

# Artificial Intelligence and Cognitive Inequality

Indira Puri<sup>1</sup> and Laura Veldkamp<sup>2</sup>

<sup>1</sup>New York University

<sup>2</sup>Columbia University \*

October 21, 2025

## Abstract

We combine insights from the medical and artificial intelligence (AI) literatures to propose a novel model, which suggests that the expansion of AI may exacerbate cognitive inequality. Information providers maximize profit by tailoring the complexity of content, offering less cognition-enhancing content to less able customers. While individuals with high cognitive abilities may benefit from this increased within-cognitive-group homogeneity, those with lower cognitive abilities—and even children—may suffer adverse effects. Anecdotal data from political discourse and cognitive skills scores are consistent with the model predictions. The findings introduce a new consideration to the debate on financial literacy and AI regulation.

How does the use of consumer data and big data technologies, like artificial intelligence (AI), affect social welfare? At its essence, AI is a prediction technology. It represents an improved ability to predict an outcome or type with a large data set. In most models, reducing prediction errors resolves a friction, improves efficiency

---

\*Thanks to Angelica Munoz Rodriguez for her outstanding research assistance and to Karthik Sastry and Michael Sockin for their helpful comments.

and improves welfare. For improved predictions to be bad, those predictions must be used to worsen some market distortion like market power or involve cognitive mistakes. Since the information we study is provided at zero price, market power is tough to prove. Therefore, we investigate the relationship between AI adoption and cognitive mistakes.

Our main argument is that adoption of AI among content providers causes internet users' cognitive ability to diverge in the short term. Cognitive ability is like a muscle. It needs to be exercised to stay strong. This simple insight from neuroscience has profound implications for modern uses of artificial intelligence (AI) to serve people with targeted information. By allowing more customized information provision, precision information targeting, enabled by AI, allows cognitive ability to diverge. It perpetuates initial inequalities in cognitive ability by serving less-able thinkers with information that fails to challenge them and fails to bring them closer to average ability. While targeted information serves consumers well in the moment, it undermines tolerance for complexity among those who need it most. This could have profound implications for economic inequality, public discourse and democracy.

We develop a model of this harm that informs us about what empirical evidence should reveal its presence. Then, we document some suggestive evidence from social media posts and student evaluations, supporting the idea that this channel is a relevant one. Finally, we use the model to predict what long-run use of information targeting might mean for human cognitive ability in the long run.

To think about the long-run, consider a parallel between the current period and the industrial revolution. As machines were able to perform much of the physical labor of human beings, the average person lost muscle mass and became weaker. If AI can reduce the complexity of information and lift some of our cognitive load, will human beings lose some of our cognitive strength? Our model tells us that this is a likely outcome of information targeting, if people are more averse to overly difficult information than they are to information that is too easy. This, combined with a little bit of type uncertainty, which is surely present, leads information providers to match people with information that is below their optimal level, and depresses the long-run cognitive ability of all. Recent results from an OECD (2024) test of

cognitive ability show a long-run cognitive decline.

Our conclusions stem from two assumptions, which serve as the cornerstone of our analysis. Therefore, Section 1 begins by laying out the evidence in support of these two assumptions. The first assumption posits that cognitive processes, particularly attention and information processing, can be likened to a muscle that requires regular exercise. We contend that engaging with complex information enhances cognitive abilities, whereas limited exposure to such stimuli may lead to cognitive decline.

The second assumption revolves around the unique capabilities of big data and algorithms like AI to tailor information complexity to individual consumers' profiles. Unlike traditional approaches where information dissemination is uniform or coarsely targeted, big data enables precise targeting based on users' precise attributes. This makes it possible to serve individuals with varying levels of cognitive aptitude with information tailored to their comprehension levels.

Section 2 integrates these two assumptions into a cohesive framework to explore the ramifications of AI for cognitive inequality. Information providers are profit maximizing. We capture the predictive power of AI in a coarse way. Without AI, firms do not know customer types. With AI, they know each customer's type exactly. While this is obviously an exaggeration, it captures the idea that AI enables improved prediction of consumers' abilities, constraints and preferences. Consumers observe the information delivered to them, in the process of consuming the good they want. This captures the practice of serving ads or targeted news feeds, in between posts, videos or podcasts that a user wants to read, watch or hear. The consumers' cognitive ability is affected by the news they consume. In a second period, consumers with higher cognitive ability achieve greater utility because they make fewer mistakes on tasks.

Testing this theory is not an easy task. The potential cognitive harm we highlight is a difficult harm to measure because, outside of a laboratory setting, it is typically difficult to observe people's cognitive ability. However, we can observe characteristics of new feeds. Section 3 derives testable predictions from the model. Specifically, it predicts that more pervasive use of AI for information targeting results in more dispersed message complexity and more dispersed cognitive ability. The model also

teaches us that AI per se is not essential for these effects. It is the more targeted matching between information and consumer. Any technological change that improves this match could trigger similar consequences. We make use of this insight in our empirical work.

Since harm to minors is especially salient in the regulatory debate, we tease out predictions for that population specifically. We find that AI content-targeting consistently harms children. By targeting them with low-complexity content, AI-driven information targeting deprives children of the types of cognitively challenging digital content that would help them develop and thrive.

Section 4 provides suggestive evidence consistent with our predictions. First, we explore news feeds. Since much of the discussion about the coarsening of public discourse is about political discourse, we measure the complexity of social media posts made by the official accounts of the two major political parties. While this is obviously a very limited scope of political speech, it is cleanly limited to online content that is purely political. We find an increasing trend in the dispersion of that message complexity, consistent with increasing use of AI to target different users with different messages, while acknowledging other possible causal mechanisms.

The final piece of evidence in Section 4.2 shows a decline in cognitive skills across most OECD countries. While a few countries experience gains, most show decline, especially for literacy-related skills. This finding motivates the dynamic model that follows.

While none of our evidence rules out other channels through which social media might have a disparate effect on grades, the two types of evidence jointly suggest that our cognitive channel might be one worth further exploration and inclusion in the debate about the social value of big data and AI.

If this mechanism is at work, what does that imply for the future? Section 5 uses the model to show that there is a force toward long-run convergence of cognitive ability, as a lower level than before. We consider offsetting equilibrium forces that would arise in a richer model. While the long-run predictions are, by their nature, speculative and untestable, they represent possibilities that are worth considering as we prepare for a future with rapid technological change.

Section 6 concludes with a discussion of the potential ramifications of cognitive inequality. The regulation of artificial intelligence (AI) has become a focal point in contemporary discourse. Our study introduces a novel dimension to the discourse of AI harms by examining its disparate impact on cognition. We use theory and data to argue that predictive data analytics may be exacerbating inequality in cognitive ability. We find evidence of this cognitive divergence in the complexity of online political communication. This raises the concern that financial literacy divergence will fuel more income divergence and that the free exchange of ideas, central to a democracy, may face yet another hurdle.

**Related literature** While many papers point to the efficiency, risk reduction and product quality benefits of data,<sup>1</sup> others argue that firms' use of personal data harms' consumers privacy, in a way that reduces utility.<sup>2</sup> A few studies have explored ways in which firms can use data to price discriminate or target weak-willed customers with harmful products.<sup>3</sup> However, information is often provided to consumers for no monetary cost and is hardly addictive or directly harmful. Yet, much of the public outrage and concern centers around customer targeting by information providers, such as social media platforms.

First and foremost, our work is related to others who study the harms of the digital economy (Bergemann et al., 2023). Some papers discuss data and others AI. Data is the fuel for AI. When data and AI are combined, predictions result. Since our focus is on the ability to predict, both types of papers are relevant. Liu et al. (2021) is a closely related paper that examines a different harm of data, that it is used to exploit temptation. Data is used to predict which consumers who are susceptible to temptation so that firms can target them with ads and sell them goods, which ultimately lower their utility. The idea that data is used to target customers is common in the management literature (Matz, 2022). Moreover, Bergemann and Bonatti (2019) highlight the emergence of social externalities resulting from the widespread

---

<sup>1</sup>For efficiency, see Jones and Tonetti (2020) and Cong et al. (2021). For risk reduction, see Eeckhout and Veldkamp (2022) and for measurement of improved product quality, see Quan (2022).

