

# Is Generative AI the Algorithmic Consumer We Are Waiting for?

Michal Gal and Amit Zac\*

Network Law Review (March 2024)

## Abstract

Most studies on the competitive effects of Generative AI focus on the supply side. Interest in consumers is generally restricted to their roles as users of the technology, as well as indirect trainers of Generative AI models through their prompt engineering. In this short article, we focus instead on the potential effects of Generative AI on competition that arise from an active use of Generative AI on the demand side, by consumers seeking goods and services. In particular, we explore the possibility that Generative AI Large Language Models (LLMs) can act as truncated algorithmic consumers, assisting consumers in deciding which products and services to purchase, thereby potentially reducing consumers' information costs and increasing competition. We explore how LLMs' unique characteristics – mainly their conversational use and the provision of an authoritative single answer, as well as spillover trust effects from their other uses – might motivate consumers to use them to search for products and services. We then analyze some of the limitations and competition concerns that might result from the use of Generative AI by consumers. In particular, we show how LLMs' modus operandi - trained to seek the most plausible next word - lead to outcome homogenization and increase entry barriers for new competitors in product markets. We also explore the potential of manipulation and gaming of LLM models. As elaborated, a combination of an LLM model with a dataset on consumers' digital profiles might potentially create a strong nudging mechanism, recreating consumer choice architecture to optimize commercial goals and exploiting consumer's behavioral biases in novel ways not envisioned.

## I. Introduction

Generative AI<sup>1</sup> is the technology *de jur*. This is not surprising, given that this general-use technology is already disrupting markets across industries and sectors,<sup>2</sup> invoking debates about the future of some businesses, jobs and education. Much has already been written on its effects on competition.<sup>3</sup> Such studies mostly emphasize the high barriers that characterize foundation models of Generative AI. These include, first and foremost, access to vast datasets needed to train the AI (to illustrate, ChatGPT

---

\*Michal Gal is Professor of law, University of Haifa and former president of ASCOLA (2016-2023). [mgalresearch@gmail.com](mailto:mgalresearch@gmail.com). Amit Zac is Assistant Professor of law, University of Amsterdam, and Research Fellow ETH Zurich [amit.zac@gess.ethz.ch](mailto:amit.zac@gess.ethz.ch). Many thanks to Thibault Schrepel and Volker Stocker for superb comments. This work was supported by ISF grant 2737/20. Any mistakes or omissions remain the authors'.

<sup>1</sup>Generative artificial intelligence (AI) describes algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images, text, simulations, and videos. McKinsey and Co., What is generative AI? January 19, 2023, <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai>. See also Competition and Markets Authority, *AI Foundation Models: Initial Report* (Sept. 18, 2023), <https://www.gov.uk/government/publications/ai-foundation-models-initial-report>.

<sup>2</sup> Daron Acemoglu, Pascual Restrepo, The wrong kind of AI? Artificial intelligence and the future of labour demand, *Cambridge Journal of Regions, Economy and Society*, Volume 13, Issue 1, March 2020, Pages 25–35, <https://doi.org/10.1093/cjres/rsz022>; Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. "Do we want less automation?." *Science* 381, no. 6654 (2023): 155-158.

<sup>3</sup> See references below.

(OpenAI) was trained on 300 Billion words and Bard (Google) was trained on 1.56 trillion words).<sup>4</sup> Entry barriers also include, inter alia, the need to update the data to provide up-to-date answers; the need to store vast amounts of data; and copyright disputes over the required data and possibly also over the output produced by the models. An interesting study performed by Gordon Phillips on all major Chinese firms, indicated that in markets in which firms reported that they use AI, the number of firms declined over time, whereas the use of a technology such as Cloud Computing increased it.<sup>5</sup> Leading researchers have already offered insights on how lessons learned with regard to supply-side entry barriers in platform and big-data based markets, can be carried over to Generative AI foundation models.<sup>6</sup>

Other studies have focused on the need to recognize and restrain some of the socio-economic harms created by Generative AI. These include, inter alia, the fact that such models can reproduce biases ingrained in the data, unless trained otherwise (e.g., autocomplete of the sentence "I talked with my Dr yesterday. I told him" (rather than her)).<sup>7</sup>

