, Mor Naaman, and Aditya Vashistha, 'AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances', 2024, https://doi.org/10.48550/ARXIV.2409.11360.

[24] Wang et al., 'Bias Amplification'.

[25] Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (London: Profile, 2019). For a balanced assessment of this viewpoint, see: Peter Königs, 'In Defense of "Surveillance Capitalism"', *Philosophy & Technology* 37, no. 4 (December 2024): 122, https://doi.org/10.1007/s13347-024-00804-1.

companies will "select for involuntary methods of capturing attention".[26] Soon, however, attention will be merely the first step in a more complex process: once a conversational agent powered by a LLM has captured the user's attention, it can direct their next action and convert their attention into intention: LLM developers and downstream firms will thus progress to "hyper-personalised manipulation via LLM-based sycophancy, ingratiation, and emotional infiltration".[27] We can thus juxtapose the old "attention economy" with the new "intention economy", where natural-language interfaces allow increasingly detailed categorisation of activity, providing unparalleled control over subsequent behaviour. Since Hayek, Ordoliberals have always viewed competition as a discovery process in which prices aggregate dispersed knowledge.[28] However, once LLMs begin shaping not just what we look at, but also what we do next, the informational preconditions of this process are fundamentally altered.

This mechanism can already be seen in the way that ChatGPT and novel AI agents are reorganising web traffic. When a user requests a product recommendation (e.g. vis Google's "AI Overviews"), rather than returning a list of links, the model now offers a narrative summary that is annotated with images and persuasive language, thereby keeping the user within the (chat) window. According to a leaked presentation, OpenAI's "publisher programme" apparently offers participating media outlets "priority placement" and "richer brand expression".[29] Although ChatGPT placements are not currently sponsored, the interface already favours members over non-members (compare, e.g., the type of sources linked in ChatGPT search), thus foreshadowing paid ranking. In effect, the LLM becomes a vertically integrated gatekeeper: it captures demand, curates supply, and can charge businesses access fees to avoid being relegated to the informational hinterland. Each redirected click represents a loss to the current online ad ecosystem[30] and a gain for the model owner's proprietary ecosystem.

From an ordoliberal perspective, the danger lies, above all, in the collapse of the functional separation between market layers that are intended to discipline one another. Traditional search engines at least made ranking criteria somewhat transparent and allowed users to click through different pages. In contrast, conversational AI agents fuse retrieval, selection, and rhetorical framing inside a single, hidden pipeline. We (typically) get one answer, via chat or voice (instead of ten blue links as in the case of the early Google website). This results in a form of "algorithmic foreclosure"[31] that eludes conventional market-share metrics: a company may control only a minority of advertising revenue, yet if its model sits at the decisive intention bottleneck, it can still influence competitive outcomes. Since LLMs learn the detailed preference vectors of each user, they might also discriminate prices with a precision that is unfamiliar to competition law. The very efficiency of the LLM-based system might therefore erode the spontaneous order that Hayek regarded as the constitutional core of a market economy.

---

[26] Christopher Hayes, *The Sirens' Call: How Attention Became the World's Most Endangered Resource* (New York: Penguin Press, 2025).

[27] Chaudhary and Penn, 'Beware the Intention Economy'.

[28] Friedrich A. Hayek, 'Competition as a Discovery Procedure', trans. Marcellus S. Snow, *The Quarterly Journal of Austrian Economics* 5, no. 3 (September 2002): 9–23. For the influence of Hayek's competition thought on ordoliberalism, see: Anselm Küsters, 'Ordering *ORDO* : Capturing the Freiburg School's Post-War Development through a Text Mining Analysis of Its Yearbook (1948–2014)', *Jahrbuch Für Wirtschaftsgeschichte / Economic History Yearbook* 64, no. 1 (25 May 2023): 55–109, https://doi.org/10.1515/jbwg-2023-0004.

[29] Mark Stenberg, 'Leaked Deck Reveals How OpenAI Is Pitching Publisher Partnerships', *Adweek*, 9 May 2024, https://www.adweek.com/media/openai-preferred-publisher-program-deck/.

[30] For the background, see, e.g.: Emmanouil Papadogiannakis et al., 'The Devil Is in the Details: Analyzing the Lucrative Ad Fraud Patterns of the Online Ad Ecosystem' (arXiv, 2023), https://doi.org/10.48550/ARXIV.2306.08418.

[31] See: Katia Schwerzmann, 'From Enclosure to Foreclosure and Beyond: Opening Ai?S Totalizing Logic', *AI and Society*, 25 November 2024.