<sup>2</sup>See Acquisti et al. (2016) for a review with many examples.

<sup>3</sup>See e.g., Galperti et al. (2024), Liu et al. (2023).

use of data. In legal scholarship, Ritter and Mayer (2017) advocate for treating data as a form of property and granting ownership rights to users, while Sokol and Comerford (2015) explore the intersection of big data and market power within the antitrust framework. However the focus on informational complexity is novel and has distinct predictions that matter for the future of labor and of democracy.

The study of complexity arises in behavioral finance, decision theory, and experimental economics (for example, Puri (2025), Hopenhayn and Saeedi (2024), and Fudenberg and Puri (2022)), where subjects are shown to be averse to content they deem too complex. This aversion underlies many shortcomings of household financial decision-making. This is related to the study of attention and rational attention (Sims, 2003; Mackowiak and Wiederholt, 2009; Caplin and Dean, 2015). What our work adds is a theory of where attention, or a tolerance for complexity, comes from.<sup>4</sup>

Many of our ideas are similar to ideas about commitment and addiction. The idea that cognitive work is unpleasant in the short run, beneficial in the long run and that consumers avoid it is similar to DellaVigna and Malmendier (2006)'s analysis of physical exercise. Our theory of this lack of commitment draws from the seminal work on rational addiction by Becker and Murphy (1988). However, the argument that the consumption of complex information falls in this category is novel, in the context of the big data privacy debate. The consequences of cognitive effort aversion differ from those of exercise or drug addiction problems.

In the data privacy literature, Seim et al. (2022) underscores three primary rationales for regulatory intervention: the absence of monetary compensation for users who generate data, the exacerbation of market power through data utilization and algorithmic superiority, and the profit-centered exploitation of data by firms. Our work explores a particular example of firms pursuing profit, at the expense of consumers.

---

<sup>4</sup>In the realm of finance, AI facilitates better prediction and parsing of market trends and investment opportunities (Goldstein et al., 2021; Veldkamp, 2023). A special issue edited and summarized by Goldstein and colleagues provides comprehensive insights into the transformative potential of AI in financial analytics, showcasing how it enables more accurate predictions and enhances decision-making processes within the financial sector.

# 1 Motivating Evidence

Our model is based on two key assumptions: that attention is like exercise and that data, combined with AI is used to target people with content matched to their attributes. Each claim has been researched extensively, the first in a medical literature, and the second in the big data and AI literature. We summarize the evidence across these literatures for each fact. We bring these disparate literatures together to develop a model with new predictions that provides a novel perspective on the heated debate surrounding the benefits and costs of AI. We present evidence of those two claims first, before we proceed to describe the model.

**Mental Exercise.** Our first key assumption is that the ability to process complex information is improved with practice. The more one reads, hears or processes complex information, the more capable a person is of performing this task. In other words, complex thinking is like a muscle that needs to be exercised. The medical literature finds evidence in favor of this key assumption: mental exercise increases cognitive ability. For example Pillai et al. (2011) and Verghese et al. (2003) find that doing crossword puzzles reduces cognitive aging. Devanand et al. (2022) document cognitive improvements for individuals with mild cognitive impairment when they train with crosswords. Sevinc et al. (2021) show that mindfulness training improves cognitive performance in older cognitively functioning adults. Ng et al. (2021) document that computerized cognitive training increase cognitive ability across adults of different ages. Petrella et al. (2023) finds weak support for crosswords helping slow the progression of Alzheimer’s disease. And Jonaitis et al. (2013) notes that game-playing preserves cognitive strengths later in life. Yan et al. (2024), using EEG, find that that mobile phone short-form video addiction tendency correlates with impaired behavioral performance on attention tasks and reduced activity in higher-order brain areas responsible for self-control. Zhang et al. (2023) found correlation between social media use and lower executive functions.

This is a non-exhaustive list of medical studies showing a link between cognitive exercise and cognitive ability across a range of cognitive types and ages. While

some studies, like Aartsen et al. (2002), find no link between cognitive ability and mental exercise, review articles conclude that the preponderance of the evidence favors mental exercise helping with a variety of cognitive outcomes (Jak et al., 2013).

Our model also assumes that people who have a low ability to process complex information tend to avoid such content. This assumption is supported by experimental evidence (Oprea, 2020). Puri (2025) finds that those with higher cognitive ability exhibit less complexity aversion.

**Big Data and Targeting.** Our second main assumption is that the rise in the abundance of data and the ability to harness it with technologies like AI has allowed providers to digital content to more accurately match content to user. As highlighted by Agrawal et al. (2018), AI has significantly enhanced predictive capabilities across diverse domains. One of the quintessential applications of AI and big data lies in predicting human traits, a task pivotal for targeted marketing strategies, political campaigns, and financial analyses.

Predictive analytics leveraging large datasets play a pivotal role in tailoring products, campaigns, information, and advertisements to specific consumer segments (Matz, 2022). For instance, Nickerson and Rogers (2014) provide an extensive analysis of the integration of big data in political campaigns, illustrating how it enables more precise targeting of individuals for various forms of engagement, including donations, volunteering, and expressing support for candidates or issues. There is a large literature on ‘versioning’ information goods, whereby firms change the qualities of products to appeal to different market segments (Bhargava and Choudhary, 2008).

Furthermore, the emergence of information firms, as documented by Dolfma and Van Der Eijk (2018), underscores the commercial value of consumer data. Firms specialize in collecting and selling consumer data to companies seeking insights into target audiences for their products or services. Such data-driven insights empower companies to optimize their marketing strategies by precisely identifying and engaging with potential consumers who are most likely to be receptive.

On the theoretical side, Saeedi et al. (2024), shows that, on a network, the quality of information is better targeted in the personalized case.

## 2 Model

Our simple model is designed to show how the combination of the two assumptions alone leads to the logical conclusion that growth in personalized content - be that through social media, digitization, or AI - can exacerbate cognitive inequality. We use as simple a structure as possible to describe that argument. After setting up the model and extracting its ideas, we examine its predictions and consider the consequences of its predictions for democracy and the economy.

**Model assumptions.** There is a high type consumer  $H$  and a low type consumer  $L$ . The high type has a higher cognitive level  $c_H$  than the low type's cognitive level  $c_L$ , where  $c_L, c_H \in (0, 1)$ .

In period 1, a monopolistic firm offers a bundle  $(v, A_i)$  to consumer of type  $i \in \{L, H\}$ , where  $v$  is a desired product (for example: email, Google search, or social interaction) and  $A_i$  denotes personalized content for type  $i$ ; more specifically,  $A_i \in (0, 1)$  denotes the complexity or informativeness of that personalized content. To isolate the channel of interest - homogeneity as driven by digitization and AI - we assume that the firm has no control over the non-personalized content  $v$  but can affect  $A_i$ , personalized content. We treat price as absorbed into  $v$  (or fixed by platform policy); under logit the price term enters additively and does not interact with the mismatch, so allowing price choice does not affect our targeting results.

The utility of a consumer of type  $i$  is:

$$U_{1,i}(v, A_i) = v - d(|c_{1,i} - A_i|) + \beta U_{2,c_{1,i}}(c_{2,i}) + \epsilon_i$$

where  $c_{1,i}$  is the consumer's period 1 cognitive type;  $v$  is the quality of the product;  $\theta > 0$  and  $d : (0, 1) \rightarrow (0, 1)$  is an increasing, strictly convex function capturing preferences over complexity. Utility in period 2,  $U_{2,i}(c_{2,i})$ , increases in cognitive ability  $c_{2,i}$  in period 2.<sup>5</sup> We assume that  $U_{2,c_{1,i}}(c_{2,i}(c_{1,i}, A_{1,i}))$  is concave and differentiable

---

<sup>5</sup>For example, higher cognitive ability may have a positive impact on human capital and wages, which in turn affects utility. Similarly, a better ability to process complex information may help the individual avoid dominated financial products (Goodman and Puri, 2025).

in  $A_{1,i}$ . Noise term  $\epsilon_i$  follows a Type 1 distribution.