What is common to all such studies of the effects of Generative AI is their focus mainly on the supply side of Generative AI markets. Interest in consumers is generally restricted to their roles as users of the technology, as well as indirect trainers of Generative AI models through their prompt engineering. In this short article, we wish to focus on some of the potential effects of Generative AI on competition that come from an active use of Generative AI on the demand side of consumer goods and services. In particular, we explore the possibility that Large Language Models of Generative AI (LLMs) can act as a truncated algorithmic consumer (chapter II). As will be shown, LLMs differs both qualitatively and quantitatively from other search tools based on big data, potentially affecting demand and market dynamics. Finally, we analyze some of the limitations and competition concerns concerning Generative AI (chapter III). In particular, we offer insights on how LLMs might affect consumer behavior as well as entry barriers for new competitors. We also explore the potential of manipulation

---

<sup>4</sup> Bard, in answer to a query "on how many words were Bard and ChatGPT trained?", double checked with Google Search (February 11, 2024). Some LLM are open source. (e.g., hugging face - [https://huggingface.co/spaces/HuggingFaceH4/open\\_llm\\_leaderboard](https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard)). Furthermore, LLM scale economies start to diminish after a certain significant volume. Petland and Schrepel, supra, p. 7. Furthermore, given that several large players have already been able to train their models on significant scales, future comparative advantages may come from other sources. Will Knight, "OpenAI's CEO Says the Age of Giant AI Models Is Already Over," Wired, April 17, 2023, <https://perma.cc/VZ8E-MGQE>.

<sup>5</sup> Yao Lu, Gordon M. Phillips, and Jia Yang, *The Impact of Cloud Computing and AI on Industry Dynamics and Competition* (June 15, 2023), [https://faculty.tuck.dartmouth.edu/images/uploads/faculty/gordon-phillips/China\\_Cloud\\_Computing.pdf](https://faculty.tuck.dartmouth.edu/images/uploads/faculty/gordon-phillips/China_Cloud_Computing.pdf). See also Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, *Generative AI at work*. No. w31161. National Bureau of Economic Research (2023).

<sup>6</sup> See, e.g., Peter Georg Picht, ChatGPT, *Microsoft and Competition Law – Nemesis or Fresh Chance for Digital Markets Enforcement?* (July 18, 2023). <https://ssrn.com/abstract=4514311>; Thibault Schrepel and Alex Pentland, *Competition between AI Foundation Models: Dynamics and Policy Recommendations* (June 28, 2023), <https://ssrn.com/abstract=4493900>; Christophe Carugati, *Competition in Generative AI Foundation Models* (June 18, 2023). Working Paper 14/2023, Bruegel, <https://ssrn.com/abstract=4553787>; FTC, *Generative AI raises Competition Concerns*, June 29, 2023, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/06/generative-ai-raises-competition-concerns> (focusing on access to resources necessary to create GenAI, and exploring the potential effects of open source models); Chen, Wei and Wang, Xiaoyu and Xie, Karen and Xu, Fasheng, *The Economics of AI Foundation Models: Transparency, Competition, and Governance* (February 14, 2024), <https://ssrn.com/abstract=4727903>;

<sup>7</sup> See, e.g., Emilio Ferrara, *Should ChatGPT be Biased? Challenges and Risks of Bias in Large Language Models* (October 26, 2023), <https://ssrn.com/abstract=4614228>.

and gaming of LLM models. As elaborated, a combination of an LLM model with a dataset on consumers' digital profiles might potentially create a strong nudging mechanism, changing a simple 'Dark Pattern' into a personalized, trustworthy, enhanced nudge, thereby recreating consumer choice architecture to optimize commercial goals and exploiting consumer's behavioral biases in novel ways not envisioned.<sup>8</sup>