European case law already hints at the stakes involved. In the *Google Android* case (T-604/18), the General Court linked Google's abusive self-preferencing to not only harm for rival search services, but also to a broader loss of "plurality in a democratic society",[32] thereby implicitly recognising that informational gatekeeping has externalities beyond price effects. If a dominant LLM were to privilege affiliated publishers or foreground subscription platforms that share revenue with the model owner, the economic harm would mirror that seen in the Android scenario, but at a deeper infrastructural layer. Although add-on gatekeeper obligations under the Digital Markets Act (DMA) may mitigate some risks (see below), technical expertise and real-time auditing capacity are lacking within EU institutions, enabling stack owners to influence emerging standards in their favour. In ordoliberal terms, therefore, *Ordnungspolitik* must extend beyond ex-post antitrust to the proactive design of transparent pipelines, open interfaces, and data access corridors that preserve contestability in the intention economy.[33]

In principle, an LLM that curates product information could also improve market efficiency by lowering search costs and matching consumers to more suitable goods. However, the potential of small efficiency gains cannot excuse a lack of governance that essentially allows a small number of firms to engineer default choice architectures unseen. This is because, as the next section demonstrates, the same persuasive capacities that drive sales can also influence political preferences.

4 Democratic repercussions: persuasion and electoral integrity

LLMs have recently crossed the threshold from being convenient but somewhat limited chatbots to becoming powerful persuaders. In a study, participants who unknowingly debated trivia, forecasting, and moral dilemmas with Claude 3.5 Sonnet, a frontier LLM, were more easily persuaded than peers who were paired with human persuaders.[34] The model's ability to guide users towards both correct and incorrect positions with equal fluency was enabled by its optimised persuasive language. The study authors conclude that "our findings suggest that AI's persuasion capabilities already exceed those of humans that have real-money bonuses tied to performance".[35] Seemingly benevolent uses can flip in an instant: an earlier experiment in which a chatbot reduced belief in conspiracy theories by 20% demonstrates the same mechanics that could be weaponised by an ill-intentioned actor against democratic norms.[36] From an ordoliberal perspective, this is important because persuasion on a large scale blurs the line between messaging in the market for ideas and biased political advocacy, creating a new form of power concentrated in the hands of firms that calibrate response style without democratic oversight. Once citizen preferences have been subtly recalibrated within a private dialogue box, the marketplace of ideas ceases to function as a check on concentrated authority. This violates the very condition of dispersed power that Walter Eucken and Franz Böhm treated as constitutional.

The influence is often invisible because it operates through linguistic framing rather than explicit argument. Preregistered experiments involving 2,582 participants demonstrated that biased autocomplete

[32] Judgment of the General Court (Sixth Chamber, Extended Composition) of 14 September 2022, Google LLC and Alphabet, Inc. v European Commission, Case T-604/18, ECLI:EU:T:2022:541.

[33] Anselm Küsters and Matthias Störring, 'Ein Ordnungsrahmen Für Die »schöne Neue Welt« der Digitalen Medien', *Wirtschaftsdienst* 105, no. 3 (2025): 155–60, https://doi.org/10.2478/wd-2025-0045. See also the policy recommendations below.

[34] Philipp Schoenegger et al., 'Large Language Models Are More Persuasive Than Incentivized Human Persuaders' (arXiv, 2025), https://doi.org/10.48550/ARXIV.2505.09662.

[35] Schoenegger et al.

[36] Thomas H Costello, Gordon Pennycook, and David Gertler Rand, 'Durably Reducing Conspiracy Beliefs through Dialogues with AI', 3 April 2024, https://doi.org/10.31234/osf.io/xcwdn.

suggestions influenced writers' expressed attitudes towards the LLM's stance. Crucially, warnings about potential biases were ineffective.[37] Cross-lingual tests confirm that "that all LLMs inhibit different moral biases to some degree and that they not only differ from the human preferences but also across multiple languages within the models themselves".[38] The line between persuasion and outright disinformation becomes blurred when hostile actors exploit the LLM supply chain through so-called data poisoning.[39] As has been reported, a Moscow-based network currently floods the open web with around 3.6 million pro-Kremlin articles per year in order to contaminate LLM training data with inaccurate information. In testing, 56 out of 450 AI chatbot responses were directly linked to these fabricated stories,[40] showing how easily propaganda can masquerade as neutral text summary once it has passed through the LLM pipeline described in Section 2. From an ordoliberal perspective, the way modern LLMs are built could enable new forms of anti-competitive behaviour or informational cartels, in the sense of efforts to undermines the competitive process by overwhelming rivals with inexpensive falsehoods. Unlike traditional media capture, which targets visible broadcasters, LLM data poisoning works upstream, manipulating the reservoir from which every downstream application draws. Like a price distortion that quietly shifts consumption, stylistic nudges re-anchor economic and political evaluations. As citizens are largely unaware of this shift, the corrective forces of open debate and counter-speech – the traditional means of preventing the misuse of communicative power – never fully engage.