From this, we can derive the probability that the agent chooses to consume the bundle versus to not consume the bundle, given their cognitive level; the utility of not consuming the bundle is fixed at zero as a normalization.

$$P^{i,A_i} := P(\text{consume}(v, A_i)) = \frac{\exp(U_{1,i}(v, A_i))}{\exp(U_{1,i}(v, A_i)) + \exp(0)} \quad (1)$$

$$= \frac{1}{1 + \exp(-v + d(|c_{1,i} - A_i|) - \beta U_{2,i}(c_{2,i}))}. \quad (2)$$

For simplicity, we express the preference over complexity as bounded above by 1 and decreasing in the distance between the complexity of the information consumed and one's cognitive ability. However, we introduce a parameter  $\theta$  to allow for the possibility that the utility cost of information that is too simple might differ from the utility cost of information that is too complex. Specifically, if the information complexity is less than  $i$ 's ability ( $A_i \leq c_{1,i}$ ), then we let  $d(|c_{1,i} - A_i|) = 1 - g(|c_{1,i} - A_i|)$ . If the information consumed is more complex than  $i$ 's cognitive ability, the utility is  $d(|c_{1,i} - A_i|) = 1 - \theta g(|c_{1,i} - A_i|)$ . In both cases,  $g : (0, 1) \rightarrow (0, 1)$  is assumed to be a decreasing, strictly concave differentiable function. If  $\theta > 1$ , this captures the idea that it might be more unpleasant to consume personalized content that is too complex for the consumer's type than it is to consume information that is too simple. However, the symmetric cost case,  $\theta = 1$ , is admissible and a special case of our results.

A key assumption in the model, taken from the medical cognitive exercise literature, is that higher complexity personalized content in period 1 leads to weakly higher cognitive ability in period 2. That is,  $c_{2,i} \in (0, 1)$  is given by  $c_{2,i} = f(A_i)$ , where  $f$  is increasing.

We conceptualize a world without targeting as a pooling equilibrium, where both consumer  $L$  and consumer  $H$  are served the same content,  $A_L = A_H = A$ . A world with digitization allows for personalized content, e.g. the firm can choose  $A_L \neq A_H$ .

**Equilibrium and welfare.** We consider two different economies, with different data regimes. In the economy without targeting, the firm chooses the content complexity to maximize  $P^{L,A} + P^{H,A}$ .<sup>6</sup> In the second economy with personalized content, the firm chooses that content  $A^L, A^H$  to maximize  $P^{L,A^L} + P^{H,A^H}$ . Long-term consumer welfare is given by period 2 utility,  $U_{2,i}$ . To put reasonable bounds on the problem, we assume that  $-d'(c_{1,H} - c_{1,L})(c_{1,H} - c_{1,L}) > \beta U'_{2,L}(f(c_{1,H}))f'(c_{1,H})$ , e.g. the complexity cost to the low type of having to process content equal to the level of the high cognitive type in period 1 exceeds the possible future cognitive benefit.

### 3 Model Predictions

Our main point is to show how AI can produce cognitive divergence. In showing this, we see how any technology that adopts or improves content targeting could generate the same result. This is an important insight for our empirical tests. Since policy makers have been concerned about social media’s effect on young people, we consider how children might be represented in this setting. We demonstrate the greater harm of AI and its accompanying targeted content on them.

For simplicity, we consider the extremes of content that is untargeted and homogeneous versus content that is exactly matched to the consumer’s type. We could instead consider noisy signals about consumer types and model AI as an increase the precision of the firm’s prediction of each consumer’s type.

We could also have, as in Saeedi et al. (2024), modeled interactions on a network instead of as an agent problem. However, for our purposes, a simple two-assumption model - without the additional assumptions needed for various network structures - is sufficient to obtain our results.

---

<sup>6</sup>Note that since there is one consumer of each type, the low and high types have equal population weights.

### 3.1 Information Targeting and Cognitive Divergence

Our first result simply describes the firms' optimal information complexity choices, in each of the two economies – the one without personalized content and the one with. For clarity, we note that we do not require firms to be myopic in deriving this result - they simply maximize the probability of a sale.

**Lemma 1.** *To maximize the probability of the consumer of type  $i$  purchasing the product by selecting product complexity  $A_i$ , it is sufficient to maximize  $1 + \exp(-v + d(|c_{1,i} - A_i|) - \beta U_{2,i}(c_{2,i}))$ ; further, this argument is differentiable and concave.*

*Proof.* The probability that consumer of type  $i$  buys the product is given by Equation 1. This function is strictly increasing in its argument; as a result, maximizing Equation 1 over  $A_i$  is the same as maximizing the argument  $1 + \exp(-v + d(|c_{1,i} - A_i|) - \beta U_{2,i}(c_{2,i}))$  over  $A_i$ .

The argument is differentiable because its components are all differentiable. The argument is concave because  $d$  is convex and increasing in its argument, and  $U_{2,i}(c_{2,i}(A_i))$  is concave and increasing in  $A_i$ .  $\square$

**Proposition 1.** *Let  $A_T^{(L)}$  denote the content complexity with targeting for the low cognitive type, and  $A_T^{(H)}$  content complexity with targeting for the high cognitive type. Let  $A^{NT}$  denote content complexity without targeting. Then  $A_T^{(L)} < A^{NT} < A_T^{(H)}$ .*

*Proof.* The probability that consumer of type  $i$  buys the product is given by Equation 1. From Lemma 1, to maximize we may take a first order condition with respect to the internal argument:

$$\frac{\partial (1 + \exp(-v + d(|c_{1,i} - A_i|) - \beta U_{2,i}(c_{2,i})))}{\partial A_i} = -d'(|c_{1,i} - A_i|)(A_i - c_{1,i}) + \beta U'_{2,i}(f(A_i))f'(A_i).$$

Notice that at  $A_i = c_{1,i}$ , the first term is zero but the second term is positive, so  $A_i^* > c_{1,i}$ . Notice next that  $A_L^* < c_{1,H}$ . This is because, by assumption,  $-d'(c_{1,H} - c_{1,L})(c_{1,H} - c_{1,L}) > \beta U'_{2,L}(f(c_{1,H}))f'(c_{1,H})$ , so that  $\frac{\partial (1 + \exp(-v + d(|c_{1,i} - A_i|) - \beta U_{2,i}(c_{2,i})))}{\partial A_i} < 0$ .

We now turn to the case without personalized content. Let  $P_g(A) = \sigma(z_g(A))$  for  $g \in \{L, H\}$ , where  $\sigma$  is strictly increasing (e.g. logistic), and  $z_g(A) = v_g - d(|c_{1,g} -$

$A|) + \beta U_{2,g}(c_{2,g}(A))$ . From Lemma 1, each  $z_g$  is strictly concave in  $A$ , so that each  $z_g$  has a unique maximizer  $A_g^*$ . Let  $A^*$  maximize  $S(A) \equiv P_L(A) + P_H(A)$ . We wish to prove that  $A_L^* < A^* < A_H^*$ .

Since  $\sigma$  is strictly increasing,  $A_g^* = \arg \max_A z_g(A) = \arg \max_A P_g(A)$  and  $z'_g(A_g^*) = 0$ , with  $z'_g$  changing sign from  $+$  to  $-$  by strict concavity. Define  $w_g(A) \equiv P_g(A)(1 - P_g(A)) \in (0, 1/4]$ . For  $A \neq c_{1,g}$ ,  $\frac{d}{dA} P_g(A) = w_g(A) z'_g(A)$ , so  $S'(A) = \sum_{g \in \{L,H\}} w_g(A) z'_g(A)$ . (At  $A = c_{1,g}$ , interpret the derivative via subgradients; this does not affect signs below.)