## I. **Generative AI as Truncated Algorithmic Consumers?**

### A. *The concept of Algorithmic Consumers*

The concept of Algorithmic Consumers was first introduced by Gal and Elkin-Koren in 2017, to describe algorithmic-based digital agents that can assist consumers in making better consumption decisions, and potentially execute entire transactions on their behalf.<sup>9</sup> This concept is based on the fundamental idea that consumers can limit the harm inflicted on them through the use of algorithms by suppliers (such as high personalized prices or coordinated prices), by using countervailing algorithms. Such countervailing algorithms might be written by consumers for their own use, or supplied by external firms.<sup>10</sup> Algorithmic consumers have two main characteristics. First, the automation of some important aspects of consumption. In the extreme version, algorithmic consumers can perform all parts of the consumption act: from using data to predict consumers' preferences, choosing the products or services to purchase, negotiating and executing the transaction, and even form coalitions of buyers to secure optimal terms and conditions.<sup>11</sup> In this extreme version, the consumer's human decision-making is completely bypassed. While such algorithms are already in use in some markets,<sup>12</sup> most algorithmic consumers assist consumers only in some parts of the consumption act, leaving the ultimate decision to the consumer (we call them "truncated algorithmic consumers"). The second condition is that the algorithmic consumer act in a fashion that furthers the consumer's welfare, rather than the supplier's or the intermediary's welfare.

As Gal and Elkin-Koren show, algorithmic consumers have the potential to significantly change market dynamics.<sup>13</sup> By outsourcing some consumption tasks to algorithms, humans can reduce the time and (decisional) energy necessary to make a transaction. But more importantly, algorithmic consumers can significantly reduce search and transaction costs, help consumers overcome biases, enable more sophisticated choices, and create or strengthen buyer power. Furthermore, algorithmic consumers can act as buying groups, increasing buying power and reducing the ability of suppliers to engage in price discrimination where the algorithmic consumer acts as a buffer that hides information regarding

---

<sup>8</sup> See, e.g., Avishalom Tor, *Illustrating a behaviorally informed approach to antitrust law: The case of predatory pricing*, ANTITRUST 18 (2003); Maurice E. Stucke, *How can competition agencies use behavioral economics?* THE ANTITRUST BULLETIN 59 (2014).

<sup>9</sup> Michal Gal and Niva Elkin-Koren, *Algorithmic Consumers*, 30(2) HARVARD JOURNAL OF LAW AND TECHNOLOGY 2 (2017).

<sup>10</sup> See, e.g., CHRISTOPHER STEINER, *AUTOMATE THIS: HOW ALGORITHMS CAME TO RULE OUR WORLD* (2012); Theo Kanter, TEDx Talks, *Ambient Intelligence*, YOUTUBE at 15:13 (Feb. 3, 2016), <https://www.youtube.com/watch?v=1Ubj2kIiKMw> [<https://perma.cc/9VAU-P2Z2>]; Don Peppers, *The Consumer of the Future Will Be an Algorithm*, LINKEDIN (July 8, 2013), <https://www.linkedin.com/pulse/20130708113252-17102372-the-consumer-of-the-future-will-be-an-algorithm>.

<sup>11</sup> See, e.g., Minghua He, Nicholas R. Jennings & Ho-Fong Leung, *On Agent-Mediated Electronic Commerce*, 15 IEEE TRANSACTIONS ON KNOWLEDGE & DATA ENGINEERING 985(2003).

<sup>12</sup> Gal and Elkin-Koren, *supra*.

<sup>13</sup> *Ibid*.

each purchaser from the supplier. Such effects may have profound impacts on suppliers' marketing strategies, trade terms, and product offers.

### *B. How Generative AI can act as a truncated algorithmic consumer*

To understand how Generative AI can potentially act as a truncated algorithmic consumer, we first briefly review the characteristics of Generative AI models that are relevant to such a task. As will be shown, Generative AI has the potential to help consumers to more easily and swiftly compare price and quality, predict price and market trends, make expedient decisions under uncertain conditions, and make better-informed choices.

The most relevant Generative AI tools, such as ChatGPT, Bard and Claude, are LLM Generative AI models.<sup>14</sup> Such models, first introduced to the general public in November 2022, are agile tools that are trained on very large language datasets to predict the most reasonable continuation of the next token (which can even be parts of a word), generating a sequence that can create whole sentences, paragraphs and essays. What is most interesting, is that such tools are based on patterns or sequences in language, rather than a theory of language. As such, the algorithm is not trained to complete an answer that serves the interests of a certain entity, nor to provide a truthful answer, but rather to provide the word that most predictably follows the previous one, given the textual data it was trained on, and the hyper parameters set by the coder, like "temperature" which sets the levels of random or less probable continuation. The core idea of a neural network, a core technology that underlies LLMs, is to create a flexible 'computing fabric', out of many simple components. The fabric is incrementally modified to learn from examples.<sup>15</sup>

The success of LLMs in performing such a task surprised even their creators. In 2023 it was estimated that such models already have an IQ equivalent to the highest percentile of the population.<sup>16</sup> GPT-4 passes the American Bar Exam, or typical Law School tests.<sup>17</sup> And it was discovered that LLMs can engage in numerous tasks that were not expected to be performed at such a high level, sometimes even better than other state-of-the-art technologies.