The democratic stakes became already tangible during the recent election season. Ahead of the 2025 German federal election, for instance, *Democracy Reporting International* found that leading chatbots were providing voters with inconsistent, and sometimes incorrect, information about registration deadlines and voting procedures.[41] (Note that this may breach the Digital Services Act's safe-design requirement.) Similarly, academic research indicates that larger frontier models, such as Llama-3-70B, tend to align more closely with left-leaning German political parties.[42] No matter which type of policy one personally prefers: When a single conversational interface shapes both voter understanding of procedural rules and programmatic policy evaluations, the ordoliberal notion of citizen sovereignty will eventually erode. Individuals cannot freely form and express political will if the informational field is systematically distorted. Traditional solutions such as plural media ownership, campaign finance disclosure, and broadcast time parity address symptoms in downstream channels, but leave the new LLM core, where political and ideological preferences are now pre-filtered and reframed, untouched.

In sum, LLMs combine persuasive power and latent ideological framing to pose a strong threat to markets and electoral integrity. The next section therefore presents an ordoliberal theory of harm that unifies these factors, arguing that the combination of economic and political influence within a single model necessitates regulatory concepts that can span both domains simultaneously.

---

[37] Sterling Williams-Ceci et al., 'Bias in AI Autocomplete Suggestions Leads to Attitude Shift on Societal Issues', 15 March 2024, https://doi.org/10.31234/osf.io/mhjn6.

[38] Karina Vida, Fabian Damken, and Anne Lauscher, 'Decoding Multilingual Moral Preferences: Unveiling LLM's Biases Through the Moral Machine Experiment' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2407.15184.

[39] Yao Qiang et al., 'Learning to Poison Large Language Models During Instruction Tuning' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2402.13459.

[40] See: https://www.newsguardrealitycheck.com/p/a-well-funded-moscow-based-global.

[41] Camila Weinmann, Duncan Allen, and Ognjan Denkovski, 'Inconsistent and Unreliable: Chatbots Provide Inaccurate Information on German Elections' (Democracy Reporting International, February 2025), https://democracy-reporting.org/en/office/global/publications/inconsistent-and-unreliable-chatbots-provide-inaccurate-information-on-german-elections#Conclusions.

[42] Rettenberger, Reischl, and Schutera, 'Assessing Political Bias in Large Language Models'.

5 Ordoliberal theory of harm: convergence of economic & political power

Ordoliberalism has always viewed the dispersion of power as a constitutional necessity rather than a mere policy preference.[43] It insists that markets as well as polities can only remain free when no single entity dominates the rules of the game, so that the "invisible hand" can function in both domains. This is what has been termed the "competition democracy nexus" of ordoliberalism,[44] a feature which has wide-ranging implications for the relationship between ordoliberalism, democracy, and populism.[45] However, the analysis above suggests that the invisible hand itself is now driven by machine learning-based language generation and thus no longer unbiased: LLMs sit at the intersection of consumer choice, advertising revenue, and public discourse; consequently, their owners will soon be able to calibrate prices, attention, and narratives. The result is not just market power but a combined ability to influence incentives and beliefs simultaneously. Ordoliberal theory, with its twin commitments to consumer and citizen sovereignty, is uniquely positioned to diagnose this threat because it recognises that economic and political orders rise and fall together.

Competition law expert Viktoria Robertson has recently called for new „[t]heories of harm that specifically take democracy-related concerns into account, be it in merger control or in unilateral conduct", as this "may allow competition authorities to more closely consider instances in which powerful companies enter the political terrain without any democratic legitimisation."[46] In this context, the first layer of harm that must be identified is the self-reinforcing economic feedback loop. Proprietary access to user data and model weights allows frontier LLM developers to improve their models' performance more quickly than newcomers can replicate – and better answers attract more queries. This generates incremental data that strengthens the next model. The leaked report cited above about ChatGPT's "priority placement" for participating publishers as well as experiments with paid product carousels[47] demonstrate how this new type of market power could be monetised and abused.[48] As the model also generates the persuasive text that accompanies these placements, the distinction between advertising and content becomes blurred. In ordoliberal terminology, the platform replaces Hayek's discovery process with a closed optimisation loop whose function is chosen privately.

A second layer, which is qualitatively different, involves shaping consumer preferences. As mentioned earlier, randomised laboratory evidence shows that LLMs are more persuasive than humans, and can guide users towards both correct and incorrect positions with equal ease.[49] Other experiments demonstrate that even subtle stylistic nudges can shift political attitudes without triggering conscious resistance, and that warnings about bias are ineffective.[50] When the same entity that allocates market

---

43 Franz Böhm, 'Economic Ordering as a Problem of Economic Policy and a Problem of the Economic Constitution', in *The Birth of Austerity: German Ordoliberalism and Contemporary Neoliberalism*, ed. Thomas Biebricher and Frieder Vogelmann (London: Rowman & Littlefield Int., 2017), 115–20.

44 Deutscher and Makris, 'Exploring the Ordoliberal Paradigm: The Competition Democracy Nexus'.

45 Malte Dold and Tim Krieger, 'Market Democracy, Rising Populism, and Contemporary Ordoliberalism' (University of Freiburg, Wilfried Guth Endowed Chair for Constitutional Political Economy and Competition Policy, 2024), https://EconPapers.repec.org/RePEc:zbw:wgspdp:281997; Gerhard Wegner, 'Ordoliberalism and Democracy: How the Interwar Period Changed the Agenda of German Liberalism', in *Ordoliberalism and European Economic Policy: Between Realpolitik and Economic Utopia*, ed. Malte Dold and Tim Krieger, Routledge Studies in the History of Economics 225 (New York: Routledge, 2019), 39–57.