To show that  $A^* < A_H^*$ , suppose for a contradiction, that  $A^* \geq A_H^*$ . By single-peakedness of  $z_H$ ,  $z'_H(A^*) \leq 0$ . Since  $A_L^* < A_H^* \leq A^*$ , strict concavity also implies  $z'_L(A^*) < 0$ . With  $w_g(A^*) > 0$ ,  $S'(A^*) = \sum_g w_g(A^*) z'_g(A^*) < 0$ , contradicting optimality of  $A^*$ . Similarly, suppose, symmetrically, that  $A^* \leq A_L^*$ . Then  $z'_L(A^*) \geq 0$  and, because  $A^* \leq A_L^* < A_H^*$ , we also have  $z'_H(A^*) > 0$ . Hence  $S'(A^*) = \sum_g w_g(A^*) z'_g(A^*) > 0$ , again contradicting optimality. Therefore neither  $A^* \geq A_H^*$  nor  $A^* \leq A_L^*$  is possible, so  $A_L^* < A^* < A_H^*$ .  $\square$

This model solution lays the foundation for our main result, which proves divergence in utility between high and low types. The divergence in utility arises because the data-driven targeting exacerbates the divergence in cognitive ability.

**Proposition 2.** *The low type experiences lower utility under the AI regime than without AI. The high type experiences higher utility under the AI regime than without AI.*

*Proof.* Let  $A_T^{(i)}$  denote the complexity of the content shown to type  $i$  under targeting, and  $A_{NT}^{(i)}$  the complexity without targeting.

Since utility in period 2 is the welfare criterion, and is an increasing function of the content complexity in the first period, and from Proposition 1  $A_T^{(L)} < A_{NT}^{(L)}$ , then  $U_{2,L}(A_T^{(L)}) < U_{2,L}(A_{NT}^{(L)})$ .

Since, from Proposition 1  $A_T^{(H)} > A_{NT}^{(H)}$ , then  $U_{2,H}(A_T^{(H)}) > U_{2,H}(A_{NT}^{(H)})$ .  $\square$

## 3.2 Adding Children

We model children as individuals whose cognitive capabilities are not yet fully developed. This means children have the same cognitive capability distribution, but shifted down: they are of type  $L$  and type  $L' = L - (H - L)$ . Since their spending capacity is lower than adults', the firm places  $\iota$  small weight on them in the absence of big data, when deciding which (single) ad to run. With AI, the firm can target ads to children costlessly.

**Proposition 3.** *Children have lower payoffs from the introduction of targeted advertising.*

*Proof.* Since  $d$  is continuous, for  $\iota$  sufficiently small, the proof of Proposition 1 continues to hold, e.g. that the single ad's complexity is greater than  $L$  in the absence of targeting.

With targeting, that complexity reduces to the child's cognitive type. Since welfare is measured by period 2 utility, this means that  $U_2(A_{NT}) > U_2(A_T^{(i)})$  for each child type  $i$ .  $\square$

The reason kids are hurt is that they are served less-sophisticated material than they would be without data targeting. This robs them of the opportunity to challenge themselves and improve their cognitive functioning.

Of course, many children have cognitive capacity that is lower than most adults. Perhaps they would be challenged by  $L$ -type information. In this simple model, they are still hurt because their skills would be improved more with the higher-complexity information. See the discussion of overly-complex information in the next Section.

## 3.3 Complexity vs. Transparency

A potential criticism of this result is the following: Language that is too complex is undecipherable. It might be so confusing to the reader that its meaning and/or educational value are lost. We consider two potential downsides to complexity.

**Costs of overly-complex information.** Perhaps the best information, for the reader’s well-being, challenges them, but not too much. Even in this case, the point still stands that the short-run information of the firm serving the ad is not to challenge the recipient’s cognitive capacity, but to maximize the probability of sale or engagement. That is best accomplished by targeting at or below their current cognitive level.

**Complex language vs. complex ideas** There is a long-standing debate about the merits of using complex language. Defenders argue that complex language expresses complex ideas with more nuance and precision than simple language could. Detractors argue that complex language reflects a writer’s lack of effort to communicate clearly. Neither our model nor our evidence will resolve this debate.

Our argument is not that one form of writing is better or worse, but that complex language is more likely to indicate complex thought. It is a proxy. Complex thought and language both challenge the human brain to exert effort and grow stronger. Even if complex language is the result of poor writing and is ineffective at persuasion, it may still challenge the cognition of the reader or listener. It is this cognitive capacity which is our focus.

## 4 Empirical Exercise

We would ideally like to test if more information targeting leads to a greater divergence in cognitive ability. This is empirically challenging for several reasons. First, there are few natural experiments where the level of targeting changes without external causes. Second, there are few objective measures of cognitive ability, especially in conjunction with targeting.<sup>7</sup>

To provide supportive evidence for our theory, we explore the complexity of online political discourse. This is a direct test of the theory’s predictions. While the findings

---

<sup>7</sup>One policy change we considered was the Plain Language Act of 2010 (US Department of Labor, 2024). However, this was implemented all at once rather than in a staggered manner, and affected only federal communications; it is therefore not clear what the right cognitive outcome variable would be.

are in line with the model’s predictions, the data turn out to be quite noisy and the scope for alternative explanations is large.

We also present evidence of a long-run decline in cognitive skills. This last piece of evidence motivates the long-run dynamic model, presented in the last section.

## **4.1 Time Trend: Variance in Complexity of Political Posts**

We here test whether, over time, the complexity of social media posts by political entities have increased dispersion.

### **4.1.1 Measuring Complexity**

There are many ways to measure language complexity. The Flesch-Kinkaid measure is probably the most well-known measure. It relies on the number of words per sentence and the number of syllables in each word. This metric is simplistic because it mis-classifies complex or less-phonetic, monosyllabic words like “fjord” or “bough,” as well as simple, common words with many syllables, like “anybody” or “congratulations.”

More recently, other measures of language complexity emerged to solve these problems. For example, The Dale–Chall formula uses a list of about 3,000 words familiar to American fourth graders. The score estimates the U.S. grade level required to understand the text, based on how many words are on that fourth-grade list. The score also factors in the average sentence length. Our results make use of this complexity metric.

### **4.1.2 Textual Data Set**

To capture information that is relevant for democratic governance, we focus on the Facebook pages of the Democratic Party (DNC) and the Republican Party (RNC). Specifically, we examine the Democratic Party page, which boasts 1.6 million followers, and the National Republican Congressional Committee page, which commands

TABLE 1: Summary Statistics

<b>Statistic</b>	<b>Democrats</b>	<b>Republicans</b>
Total Number of Posts	10,080	1,946
Average Number of Posts per Month	140	27.03
Min Dale–Chall Readability Score	0	0
Max Dale–Chall Readability Score	35.27	19.62
Average Word Count per Post	26.95	23.04
Standard Deviation of Word Count per Post	16.60	31.26

1 million followers.<sup>8</sup> These figures are similar for RNC/Republican Party.

Our methodology involves analyzing the textual content of each post on these pages since their inception. We exclude any URLs or links from our analysis, as these will be examined separately at a later stage. Within each month of a given year, we measure the dispersion in the complexity of the posts. This involves applying complexity measures to assess the readability and linguistic sophistication of the text. By analyzing the variance in grade level readability across posts within a month, we gain insights into how the complexity of messaging changes over time.

Using the Dale–Chall complexity measure described above, we give each post a score. Table 1 provides summary statistics for each page.

Of course, this text is not always bundled with a purchase, as described in the model. However, it is text that is likely to show up in a Facebook user’s news feed, interspersed with other posts. The post may be more or less effective in achieving its persuasive goals, depending on how aligned its messaging is with user characteristics. The crucial importance of matching political communication with user attributes was revealed in the Cambridge Analytica scandal, described by Matz (2022).

### 4.1.3 Divergence of Complexity in Political News

To test our theory, we examine online communications of the major US political parties. Measuring divergence in political news is an important endeavor, especially

---

<sup>8</sup>Note that the National Republican Congressional Committee, has more followers (1 million) than the Republican Party Page (21k). We therefore choose the page with more followers.

in an era characterized by increasing polarization and the proliferation of partisan media outlets. This exercise supports the theory and shows that it operates in a realm of great importance to society.


The hypothesis is that targeting of information causes the communications to differ more in their complexity. In other words, we are looking for dispersion in linguistic complexity, across posts of the same source. If data and algorithms are getting more precise at targeting recipients with communications designed for their cognitive level, then this dispersion should rise over time.