One of the functions current chat-like LLM models serve is rather similar to existing search engines: both can assist humans in searching for answers to different queries. Yet they differ in important aspects. First, common search models such as Google Search or Yahoo provide links to webpages that the algorithm deems most relevant to the query, requiring the user to search further by opening the links and probing for the answer within them. While such search engines often also highlight some

---

<sup>14</sup> Some LLMs are now multimodal, include other types of Generative AI as well (e.g., ChatGPT Plus includes Dall-E and Sora).

<sup>15</sup> S. Wolfram. *What is ChatGPT doing and why does it work?* Technical report, Wolfram Mathematica, 2023.

<sup>16</sup> Eka Roivainen, *I Gave CHatGPT an IQ Test. Here's What I Discovered*, SCIENTIFIC AMERICAN (March 28, 2023), <https://www.scientificamerican.com/article/i-gave-chatgpt-an-iq-test-heres-what-i-discovered/>. IQ tests measure crystallized intelligence - the ability to deduce secondary relational abstractions by applying previously learned primary relational abstraction, rather than fluid intelligence, which is the ability to solve novel reasoning puzzles. Wikipedia, Fluid and Crystallized Intelligence (accessed Feb. 20, 2024), [https://en.wikipedia.org/wiki/Fluid\\_and\\_crystallized\\_intelligence](https://en.wikipedia.org/wiki/Fluid_and_crystallized_intelligence). Computer scientists are seeking ways to strengthen LLM's logic.

<sup>17</sup> Katz, Daniel Martin, Michael James Bommarito, Shang Gao, and Pablo Arredondo. "Gpt-4 passes the bar exam." *Available at SSRN 4389233* (2023).; Choi, Jonathan H., Kristin E. Hickman, Amy B. Monahan, and Daniel Schwarcz. "ChatGPT goes to law school." *J. Legal Educ.* 71 (2021): 387.

relevant parts in one of the first links, the user might still be required to read the source, to ensure the context is the relevant one, or to seek other parts of the relevant information. LLMs, on the other hand, provide full answers, without any need to delve deeper, unless the user wishes to ask a follow-up question or the user cannot rely on the verity of the answer, a point we return to below. Links to sources can be provided upon a specific prompt of the user to the LLM. The difference between the two can be likened to a "librarian" versus a "stochastic parrot". Second, LLM models change the nature of human-computer interaction, basing it on a conversational use (closer to the smartphones' personal assistants), which makes it easier and more user-friendly for a lay person (that is indicated, for example, in the fact that most users start with "please..." when asking the LLM for an answer). They mimic human interactions,<sup>18</sup> reflecting their advantage in creating a better AI-human symbiosis to non-specialist users. They have an effect which can be likened to moving away from computer punch-cards that were used to display computer-generated data, to computers with graphic interfaces that made it much clearer to lay persons to use and understand. The new interface of chat-like LLMs can impact consumer trust in the AI recommendations,<sup>19</sup> in ways non-interactive search engines will struggle to follow. We return to the limitations of this feature later.