46 Viktoria H.S.E. Robertson, 'Protecting Democracy in the Digital Era', 19 January 2025, https://doi.org/10.59704/6162c06c70bb5540.

47 See: https://arstechnica.com/ai/2025/04/chatgpt-goes-shopping-with-new-product-browsing-feature/.

48 Again, for this point, see: Chaudhary and Penn, 'Beware the Intention Economy'.

49 Schoenegger et al., 'Large Language Models Are More Persuasive Than Incentivized Human Persuaders'.

50 Williams-Ceci et al., 'Bias in AI Autocomplete Suggestions Leads to Attitude Shift on Societal Issues'.

opportunities also shapes the cognitive framework through which choices are evaluated, consumer sovereignty collapses. The situation is no better for citizen sovereignty: chatbots that provided incorrect procedural guidance prior to elections[51] exemplify how informational asymmetry might directly impact, or even erode, electoral participation.

Thirdly, conventional antitrust metrics struggle to capture this dual power, as market share in advertising or cloud services tells us little about how effectively a single LLM can influence collective beliefs. In the *Google Android* case (2022) mentioned earlier, the EU General Court found that Google's abusive conduct was harming users' interests in accessing multiple sources of information online. These interests, the Court held, are "not only consistent with competition on the merits, [but] also necessary in order to ensure plurality in a democratic society".[52] In this sense, the LLM theory of harm outlined here fulfils exactly Roberton's call for "democracy-related concerns", extending merger review and unilateral conduct doctrines into the informational realm. In other words, *Ordnungspolitik* must invent new metrics that register invisible leverage – the ability to influence agendas and perspectives, for instance – because price effects no longer reveal the true social cost of concentration. In June 2025, the Dutch Authority for Consumers and Markets showed how to build on this type of reasoning when imposing structural remedies during DPG Media's takeover of RTL Nederland that were aimed at safeguarding media pluralism rather than prices.[53] These remedies included independent editorial foundations, perpetual free-news access, and veto rights on changes to the company's mission. This case demonstrates how competition authorities can address information-related concerns in their merger reviews and hints at how *Ordnungspolitik* could rely on new metrics that register the ability to influence agendas and perspectives. While the ACM's decision illustrates how such societal values could be incorporated into traditional merger assessments, it also raises questions about legal certainty and whether competition authorities have the necessary expertise to conduct such fundamentally non-economic evaluations.

Finally, it is important to emphasise that algorithmic capture by LLMs might structure civic debate before regulators or citizens even realise it. As AI scholar Louise Amoore reminds us: "The processes and arrangements of weights, values, bias, and thresholds in neutral nets are, I think we can safely say, not part of our statutory political domain. And yet, I suggest that they must be presented as questions and political claims in the world."[54] In other words, the LLM design choices are constitutional choices in all but name.[55] To tackle this challenge, the theory of harm developed here spans four interlocking steps: self-reinforcing data monopolies; preference shaping that undermines sovereignty; the inadequacy of price-based metrics; and the portrayal of algorithmic capture as a new form of private constitutional

---

51 Weinmann, Allen, and Denkovski, 'Inconsistent and Unreliable: Chatbots Provide Inaccurate Information on German Elections'.

52 Judgment of the General Court (Sixth Chamber, Extended Composition) of 14 September 2022, Google LLC and Alphabet, Inc. v European Commission, Case T-604/18, ECLI:EU:T:2022:541.

53 In particular, the ACM imposed extensive structural conditions to preserve media pluralism, including: independent foundations with priority shares overseeing RTL Nieuws and NU.nl, with veto powers over editor-in-chief appointments and changes to the company's mission; expanded rights for the Democracy & Media Foundation to prevent the sale or termination of national news brands without approval; mandatory free access to news sites, with prohibitions on paywalls; editorial independence protections through strengthened editorial statutes and separate newsrooms; and a binding charter committing all shareholders to journalistic principles. See: https://www.acm.nl/en/publications/acm-attaches-strict-conditions-acquisition-rtl-nederland-dpg-media.

54 Louise Amoore, *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others* (Durham: Duke University Press, 2020), 81.

55 Küsters and Wörsdörfer, 'Exploring Laws of Robotics: A Synthesis of Constitutional AI and Constitutional Economics'.

power. This diagnosis paves the way to the policy sections of this paper, which consider whether Europe's existing policy tools can neutralise this combined power before it solidifies.