For each month, we plot the dispersion in complexity scores observed on a given website. This graphical representation, depicted in Figure 1, provides a visual understanding of how the complexity of posts on the DNC and RNC pages evolves over time. The variance in readability of posts is largely increasing over time: bootstrapped standard errors give a statistically significant slope for each complexity measure on each page: regressing the variance in the readability measure on the number of months since the page’s inception gives a point estimate of 0.13 for Democrats (95% confidence interval of [0.11,0.16]) and 0.12 for Republicans (95% confidence interval of [0.08, 0.16]).


While this finding is consistent with the theory, it is far from identifying a causal effect. It does tell us that something is happening to the complexity of political discourse. The model offers an explanation, but not the only possible explanation, for this trend.

## 4.2 The Decline in Cognitive Ability

An OECD study documents a decline in cognitive abilities across the adult population in many countries. The 2023 Survey of Adult Skills, conducted as part of the OECD’s Programme for the International Assessment of Adult Competencies (PIAAC), evaluated the proficiency of adults aged 16 to 65 across 38 countries, including 31 OECD and 7 partner economies. The assessment focused on three core information-processing skills: literacy, numeracy, and adaptive problem-solving in technology-rich environments. Data were collected between 2017 and 2022, depend-



(A) Democrats



(B) Republicans

FIGURE 1: Variance of Complexity of Posts Over Time: Democrats and Republicans

*Note:* This figure plots bootstrapped point estimates and standard errors for the variance within a given month of readability scores for posts on the Democratic Party page (<https://www.facebook.com/democrats>) and the Republican Party page (<https://www.facebook.com/republicaninsider>). The readability score we consider is Dale-Chall. Standard errors are bootstrapped. The Republican NRCC page did not post from Dec 2022 - March 2023.

ing on the country, with the most recent round concluding in 2023. In total, over 220,000 adults participated in the survey, providing nationally representative samples for each participating country.

Participants completed direct assessments measuring their ability to comprehend and work with written texts, use numerical information, and solve problems. Literacy questions required respondents to understand, evaluate, and use written texts encountered in everyday life. Tasks included reading emails, interpreting instructions, or extracting information from newspaper articles. The focus was on the ability to process text, determine relevance, and make inferences, rather than simply decoding words. Numeracy questions involved solving problems that required understanding, interpreting, and using numerical and mathematical concepts in practical contexts. Examples included calculating discounts, comparing unit prices, interpreting graphs and tables, and understanding probabilities—tasks reflecting the quantitative reasoning needed in daily and workplace settings. Problem-solving questions assessed individuals’ capacity to use digital tools to access and evaluate information, navigate software interfaces, and resolve goal-oriented problems. Respondents were asked to complete tasks such as managing emails, filling out online forms, or finding information on a website, requiring them to interact with simulated digital environments. Scores are reported on a 0–500 point scale for each skill domain. In literacy, 226–275 points is the minimum level generally considered necessary to meet the demands of everyday life and work in a knowledge-based economy.

Figure 2 illustrates the trends in cognitive ability for three representative countries in the sample and then also for the group as a whole. The increasing lines indicate an increasing fraction of adults who have a low level of cognitive ability. Table 2 in the Appendix reports the level and change in scores for literacy and numeracy, for each country in the sample.

While a few countries, like Finland for example, experience skill improvements, the majority of countries experience a decline in cognitive skills. The decline in literacy skills is particularly striking. For most countries, cognitive ability was improving between 2004 and 2010. It was only after 2010 that aggregate cognitive ability worsened. Facebook went public in 2012. Thus, the post-2010 period corresponds to a

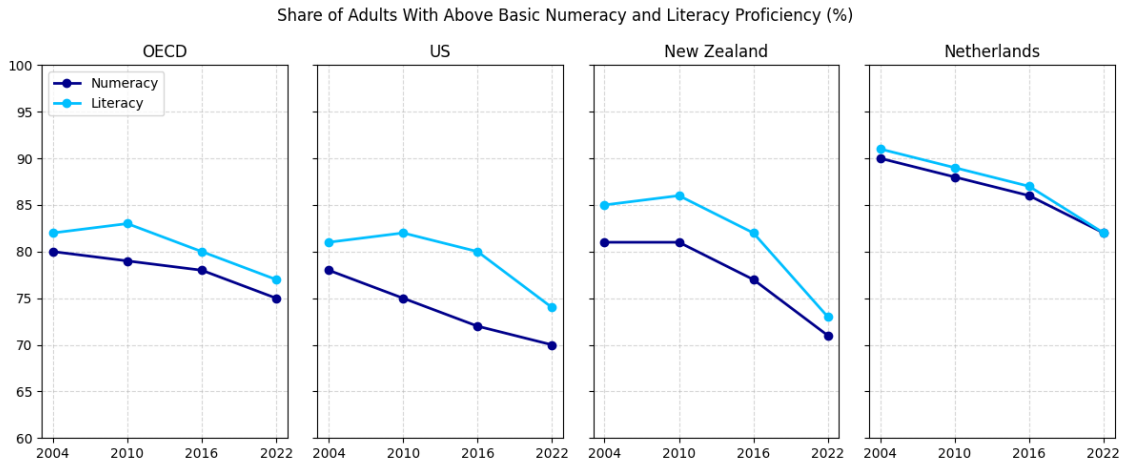


FIGURE 2: Share of Adults with Above Basic Numeracy (dark line) and Literacy (light line) (%). *Note:* Source: OECD 2023 Survey of Adult Skills. Basic numeracy means struggles to use mathematical reasoning when evaluating simple claims. Basic literacy means struggles with inference or integrating multiple pieces of information.

time of rapid expansion in social media use.

A 10–15 point difference in average scores is interpreted as an effect comparable to the difference in average skill levels between people with one to two additional years of education. In many countries with the worst outcomes, the decline in literacy is equivalent to two to four fewer years of education.

### 4.3 Direct Studies on the Impact of AI

We summarize below the findings of five studies on the impact of AI on output and thinking. These are consistent with our model predictions that AI reduces cognitive ability, in the short term for low cognitive ability individuals, and in the longer term for both low and high cognitive ability individuals.

**Undergraduate Essay-Writing (Kosmyna et al., 2025).** Fifty-four undergraduates at MIT were asked to compose a 20-minute essay on a given prompt, either unaided or with ChatGPT’s assistance. EEG recordings revealed that AI users showed significantly lower neural activation in regions associated with creative

thought and attention.

**Knowledge Workers Survey (Lee et al., 2025).** A panel of 319 professionals who used generative AI at least weekly reported on over 900 real-world tasks. Only about 60% of these tasks required substantive critical review of AI outputs; the remainder were effectively ‘mindless,’ and most respondents noted reduced cognitive effort.

**Critical-Thinking Assessment (Gerlich, 2025).** Six hundred sixty-six British adults completed a standardized critical-thinking battery after reporting their frequency of AI use. Heavy AI users scored lower on every subtest.

**Divergent Thinking (Kumar et al., 2025).** In a randomized experiment with 460 participants, exposure to AI-generated ideas for everyday objects led subjects to produce less novel and less varied uses themselves.

**Novice Programmers (Prather et al., 2024).** Twenty-one coding novices attempted a basic programming exercise with or without an AI assistant that interjected both solutions and ‘provocations.’ Novices receiving repeated prompts performed worse overall.

Together, these diverse samples and studies are consistent with the core channel of our theory: differential matching of content complexity to user ability, amplified by AI’s targeting precision, drives short-run productivity gains for high-ability individuals and cognitive offloading (and eventual skill decay) for lower-ability users.

## 5 Long-Run Cognitive Convergence

Given that some modest evidence is consistent with our mechanism, what might this teach us about what is to come? We use the model to show that a likely outcome is that after the divergence, we should expect to see long-run convergence, but at a lower cognitive level.

## 5.1 Model

To analyze long-run dynamics, we extend the static model by introducing (i) generational dynamics, (ii) type uncertainty, and (iii) a lower bound on cognitive ability. These features were omitted from the short-run analysis but are essential for studying intertemporal effects.