The third main difference is that LLMs can balance a multitude of factors as required by consumers in their queries (such as budget, brand preferences, features), tailoring the answer to the specific requirements of the user - a task which is more difficult for traditional search algorithms. Fourth, LLMs, by definition, have an endless path for learning. Not in the same indirect way search algorithm do (e.g., trying to estimate how content a consumer is by his activity following the search). One feature that made LLM technology so successful in comparison to other machine learning models, is that it can simply learn from natural language text.<sup>20</sup> A feedback is embedded in each new natural language text. Moreover, users can give direct feedback on the generated answer (ChatGPT for example asks, 'Is this conversation helpful so far?' with thumb up or down as a quick direct feedback). As users get better in generating prompts for answers, and the LLMs can learn to generate better answers, becoming even more useful for them.<sup>21</sup> Finally, despite the fact that LLMs are language models, they are capable of conducting data analysis tasks, coding tasks and specifically in our context analyze pricing trends.<sup>22</sup> These capabilities, which will grow in scope and quality will outperform search

---

<sup>18</sup> Cai, Zhenguang G., David A. Haslett, Xufeng Duan, Shuqi Wang, and Martin J. Pickering. "Does ChatGPT resemble humans in language use?" arXiv preprint arXiv:2303.08014 (2023).

<sup>19</sup> Ye, Yang, Hengxu You, and Jing Du. "Improved trust in human-robot collaboration with ChatGPT." *IEEE Access* (2023) found that: "incorporating ChatGPT in robots significantly increased trust in human-robot collaboration, which can be attributed to the robot's ability to communicate more effectively with humans. Furthermore, ChatGPT's ability to understand the nuances of human language and respond appropriately helps to build a more natural and intuitive human-robot interaction'. We believe such findings will generalize well to other Human-AI interactions.

<sup>20</sup> Multimodal input might make it even better.

<sup>21</sup> Much depends on the quality of the feedback LLMs receive relative to search engines. LLMs can factor in what users seek and whether they seek further information, while Search engines can analyze which result on the list users prefer. Thibault Schrepel, *Competition Is One Prompt Away*, Network Law Review, Fall 2023.

<sup>22</sup> Jaimovitch-López, Gonzalo, Cèsar Ferri, José Hernández-Orallo, Fernando Martínez-Plumed, and María José Ramírez-Quintana. "Can language models automate data wrangling?." *Machine Learning* 112, no. 6 (2023): 2053-2082; Sarkar, A. (2023, October). Will Code Remain a Relevant User Interface for End-User Programming with Generative AI Models?. In *Proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (pp. 153-167).

engines as they offer a bundle of consumer related services, which in their turn, will enhance trust in the recommendation itself.

Why is this important for competition? Because LLMs can better intermediate the vastness of online information to consumers. While everyone likes a bargain, oftentimes consumers do not have much time to seek them. Yet by using an LLM, a consumer seeking the highest quality product of a certain kind can potentially receive such information in a fast and conversational matter. The combination of an LLM with real-time mapping of online data can further assist consumers in making better decisions. This can be illustrated by the many versions of GPTs for Travel, offered under the ChatGPT interface, which seek to find best offers for flights at any given time, plan your itinerary and more.

## II. ***Limitations and competition concerns concerning Generative AI***

Of course, LLM models, as they currently stand, are not a panacea for all consumption ills. Below we explore some limitations.

First, as studies have shown, LLMs are not similarly efficient in all languages. This implies that consumers searching in English are more likely to enjoy the benefits such technology has to offer, relative to those speaking, for example, Catalan. This is a direct result from the fact that most of the training data fed into existing LLMs is in the most common languages. Furthermore, it results from the fact that most users interact with LLM models in common languages, thereby training the algorithm as mentioned. Every time a user changes a request in a string to make it more accurate or specific, the interaction provides data from which the algorithm can learn for future queries. Such interaction is, in fact, a form of supervised learning, and if it can be generalized, the algorithm becomes better. Yet this problem can be partly solved by the fact that foundation models may enable some degree of transfer learning from one language to another (adaptation).

Second, Generative AI is not trained to provide truthful answers. Rather, it may engage in what computer scientists call hallucinations, which are answers that make sense linguistically, but are incorrect factually. In such cases, the consumer might be given incorrect information about some possibilities. Yet the risk of reliance here is not high, since to complete the transaction the consumer would have to attempt to buy the product/service suggested. Furthermore, creators of LLMs are working on limiting the possibility and frequency of hallucinations, inter alia by constraining filter output by using other machine learning methods or by comparing some answers on veritable sources.<sup>23</sup>

The next two limitations are much more fundamental and carry much more weight. By their nature, LLM models are trained to seek the most plausible next word. This feature of LLMs provides more mainstream answers, leading to what some researchers call "outcome homogenization".<sup>24</sup> For example, a recent study performed by Shur-Ofry on cultural questions (such as "name the three most

---

<sup>23</sup> A fact which, by itself, might give firms like Google who have already mapped the internet, a comparative advantage.