6 Europe's existing arsenal: competition law, DMA, DSA, and AI Act

European competition law already contains some ordoliberal notions that may address the informational gatekeeping power of LLMs.[56] While Art. 102 TFEU remains the most important rule for tackling exploitative or exclusionary conduct of dominant firms, the preceding analysis cautions that market-share metrics calibrated for price competition (which are used to identify a "dominant" undertaking in a "relevant" market) undervalue the harm of intention capture. Enforcement of Art. 102 TFEU and effective merger control thus need to consider whether combining datasets, model weights, or guardrail stacks in a dominant LLM would create an informational "essential facility".[57] In the case of *Android Auto*, the CJEU ruled that a dominant platform's refusal to interoperate with a rival app can be abusive, even when the platform is not strictly "indispensable" but merely makes the rival's service more attractive to users.[58] Commentators have dubbed the judgment "the end of the essential-facility doctrine as we know it",[59] as the Court replaced *Bronner*'s rigid indispensability test with a much looser convenience threshold, emphasising the platform's role as an access point for downstream innovation. This feature of the case might be significant for LLMs: if a large-scale model controls the interface through which users search, create, and verify information, it may also be considered an informational essential facility. In this respect, the existing arsenal might be conceptually fit for purpose, but only if the authorities are willing to expand the notion of consumer welfare beyond price to encompass attention and discursive plurality.

While traditional competition law enforcement is reactive, the Digital Markets Act (DMA) provides a proactive complement.[60] By designating "gatekeepers" in advance, the DMA requires dominant intermediaries to refrain from favouring themselves, to maintain data portability and, crucially for LLMs, to provide business users with fair, reasonable, and non-discriminatory access to ranking parameters. However, effective oversight requires significant investment in additional resources and technical expertise within EU institutions. According to observers, DMA stakeholder workshops on technical implementation are already dominated by the very Big Tech firms that will be regulated,[61] raising the ordoliberal spectre of regulatory capture. One idea would therefore be to amend the DMA to create explicit data-access corridors for independent researchers – an idea that mirrors transparency mandates in the Digital Services Act (DSA), but which is currently absent from the DMA. Moreover, regulators must design participatory processes to prevent tech giants from gaining implicit control over regulation during stakeholder sessions. An early DMA amendment should clarify that gatekeeper status can arise at the model layer, where intentions are captured, rather than only at the traditional interface

---

[56] Anselm Küsters and Isabel Oakes, 'Taming Giants: How Ordoliberal Competition Theory Can Address Power in the Digital Age', *Schmollers Jahrbuch – Journal of Contextual Economics* 141, no. 3 (2021): 149–88.

[57] Rok Dacar, 'The Essential Facilities Doctrine, Intellectual Property Rights, and Access to Big Data', *IIC - International Review of Intellectual Property and Competition Law* 54, no. 10 (November 2023): 1487–1507, https://doi.org/10.1007/s40319-023-01396-7.

[58] Judgment of the Court (Grand Chamber) of 25 February 2025, Alphabet Inc. and Others v Autorità Garante della Concorrenza e del Mercato (AGCM), Case C-233/23, ECLI:EU:C:2025:110.

[59] See: https://competitionlawblog.kluwercompetitionlaw.com/2025/03/13/android-auto-the-end-of-the-essential-facility-doctrine-as-we-know-it/.

[60] Moreno Belloso, 'The EU Digital Markets Act (DMA): A Summary', *SSRN Electronic Journal*, 13 May 2022, 1–4; Jacques Crémer et al., 'Fairness and Contestability in the Digital Markets Act', *SSRN Electronic Journal*, 2021, https://doi.org/10.2139/ssrn.3923599.

[61] See: https://corporateeurope.org/en/2024/10/uncovering-big-techs-hidden-network.

or operating system layer. This would close a loophole that otherwise allows LLM providers to give preferential treatment to affiliated content without formal designation.

Meanwhile, the DSA addresses a different aspect of the problem: the systemic risks that arise from disseminating illegal content and disinformation.[62] The duties imposed on "very large online platforms" with regard to risk assessment and safe design proved to be prescient in the run-up to Germany's 2025 federal election, when chatbots provided contradictory or simply false information about voter registration deadlines. This behaviour was identified by *Democracy Reporting International* as a potential violation of the DSA.[63] Unlike the DMA, the DSA incorporates a right of access to research data, which could, in principle, enable auditors to trace how prompts, retrieval layers, and safety filters interact. However, the Act focuses on the legality of content rather than the more subtle issue of algorithmic framing that influences preferences without crossing the threshold of illegality. Without supplementary guidance, LLM-driven platforms may demonstrate formal compliance with the DSA (if at all) while leaving competitive and democratic distortions untouched.