As before, there are two consumer types,  $H$  and  $L$ , with initial cognitive abilities always greater than the cognitive lower bound and satisfying  $c_{-1,H} > c_{-1,L} > \underline{c}$ . Cognitive ability evolves according to:

$$c_{i,t} = \gamma c_{i,t-1} + (1 - \gamma)A_{i,t}, \quad (3)$$

where  $\gamma \in (0, 1)$  is a persistence parameter and  $A_{i,t}$  is the complexity of the content consumed. In this overlapping generations model, the law of motion reflects intergenerational transmission of ‘tolerance for complexity,’ whereas in a single-agent dynamic (Section 2) it reflects within-person skill carryover.

In the static setting, firms with AI and large data sets were assumed to observe consumer types perfectly—an unrealistic assumption. In practice, firms observe noisy signals of type, with errors that decrease as algorithms improve. We introduce this type uncertainty by assuming firms observe a noisy signal of type. Firms start out in period 1 with a moderately noisy signal:  $s_{i,1} = c_{i,0} + \varepsilon_1$ ,  $\varepsilon_1 \sim \mathcal{N}(0, \sigma_1^2)$ . In each subsequent period  $t$ , signals become increasingly precise:

$$s_{i,t} = c_{i,t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_t^2), \quad \sigma_t^2 \leq \sigma_{t-1}^2, \quad \lim_{t \rightarrow \infty} \sigma_t > 0.$$

The prior mean of the firm for the ability of for consumer  $i$  in period  $t$  is  $c_{i,t-1}$ , its prior variance is  $\eta_{i,t}$ , bounded away from zero in all periods  $t$ , and its prior distribution is a normal.

Each consumer of type  $i$  in period  $t$  derives utility:

$$U_{t,i}(v, A_{i,t}) = v - d(|c_{i,t} - A_{i,t}|) + \beta W(f(c_{i,t+1})) + \varepsilon_{t,i}, \quad (4)$$

where  $\varepsilon_{t,i}$  follows a Type I extreme value distribution,  $v$  is product quality, and  $d(\cdot)$  captures the cost of a mismatch between content complexity and consumer ability. We assume  $W(f(\cdot))$  is strictly concave and differentiable. The mismatch cost is given by:

$$d(|c_{i,t} - A_{i,t}|) = \begin{cases} 1 - g(|c_{i,t} - A_{i,t}|), & \text{if } A_{i,t} \leq c_{i,t}, \\ 1 - \theta g(|c_{i,t} - A_{i,t}|), & \text{if } A_{i,t} > c_{i,t}, \end{cases} \quad (5)$$

where  $\theta > 1$  and  $g : (0, 1) \rightarrow (0, 1)$  is differentiable, strictly concave and decreasing. This captures the idea that content that is too complex is more harmful than content that is too simple. The outside option of not consuming delivers zero utility. To add more structure to this multi-period model, we further assume that  $g = f^{-1}$ , so that the cost of mismatch is exactly equal to the inverse of the future cognitive gain from overshooting; that  $f', W'$  are bounded above by 1; and that  $\theta > 3$  for a sufficiently strong complexity aversion discontent.

**Optimum** Given the distribution of  $\varepsilon_{t,i}$ , the probability that a consumer of type  $i$  consumes the content is:

$$P_{i,A_{i,t}} = \frac{1}{1 + \exp(-v + d(|c_{i,t} - A_{i,t}|) - \beta W(c_{i,t+1}))}. \quad (6)$$

In each period, the firm maximizes the probability of consumption by solving:

- Without AI (no targeting):  $\max_{A_0} [P_{L,A_0} + P_{H,A_0}]$
- With AI and  $t \geq 1$  (with targeting):  $\max_{A_{t,i}} P_{i,A_{t,i}}$  for each type  $i$

## 5.2 Optimal Targeting with Type Uncertainty

The next result shows that, even with AI, the firm sets the content complexity  $A_H^*$  below the ability of high-type content consumers  $c_H$ , in expectation. This result does not depend on the dynamic structure of the model. It simply describes how the static optimum depends on type uncertainty.

**Proposition 4** (Optimal Targeting Under Type Uncertainty). *In the AI model with type uncertainty, the firm's optimal targeted content complexity  $A_i^*$ ,  $i \in \{L, H\}$  satisfies*

$$A_i^* \leq \mathbb{E}[c_i \mid s],$$

with strict inequality in expectation.

*Proof.* We illustrate the argument for type  $H$  (the case for type  $L$  is analogous).

The firm maximizes the probability of purchase, which requires maximizing expected utility:

$$\max_{A_H \in (0,1)} \mathbb{E}_{c_H|s} \left[ v - d(|c_H - A_H|) + \beta W(f(A_H)) \right].$$

By the assumptions made in Section 5.1 that  $g = f^{-1}$  and  $f' \leq 1$ , we have  $g'(y) = 1/f'(f^{-1}(y)) \geq 1$ .

The first-order condition is:

$$\mathbb{E}_{c_H|s} \left[ \mathbf{1}_{\{c_H \geq A_H\}} g'(c_H - A_H) - \mathbf{1}_{\{c_H < A_H\}} \theta g'(A_H - c_H) \right] + M(A_H) = 0,$$

where  $M(A_H) = \beta W'(f(A_H))f'(A_H) \leq \beta$  by the bounds on  $W'$  and  $f'$  in Section 5.1.

At  $A_H = s_0 := \mathbb{E}[c_H \mid s]$ , the symmetry of the normal posterior implies:

$$\mathbb{E} \left[ \mathbf{1}_{\{c_H \geq s_0\}} g'(c_H - s_0) - \mathbf{1}_{\{c_H < s_0\}} \theta g'(s_0 - c_H) \right] = (1 - \theta) \cdot \mathbb{E}[g'(s_0 - c_H) \cdot \mathbf{1}_{\{c_H < s_0\}}].$$

Since  $g'(z) \geq 1$  for all  $z \geq 0$ :

$$\mathbb{E}[g'(s_0 - c_H) \cdot \mathbf{1}_{\{c_H < s_0\}}] \geq \frac{1}{2},$$

so the expected marginal cost at  $s_0$  is  $\leq -(\theta - 1)/2$ .

The total derivative at  $s_0$  is therefore:

$$\left. \frac{\partial}{\partial A_H} \right|_{A_H=s_0} \leq -\frac{\theta - 1}{2} + \beta < 0,$$

where the inequality follows from  $\theta > 3$  and  $\beta < 1$ .

Since the objective is strictly concave and the derivative at  $s_0$  is strictly negative, we have  $A_H^* < s_0 = \mathbb{E}[c_H | s]$ .

Taking expectations over signals:  $\mathbb{E}_s[A_H^*(s)] < \mathbb{E}_s[s_0] = \mathbb{E}[c_H]$ . □

This teaches us that uncertainty lowers the firm's optimal choice of complexity. The reason is that the risk of overly-complex information exceeds the costs of overly-simple information. This assumption is consistent with the experimental aversion to overly-complex objects in Puri (2025). This difference creates concavity in the firm's objective function. With a concave objective function in complexity, uncertainty lowers the optimal choice of complexity. The firm chooses to hedge the greater risk. The greater risk is that of too much complexity. Therefore, an uncertain firm chooses a simpler message.

Whether targeting harms all consumers, or benefits a subset, depends on the amount of type uncertainty. It also hinges on an assumption we made earlier, that overly-complex information has a greater utility cost than overly-simple information. Easy tasks are less unpleasant than ones that strain our abilities. Our bodies are designed to conserve energy for survival. This includes energy used by the brain. We are designed to avoid tasks that are overly difficult, even if doing them might benefit us in the long run. This asymmetry is what makes it likely that, while cognitive ability may diverge in the short run, it will eventually fall and converge, for all human kind.

### 5.3 Short-Term Divergence, Long-Term Convergence

Under the model of Section 5.1, we now show that the cognitive gap widens after Period 1, when the firm receives moderately precise signals, relative to the cognitive gap after Period 0, when the firm had no type information.

**Proposition 5.** *In Period 1 (moderately precise signals), relative to period 1' (no information), the cognitive gap is wider:  $c_{1,H} - c_{1,L} > c_{1',H} - c_{1',L}$ .*

*Proof.* From the proof of Proposition 1, in Period 1', the firm selects a pooling complexity content,  $A_0$ . The gap after Period 1' is therefore  $c_{1',H} - c_{1',L} = \gamma(c_{0,H} - c_{0,L})$ .