<sup>24</sup>Rishi Bomassani et al., *Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?*, NEURIPS, 1 (2022), <https://arxiv.org/abs/2211.13972v1> ("the phenomenon of individuals (or groups) exclusively receiving negative outcomes from all decision-makers they interact with").

important people in the nineteenth century"), has shown that LLMs provide answers that "are likely to be geared toward the popular and reflect a mainstream and concentrated worldview, rather than a multiplicity of contents and narratives."<sup>25</sup> While similar studies have not yet been performed – to our knowledge – on wider consumer markets (this is our next step in our research agenda),<sup>26</sup> it is easy to imagine similar outcomes in the consumption sphere. For example, if a consumer seeks the best washing machine of a certain type, the answer is most likely going to be the one that most websites overall – and over a long temporal period – have pointed to, rather than a newcomer brand for which new and good reviews are only beginning to accumulate. This result depends, of course, on the dataset: whether it includes enough new text to alter the answer. It also depends on how the consumer structures his question, and whether the model in its predictive function a user's previous revealed preferences. Yet this ingrained pull of the algorithm towards mainstream answers, resulting from its technological features rather than a bug in its code, could have adverse competitive effects, by hiding the much-needed information for consumer choices, limiting the ability of newcomer firms and best offers to be recognized, and negatively affecting market dynamics.

This effect may be exacerbated by five elements, some mentioned as features earlier. The first is limited digital literacy. The more consumers come to rely on Generative AI as a reliable source, not understanding the way it operates or at least how to use it to maximize their benefit, the more it might generate consumption misperceptions. These misperceptions may be strengthened by the fact that many Generative AI LLMs are a general-purpose technology<sup>27</sup> ("Ask me anything..."). Digital literacy-including "prompt engineering", the art and science of learning how to ask the right query- is thus extremely important in our day and age. The second, connected element, is the increased trust of consumers in answers provided by such algorithms (automation bias) which can be aggravated by the nature of the interaction.<sup>28</sup> The fact that consumers are often offered one suggestion, the phrasing of the output in a clear and often confident tone that generates an aura of authority,<sup>29</sup> and the fact that more consumers are likely to "consume" the answers via audio rather than via images (e.g., via voice-activated personal assistants such as Alexa and Siri),<sup>30</sup> exacerbate the tendency to rely on the algorithm. As Blickstein-Shchory and Gal have argued in the context of voice shoppers, this creates "choice gaps".<sup>31</sup> Both elements can lead to a third one, which is lower engagement of individual consumers in the provision of information to other consumers regarding their experience with products, limiting this essential dialogue. This is partly because the likelihood of other human consumers reading their reviews and signaling their appreciation of them is lower in a world where many consumers use an AI intermediary. Fourth, the prevalent use of LLMs could lead to AI Echo-chambers that result from feedback loops, whereby the texts generated by LLMs will percolate back

---

<sup>25</sup> Michal Shur-Ofry, *Multiplicity as an AI Governance Principle* (May 10, 2023), <https://ssrn.com/abstract=4444354>.

<sup>26</sup> Shur-Ofry has also studied recommendations for TV shows and some types of recipes. *Ibid.*

<sup>27</sup> General purpose technologies are characterized by three properties: they have pervasive application across many sectors; they spawn further innovation in application sectors, and they themselves are rapidly improving. See, e.g., Timothy F. Bresnahan & Manuel Trajtenberg, *General Purpose Technologies 'Engines of Growth'?* 65 *J. of Econometrics* 83 (1995).

<sup>28</sup> Shur-Ofry, *id.*, p. 31.

<sup>29</sup> Shur-Ofry, *id.*, p. 12.

<sup>30</sup> Noga Blickstein Shchory & Michal S. Gal, *Voice Shoppers: From Information Gaps to Choice Gaps in Consumer Markets*, 88 *BROOKLYN LAW REVIEW* 111 (2022).