The EU AI Act is the latest addition to Europe's regulatory framework for addressing the power of LLMs, introducing provisions that complement competition law, the DMA, and the DSA.[64] The Act establishes a comprehensive legal framework for AI regulation, paying particular attention to general-purpose AI (GPAI) models, including LLMs.[65] Essentially, the AI Act makes a key distinction between AI models and systems, classifying GPAI models based on systemic risk rather than traditional market-based assessments. Under the law, a GPAI model is classified as having "systemic risk" if it meets specific computational thresholds – notably, models requiring more than 10^{25} floating-point operations for training are presumed to have high-impact capabilities. While this computational threshold offers a technically grounded approach to identifying potentially problematic LLMs, it has been shown to be too narrow.[66] Modern LLMs can capture user intentions and influence information access patterns, regardless of their initial "compute". The AI Act's transparency requirements for GPAI models directly address the informational opacity concerns raised in the ordoliberal analysis. Providers of general-purpose AI models must maintain technical documentation, provide summaries of training content, and comply with copyright laws. Additional obligations for models with systemic risk include conducting model evaluations, assessing and mitigating systemic risks, tracking serious incidents, and ensuring cybersecurity.

However, the AI Act faces institutional capacity challenges similar to those identified with the DMA. The development of a Code of Practice for GPAI models involves stakeholder consultations which are dominated by the very tech giants that are being regulated.[67] Technical complexity continues to give incumbents structural advantages in shaping implementation standards. Furthermore, there are still gaps in addressing the subtle forms of intention capture identified in the ordoliberal theory of harm discussed above. Although the AI Act requires risk assessment and mitigation, it does not explicitly address how LLMs might influence user preferences by presenting information or ordering search

---

62 Aina Turillazzi et al., 'The Digital Services Act: An Analysis of Its Ethical, Legal, and Social Implications', *Law, Innovation and Technology* 15, no. 1 (2 January 2023): 83–106, https://doi.org/10.1080/17579961.2023.2184136.

63 Weinmann, Allen, and Denkovski, 'Inconsistent and Unreliable: Chatbots Provide Inaccurate Information on German Elections'.

64 In general, see: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai.

65 Nicola Fabiano, 'AI Act and Large Language Models (LLMs): When Critical Issues and Privacy Impact Require Human and Ethical Oversight' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2404.00600.

66 Sara Hooker, 'On the Limitations of Compute Thresholds as a Governance Strategy' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2407.05694.

67 See: https://corporateeurope.org/en/2025/01/setting-rules-their-own-game-how-big-tech-shaping-ai-standards.

results. Clear guidance on how systemic risk assessments should evaluate competitive distortions would be required for this. Overall, the effectiveness of the new EU AI Act in addressing the informational gatekeeping power of LLMs will depend on the quality of its implementation and the development of sufficient institutional capacity to match the technical sophistication of the entities being regulated. *Ordnungspolitik* must therefore treat institutional capacity as part of the competitive landscape.

Recent market developments suggest that the EU's current toolkit, especially with respect to antitrust, may be more effective than previously thought. Just two years ago, OpenAI's position seemed unassailable. However, the "DeepSeek shock"[68] and the rapid advancement of more open alternatives, such as Llama 3, Mistral and associated projects, suggest that LLM technology is becoming commoditised. As baseline inference capabilities become interchangeable inputs, market power is increasingly exercised once more through classic competitive levers such as vertical integration, bundling, and switching costs, rather than through exclusive control of model weights. This implies that conventional doctrines on tying, foreclosure, and leveraging (via, e.g., Art. 102 TFEU, the DMA's ban on self-preferencing, and merger control focused on conglomerate effects) can more easily be used against practices such as coupling Microsoft's cloud and productivity suites with proprietary models. However, even commodified models can function as informational chokepoints within closed ecosystems, so further, more unconventional measures are needed to complement competition law enforcement. Unless regulators can examine the inner workings of LLMs and LLM-driven platforms in real time, formal rules will always lag behind the substantive concentration of economic and political power mapped out in this paper.

7 Policy blueprint: utility engineering and Constitutional AI

Recognising that LLMs already have implicit, structured worldviews that influence the "invisible hand" in both markets and politics, ordoliberals are confronted with a governance challenge: how to guide these worldviews without replacing one private arbiter with a public censor. Given the gaps left by traditional competition policy, policymakers should treat LLM value alignment as a constitutional issue – one that belongs in the same normative space as competition rules and fundamental rights, rather than in the technical background of model training. The goal is neutrality in the classical ordoliberal sense: an institutional framework that channels private initiative, both in the economic and in the political domain, without embedding hidden coercion. As the LLM pipeline combines economic and political power at an early stage, neutrality must also be established at this stage, i.e. inside the model's latent value function, before downstream solutions such as abuse control via Art. 102 TFEU or content moderation via the DSA can be effective. But how can this be achieved in practice?