For Period 1, when the firm receives a noisy signal, for each type  $i$ , the firm maximizes  $E_{c_i|s_{i,1}}[P_{i,A_{i,1}}]$ . The first-order condition is:

$$E_{c_i|s_{i,1}}[P_{i,A_{i,1}}(1 - P_{i,A_{i,1}})[-d'(|c_{i,0} - A_{i,1}|)|c_{i,0} - A_{i,1}| + \beta U'_{i+1}(c_{i+1,1})(1 - \gamma)]] = 0 \quad (7)$$

Due to the asymmetric cost parameter  $\theta > 1$ , the optimal targeting satisfies  $A_{i,1} < E[c_{i,0}|s_{i,1}]$  for all  $i$ . This is similar to the proof of Proposition 4.

The cognitive gap becomes  $c_{1,H} - c_{1,L} = \gamma(c_{0,H} - c_{0,L}) + (1 - \gamma)(A_{1,H} - A_{1,L})$ . This is larger than  $c_{1',H} - c_{1',L}$  when  $A_{1,H} - A_{1,L} > 0$ . But this is always true, since the firm receives a single signal  $\epsilon_1$ , and therefore updates its priors  $E[c_{0,H}] > E[c_{0,L}]$  in the same direction.  $\square$

We next show that, in the long term, due to the lower bound on cognition  $\underline{c}$ , cognitive ability for the low and high types coverage, to a lower level than they were in the world where targeting was never introduced.

**Proposition 6.** *As  $t \rightarrow \infty$ , the cognitive gap narrows: for  $t$  large enough,  $c_{t,H} - c_{t,L} < c_{0,H} - c_{0,L}$ . Further,  $c_{t,L}, c_{t,H} \rightarrow \underline{c}$ .*

*Proof.* From the proof of Proposition 4, with noisy signals, the firm sets  $A_{i,t} < c_{i,t-1}$  for both types. This implies that cognitive ability decreases over time for both types: since  $c_{i,t} = \gamma c_{i,t-1} + (1 - \gamma)A_{i,t} < \gamma c_{i,t-1} + (1 - \gamma)c_{i,t-1} = c_{i,t-1}$ .

Since  $\{c_{i,t}\}$  is a decreasing sequence bounded below by  $\underline{c}$ , it must converge to some limit  $L_i \geq \underline{c}$ .

At any limit point  $L_i$ , the dynamics require  $L_i = \gamma L_i + (1 - \gamma)A_{i,\infty}$ , which implies  $A_{i,\infty} = L_i$ . However, from Proposition 4 and its proof, we know that the firm's optimal choice satisfies  $A_i^* \leq \mathbb{E}[c_i|s] - \xi$  for some uniform gap  $\xi > 0$ . This gap exists because: (i)  $g'(y) \geq 1$  for all  $y$  (as established in the proof of Proposition 4), (ii)  $\theta > 3$ , and (iii) the derivative of the firm's objective at  $\mathbb{E}[c_i|s]$  is bounded below by  $-(\theta - 1)/2 + \beta < -1/2 < 0$ .

Therefore, if  $L_i > \underline{c}$ , we would have  $A_{i,\infty} \leq L_i - \xi < L_i$ , contradicting the requirement that  $A_{i,\infty} = L_i$ . Hence, both types must converge to exactly  $\underline{c}$ :  $\lim_{t \rightarrow \infty} c_{i,t} = \underline{c}$  for  $i \in \{L, H\}$ .

As a result, the cognitive gap vanishes in the long run:  $\lim_{t \rightarrow \infty} (c_{H,t} - c_{L,t}) = 0 < c_{0,H} - c_{0,L}$ .  $\square$

**The effect of equilibrium wages** As high-cognitive types become scarcer, the returns to higher cognitive ability may rise, inducing some to work hard to retain their cognitive ability. One could extend the model with a equilibrium labor market and various skill types to achieve such an outcome. At the same time, the use of AI in the workplace could reduce the need for complex cognition there as well.

If we return to the industrial revolution analogy from the start, it suggests another force that might prevent the collapse of cognition. Perhaps cognitive ability becomes something that people do for recreation and health. We might derive longevity benefits or utility from maintaining cognitive ability, just like many gym-users do today by maintaining their physical strength.

While these and many unforeseen forces may arise to prevent the extreme convergence our results predict, most equilibrium effects do not fully counter the impulse that generated them. In this case, it may be that none fully offset the cognitively depressive effects of AI. While this is speculative, good planning for the future requires a degree of speculation and contingency planning, in times of rapid change.

## 6 Conclusion: Social Consequences

If cognitive ability is like a muscle that needs to be exercised, and if data is used to target people with different characteristics and deliver them different content, the natural outcome is that profit-maximizing firms will deliver content that exacerbates cognitive inequality. AI-driven content provision is like placing candy on the plate of every diner that has a hard time finishing their vegetables. The candy provides a short-run rush, makes them less healthy and happy later on and is most harmful to those who were already struggling to build healthy habits.

Complexity diverges in political discourse on social media and increases and rises in conjunction with information targeting technologies. We document these patterns and show that they are consistent with our cognitive divergence mechanism.

Divergence in cognitive capabilities among Americans carries significant economic and political ramifications. Economically, cognitive disparities can exacerbate existing inequalities in income, wealth, and employment opportunities. Those with higher cognitive abilities are often better equipped to excel in complex and lucrative fields such as technology, finance, and management, thereby widening the income gap between cognitive elites and those with limited cognitive capacities. Since the highly able are often born to parents with similar ability, cognitive stratification inhibits those with lower initial ability from catching up, hindering inter-generational social mobility. Long-term stratification can also contribute to social unrest and political polarization, as disenfranchised individuals feel excluded from the benefits of technological progress.

Furthermore, cognitive disparities may intensify ideological polarization. Cognitive elites may gravitate towards political ideologies and movements that offer nuanced and multi-faceted solutions, advocating for policies that favor intellectual achievement and economic competitiveness. Conversely, individuals with lower cognitive abilities may be drawn to populist or authoritarian narratives that promise simplistic solutions to complex societal problems, such as nativism, protectionism, or anti-establishment rhetoric. This ideological polarization can undermine social cohesion, compromise the functioning of democratic institutions, and impede efforts to address pressing social challenges.

Of course, big data and new data algorithms offer many benefits, which were not the focus of this study. Firms can predict demand and supply more efficiently, consumers find products they prefer and innovation may accelerate. Throughout history, similar advances in technology have lifted people out of poverty and improved lives. But such advances usually have some adverse effects that societies need to understand and to mitigate. One such adverse effect may be the reshaping of communication and cognition.

## References

- Aartsen, M. J., C. H. M. Smits, T. Van Tilburg, K. C. P. M. Knipscheer, and D. J. H. Deeg.** 2002. “Activity in Older Adults: Cause or Consequence of Cognitive Functioning? A Longitudinal Study on Everyday Activities and Cognitive Performance in Older Adults.” *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 57: P153–P162.
- Acquisti, Alessandro, Curtis Taylor, and Liad Wagman.** 2016. “The Economics of Privacy.” *Journal of Economic Literature*, 54: 442–92.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb.** 2018. *Prediction Machines*. Boston, MA:Harvard Business Review Press.
- Becker, Gary S, and Kevin M Murphy.** 1988. “A Theory of Rational Addiction.” *Journal of Political Economy*, 96: 675–700.
- Bergemann, Dirk, and Alessandro Bonatti.** 2019. “The Economics of Social Data.” *SSRN Electronic Journal*.
- Bergemann, Dirk, Jacques Crémer, David Dinielli, Carl-Christian Groh, Paul Heidhues, Maximilian Schäfer, Monika Schnitzer, Fiona Scott Morton, Katja Seim, and Michael Sullivan.** 2023. “Market design for personal data.” HAL Post-Print hal-04213208.
- Bhargava, Hemant K., and Vidyanand Choudhary.** 2008. “When Is Versioning Optimal for Information Goods?” *Management Science*, 54: 1029–1035.
- Caplin, Andrew, and Mark Dean.** 2015. “Revealed Preference, Rational Inattention, and Costly Information Acquisition.” *American Economic Review*, 105: 2183–2203.
- Cong, Lin William, Danxia Xie, and Longtian Zhang.** 2021. “Knowledge accumulation, privacy, and growth in a data economy.” *Management Science*, 67: 6480–6492.