<sup>31</sup> *Id.*

into the web, and serve as training materials for the next generation of LLMs.<sup>32</sup> As OpenAI conceded, AI systems have the potential to "reinforce entire ideologies, worldviews, truths and untruths, and to cement them or lock them in, foreclosing future contestation, reflection, and improvement..." thereby reinforcing conformity of outcomes rather than diversity and multiplicity.<sup>33</sup> Finally, as Shur-Ofry argues, several design and technological traits increase the power of LLMs to influence users' perceptions, offering algorithmic-based content recommendations that do not necessarily meet consumers' true preferences, whether intentionally or not.<sup>34</sup> This is also a result of the distancing of the consumer from the source materials: LLMs generate new text, not exposing the consumer to the original sources, and the model's decisional process (such as the training datasets, the hyperparameters, the human feedback and the values assigned by human trainers). As a result, as Gal argued elsewhere, "a user who is unaware of the algorithm's limitations, would likely not be aware of choices he has forgone".<sup>35</sup>

A final concern relates to the gaming of LLMs. Here we focus on three possibilities. First, LLMs can potentially be tricked by those supplying the information, what consumer scientists call "Adversarial Machine Learning". Studies have already shown that the data on which the LLM is trained affects the quality of its answers. The payoff from such gaming might be especially high, given that the LLMs' decisional mechanism is not transparent, and where trust in Generative AI was generated by uses of the same technology for other purposes (spillover effects). Yet LLM developers may have incentives to limit at least part of these gaming efforts. For example, ChatGPT was trained using human coders who detected toxicity in the training data, making sure OpenAI filter it out before it ever reached the user.<sup>36</sup> Since the model continues to develop over time, such screeners can be used for commercial reasons as well. Second, the LLM can be gamed by users through their queries, given that the model is also trained on user responses.

Third, and most disconcerting, gaming may be performed by the creator of the AI. Output can be governed by heuristics or AI-filter (e.g., trained classifiers), not only the model's parameters. Such filtering can be a way to combat misinformation, but can also amount to self-preferencing, causing the model to emphasize (or ignore) sets of potential outputs in response to specific prompts.<sup>37</sup> Indeed, Generative AI, coupled with data on consumers' preferences, can potentially be used to exploit behavioral biases or heuristics and nudge consumer into choices they would not have otherwise make. Such choice architecture or user interfaces are known in the literature as Dark Patterns.<sup>38</sup> In a forthcoming paper we explore how market dynamics – especially competition among suppliers of

---

<sup>32</sup> Eric Ulken, *Generative AI Can Bring Wrongness at Scale*, NiemanLab, <https://www.niemanlab.org/2022/12/generative-ai-bringswrongness-at-scale/>; Shur Ofry, p. 12.

<sup>33</sup> GPT-4 Technical Report, OpenAI (2023), <https://cdn.openai.com/papers/gpt-4.pdf>, at 49.

<sup>34</sup> Shur-Ofry, *supra*, p. 31.

<sup>35</sup> Michal Gal, *Algorithmic Challenges to Autonomous Choice*, 25 MICHIGAN TECHNOLOGY LAW REVIEW 59, 95 (2018) .

<sup>36</sup> <https://time.com/6247678/openai-chatgpt-kenya-workers/>.

<sup>37</sup> <https://platform.openai.com/docs/guides/moderation/overview> (OpenAI moderation tools 'check whether content complies with OpenAI's usage policies. Developers can thus identify content that our usage policies prohibits and take action, for instance by filtering it.')

<sup>38</sup> Jamie Luguri and Lior Jacob Strahilevitz, *Shining a light on dark patterns*, 13(1) JOURNAL OF LEGAL ANALYSIS 43 (2021); Amit Zac, Yun-Chun Huang, Amédée von Moltke, Christopher Decker, and Ariel Ezrachi, *Dark patterns and online consumer vulnerability*, Available at SSRN 4547964 (2023).

search services, including those based on LLM models, as well as consumer reactions to self-preferencing conduct – affect the motivation and the ability to engage in such conduct.

## **Conclusion**

How might these traits affect market dynamics? In performing their role as infomediaries, LLMs might steer some consumers away from best options. But, more importantly, they can create systemic effects on competition, increasing entry barriers to newcomers or to those who have significantly increased the quality of their products. While LLMs can bring about many benefits, they do not provide the ultimate know-it-all algorithmic consumer.