Computer scientists have started developing an algorithmic tool for this initial task under the name of "Utility Engineering". In a seminal xAI study, researchers first elicited an LLM's revealed preferences by presenting it with thousands of moral and policy dilemmas.[69] They then re-weighted these preferences to align with a democratically selected target distribution. This proof of concept used US census data and demonstrated that, after a single fine-tuning process, the model's responses became measurably closer to those of the median Trump voter than the median Biden voter. Without normatively judging this shift, the test certainly shows both the malleability of emergent values and the feasibility of precise adjustment. From an ordoliberal perspective, there might be a helpful analogy with price

[68] See: https://cyber.fsi.stanford.edu/publication/taking-stock-deepseek-shock.
[69] Mantas Mazeika et al., 'Utility Engineering: Analyzing and Controlling Emergent Value Systems in AIs', 2025.

theory: just as a well-designed tax can correct an externality without dictating individual choices, Utility Engineering aims to correct ideological biases while leaving the competitive discovery process intact. Since larger LLMs exhibit more internally consistent value systems, early intervention is key: once a frontier model has gained a dominant market share, minor adjustments may no longer be enough to restore neutrality. In other words, the window for effective *Ordnungspolitik* closes quickly: any delay risks allowing emergent private constitutions to solidify into "self-made law" *(selbstgeschaffenes Recht der Wirtschaft*), as feared by Großmann-Doerth.[70]

Another line of research experiments with constitutional AI and, more specifically, ordoliberal system prompts. System prompts are instructions provided to LLMs that define their behaviour, role, and operational parameters before they engage with user queries. Together with Wörsdörfer, this paper's author compared the effect of a short Freiburg School-style system prompt on the content-moderated Llama-2 model with that of its uncensored counterpart.[71] The guarded version produced only marginal shifts in the content-moderated model, suggesting that corporate safety layers already saturate the value space. In contrast, the uncensored version yielded visibly more rule-consistent answers across 400 ethical-dilemma iterations. Accordingly, system prompts can act as lightweight constitutional norms, but their efficacy depends on transparent access to the base model – an access that incumbent providers are increasingly restricting.

Moreover, these types of technical alignment alone cannot confer legitimacy, which is why public input should play an important role, too.[72] For instance, deliberative bodies that sample the population by lot and deliberate under expert guidance might help to bring about "constitutional AI". According to modern ordoliberal scholars, "mini-publics" can translate dispersed preferences into actionable norms without succumbing to lobby capture.[73] Integrating these assemblies into the LLM alignment loop would enable the public to author the utility functions and system prompts that guide frontier models. In practice, an assembly could select the weight vector for a Utility-Engineered model or approve the text of an ordoliberal system prompt. This would replace corporate policy with a transparent, democratically grounded mandate. In ordoliberal terms, not only must the invisible hand be free of bias, it must also be seen to act under rules authored by the citizenry it ultimately serves.

However, long-term resilience against the cumulative power of a small number of LLMs cannot rely on supervisory vigilance alone; it requires an innovation ecosystem capable of contesting incumbents on their merits. In this regard, targeted industrial policy investments in a "EuroStack" could help reduce Europe's digital dependency: Currently, over 80% of digital technologies and infrastructures are imported, and 70% of foundational AI models originate from the US.[74] Backed by a cross-party coalition in the European Parliament, the EuroStack framework aims to establish Europe as a digitally sovereign power by developing its own technological capabilities. At its core, it encompasses interconnected layers of advanced technologies, including semiconductors, networks, satellites, software, cloud computing, quantum technology, the Internet of Things and, critically, data and AI systems. These layers are connected through common services for all EU citizens, such as federated data spaces, an EU-wide

[70] Großmann-Doerth, 'Selbstgeschaffenes Recht der Wirtschaft und staatliches Recht'.

[71] Küsters and Wörsdörfer, 'Exploring Laws of Robotics: A Synthesis of Constitutional AI and Constitutional Economics'.

[72] Saffron Huang et al., 'Collective Constitutional AI: Aligning a Language Model with Public Input', in *The 2024 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '24: The 2024 ACM Conference on Fairness, Accountability, and Transparency, Rio de Janeiro Brazil: ACM, 2024), 1395–1417, https://doi.org/10.1145/3630106.3658979.

[73] Dold and Krieger, 'Market Democracy, Rising Populism, and Contemporary Ordoliberalism'.

[74] Francesca Bria, Paul Timmers, and Fausto Gernone, *EuroStack – A European Alternative for Digital Sovereignty* (Gütersloh: Bertelsmann Stiftung, 2025).

digital ID, and the digital euro. Relying on open-source innovation, the EuroStack initiative would directly challenging the concentrated power of proprietary LLM providers through various means.[75] Firstly, open source enables innovation to occur much more rapidly than in closed, proprietary systems. Secondly, open source fundamentally democratises access to cutting-edge AI technologies by eliminating the barriers that typically limit access to leading-edge LLMs. Unlike closed AI models, where data and training methods are closely guarded secrets, open-source models allow anyone to examine the underlying systems, reducing potential hazards and biases. Similarly, open Small Language Models (SLMs) require significantly less computational power and memory, enabling broader participation in AI development and preventing the concentration of capabilities among well-funded tech giants.[76] Moreover, SLMs can be fine-tuned for specific tasks or domains, often resulting in superior performance in targeted applications compared to general-purpose large models. In sum, the EuroStack initiative could thus focus the continent's next investments on building diverse data centres optimised for inferencing domain-specific SLMs, creating a more distributed and resilient AI ecosystem. This strategic approach serves two ordoliberal purposes: it creates competitive pressure on leading LLM providers while simultaneously countering the biases that currently characterises AI development, such as a strong focus on English-language training data. By fostering an innovation ecosystem capable of competing with incumbents on their merits, EuroStack offers a pathway to domestic LLMs that are rooted in European values.

8 Conclusion: reasserting the invisible hand in the age of LLMs

This study began with the simple observation that conversational AI is increasingly replacing Google searches, thereby becoming an infrastructural layer which already mediates a quarter of all search behaviour in the US, with Europe not far behind. By taking over the moment when attention turns into intention, large language models combine three forms of leverage – market dominance, informational asymmetry, and norm-setting capacity – into a single "algorithmic hand". To encapsulate the ordoliberal concern in one sentence: once machine-generated language drives Adam Smith's invisible hand, it ceases to be impartial. The analysis traced how this platform shift threatens to distort competition through priority placement and paid ranking schemes, reshapes political discourse by influencing attitudes and providing voters with unreliable advice, and enables hostile actors to contaminate the data sources that AI chatbots rely on. The cumulative effect is a new form of private constitutional power that conventional antitrust metrics focused on price may struggle to register.

However, ordoliberalism's constitutional vision might offer an alternative by combining robust competition with democratic self-rule. In particular, the policy blueprint sketched in the second half of this paper proposes a two-pronged approach. Firstly, it seeks to strengthen existing "classic" competition policy instruments, such as Art. 102 TFEU, merger control, the DMA, the DSA, and the AI Act by providing enforcers with the technical capability to inspect model pipelines and by closing loopholes that allow gatekeeping to shift to the model layer. Secondly, new institutions are to be built that tackle bias upstream, including the new AI Office and targeted industrial-policy funding for open-source and small language models to maintain contestable markets. There is also great potential for democratically mandated LLM alignment tools, such as "Utility Engineering" and ordoliberal system prompts.

[75] For more details on these suggestions, see: Anselm Küsters, 'Future-Proofing the EU: Ordoliberal Governance and Algorithmic Regulation', *Constitutional Political Economy*, 1 May 2025, https://doi.org/10.1007/s10602-025-09467-2.

[76] Fali Wang et al., 'A Comprehensive Survey of Small Language Models in the Era of Large Language Models: Techniques, Enhancements, Applications, Collaboration with LLMs, and Trustworthiness' (arXiv, 2024), https://doi.org/10.48550/ARXIV.2411.03350.

Together, these measures aim to restore the invisible hand and establish a governance framework in which LLMs serve as a foundation for freedom rather than a source of hidden coercion.

However, before turning to implementation, policymakers require an explicit framework for distinguishing between *necessary* and *sufficient* conditions in digital regulation, particularly since EU digital policymakers are increasingly confronted with trade-offs or even trilemmas.[77] Some of the interventions proposed above, such as effective enforcement of existing competition law rules, and establishing democratic legitimacy for future LLM value function regulation via citizens' assemblies, are necessary prerequisites; without them, any broader governance architecture would eventually collapse. Others, such as direct model-inspection powers or subsidised computing power for open-source projects in Europe, may not be strictly necessary, but they could be sufficient, alongside the core safeguards, to ensure contestability and accountability. Mapping measures in the rapidly growing literature on LLM alignment against this logic would help clarify priorities in a resource-constrained environment. In particular, regulators should first guarantee the conditions that must be met on every plausible path towards fair and transparent AI markets. Only then should they introduce instruments that accelerate or stabilise the desired equilibrium or privilege domestic industrial participants.

The coming months and years will test Europe's capacity for rapid, smart governance. Regulators will need engineers at their side. Both will also require citizens who are willing to inform or even author the rules that will, in turn, shape their informational horizons. This is a collective constitutional moment: it belongs as much to parliaments and competition agencies as to open-source communities and everyday users, whose conversations now inform the next LLM training run. If Europe rises to this challenge, the algorithmic hand can be made transparent, contestable, and accountable to the public it serves. The alternative is an unspoken constitutional shift in which private code sets public horizons. The choice remains open – but not indefinitely.

[77] For this point, see: Küsters and Sottilotta, „Trade-Offs and Risks in EU Digital Policy, cepInput 2/2025 (01.14.2025), https://www.cep.eu/eu-topics/details/trade-offs-and-risks-in-eu-digital-policy.html.

Centrum für Europäische Politik
FREIBURG | BERLIN

Author:

Dr. Anselm Küsters, LL.M.
Head of Department Digitalisation and New Technologies, cep Berlin
Interim Professor for Digital Humanities, University of Stuttgart
kuesters@cep.eu

Centrum für Europäische Politik FREIBURG | BERLIN
Kaiser-Joseph-Straße 266 | D-79098 Freiburg
Schiffbauerdamm 40 Räume 4205/06 | D-10117 Berlin
Tel. + 49 761 38693-0