- DellaVigna, Stefano, and Ulrike Malmendier.** 2006. “Paying Not to Go to the Gym.” *American Economic Review*, 96: 694–719.
- Devanand, D. P., Terry E. Goldberg, Min Qian, Sara N. Rushia, Joel R. Sneed, Howard F. Andrews, Izael Nino, Julia Phillips, Sierra T. Pence, Alexandra R. Linares, Caroline A. Hellegers, Andrew M. Michael, Nancy A. Kerner, Jeffrey R. Petrella, and P. Murali Doraiswamy.** 2022. “Computerized Games versus Crosswords Training in Mild Cognitive Impairment.” *NEJM Evidence*, 1.
- Dolfsma, Wilfred, and Rene Van Der Eijk.** 2018. “How info-firms use big data to target customers.” *Journal of Business Strategy*, 39: 50–55.
- Eeckhout, Jan, and Laura Veldkamp.** 2022. “Data and Market Power.”
- Fudenberg, Drew, and Indira Puri.** 2022. “Simplicity and Probability Weighting in Choice under Risk.” *AEA Papers and Proceedings*, 112: 421–425.
- Galperti, Simone, Aleksandr Levkun, and Jacopo Perego.** 2024. “The value of data records.” *Review of Economic Studies*, 91: 1007–1038.
- Gerlich, Michael.** 2025. “AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking.” *Societies*, 15: 6. Publisher: MDPI AG.
- Goldstein, Itay, Chester S Spatt, and Mao Ye.** 2021. “Big Data in Finance.” *The Review of Financial Studies*, 34: 3213–3225.
- Goodman, Aaron, and Indira Puri.** 2025. “Overvaluing simple bets: Evidence from the options market.” *Journal of Financial Economics*, 172: 104140.
- Hopenhayn, Hugo, and Maryam Saeedi.** 2024. “Optimal Simple Ratings.” *Working Paper*.
- Jak, Amy J., Adriana M. Seelye, and Sarah M. Jurick.** 2013. “Crosswords to Computers: A Critical Review of Popular Approaches to Cognitive Enhancement.” *Neuropsychology Review*, 23: 13–26.

- Jonaitis, Erin, Asenath La Rue, Kimberly D. Mueller, Rebecca L. Koscik, Bruce Hermann, and Mark A. Sager.** 2013. “Cognitive activities and cognitive performance in middle-aged adults at risk for Alzheimer’s disease.” *Psychology and Aging*, 28: 1004–1014.
- Jones, Charles I, and Christopher Tonetti.** 2020. “Nonrivalry and the Economics of Data.” *American Economic Review*, 110: 2819–58.
- Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes.** 2025. “Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task.” Version Number: 1.
- Kumar, Harsh, Jonathan Vincentius, Ewan Jordan, and Ashton Anderson.** 2025. “Human Creativity in the Age of LLMs: Randomized Experiments on Divergent and Convergent Thinking.” 1–18. Yokohama Japan:ACM.
- Lee, Hao-Ping (Hank), Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson.** 2025. “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers.” 1–22. Yokohama Japan:ACM.
- Liu, Zhuang, Michael Sockin, and Wei Xiong.** 2021. “Data Privacy and Temptation.” Princeton University. Economics Department. Working Papers 2021-77.
- Liu, Zhuang, Michael Sockin, and Wei Xiong.** 2023. “Data privacy and algorithmic inequality.” National Bureau of Economic Research.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2009. “Optimal Sticky Prices under Rational Inattention.” *American Economic Review*, 99: 769–803.
- Matz, Sandra.** 2022. *The Psychology of Technology: Social Science Research in the Age of Big Data.* . 1st ed., American Psychological Association.

- Ng, Nicole F., Allen M. Osman, Kelsey R. Kerlan, P. Murali Doraiswamy, and Robert J. Schafer.** 2021. “Computerized Cognitive Training by Healthy Older and Younger Adults: Age Comparisons of Overall Efficacy and Selective Effects on Cognition.” *Frontiers in Neurology*, 11: 564317.
- Nickerson, David W., and Todd Rogers.** 2014. “Political Campaigns and Big Data.” *Journal of Economic Perspectives*, 28: 51–74.
- OECD.** 2024. “Survey of Adult Skills (PIAAC) 2023: Insights and Interpretations.”
- Oprea, Ryan.** 2020. “What Makes a Rule Complex?” *American Economic Review*, 110: 3913–3951.
- Petrella, Jeffrey R., Andrew M. Michael, Min Qian, Adaora Nwosu, Joel Sneed, Terry E. Goldberg, Davangere P. Devanand, and P. Murali Doraiswamy.** 2023. “Impact of Computerized Cognitive Training on Default Mode Network Connectivity in Subjects at Risk for Alzheimer’s Disease: A 78-week Randomized Controlled Trial.” *Journal of Alzheimer’s Disease*, 91: 483–494.
- Pillai, Jagan A., Charles B. Hall, Dennis W. Dickson, Herman Buschke, Richard B. Lipton, and Joe Verghese.** 2011. “Association of Crossword Puzzle Participation with Memory Decline in Persons Who Develop Dementia.” *Journal of the International Neuropsychological Society*, 17: 1006–1013.
- Prather, James, Brent Reeves, Juho Leinonen, Stephen MacNeil, Arisoa S. Randrianasolo, Brett Becker, Bailey Kimmel, Jared Wright, and Ben Briggs.** 2024. “The Widening Gap: The Benefits and Harms of Generative AI for Novice Programmers.” Version Number: 1.
- Puri, Indira.** 2025. “Simplicity and Risk.” *The Journal of Finance*, 80: 1029–1080.
- Quan, Junjun.** 2022. “Tracing Out International Data Flow: The Value of Data and Privacy.” Available at SSRN 4185169.
- Ritter, Jeffrey, and Anna Mayer.** 2017. “Regulating Data as Property.” *Duke Law and Technology Review*, 16: 220–277.

- Saeedi, Maryam, Yikang Shen, and Ali Shourideh.** 2024. “Getting the Agent to Wait.” *Working Paper*.
- Seim, Katja, Dirk Bergemann, Jacques Cre, David Dinielli, Carl-Christian Groh, Paul Heidhues, Maximillian Schaefer, Monika Schnitzer, Ludwig-Maximilians-University Munich, Fiona M Scott Morton, and Michael Sullivan.** 2022. “Market Design for Personal Data.” *The Tobin Center at Yale University: Digital Regulation Project*.
- Sevinc, Gunes, Johann Rusche, Bonnie Wong, Tanya Datta, Robert Kaufman, Sarah E. Gutz, Marissa Schneider, Nevyana Todorova, Christian Gaser, Götz Thomalla, Dorene Rentz, Bradford D. Dickerson, and Sara W. Lazar.** 2021. “Mindfulness Training Improves Cognition and Strengthens Intrinsic Connectivity Between the Hippocampus and Posteromedial Cortex in Healthy Older Adults.” *Frontiers in Aging Neuroscience*, 13: 702796.
- Sims, Christopher A.** 2003. “Implications of rational inattention.” *Journal of Monetary Economics*, 50: 665–690.
- Sokol, D. Daniel, and Roisin Comerford.** 2015. “Antitrust and Regulating Big Data.” *George Mason Law Review*, 23.
- US Department of Labor.** 2024. “Plain Language.” United States Government.
- Veldkamp, Laura.** 2023. “Valuing data as an asset.” *Review of Finance*, 27: 1545–1562.
- Vergheese, Joe, Richard B. Lipton, Mindy J. Katz, Charles B. Hall, Carol A. Derby, Gail Kuslansky, Anne F. Ambrose, Martin Sliwinski, and Herman Buschke.** 2003. “Leisure Activities and the Risk of Dementia in the Elderly.” *New England Journal of Medicine*, 348: 2508–2516.
- Yan, Tingting, Conghui Su, Weichen Xue, Yuzheng Hu, and Hui Zhou.** 2024. “Mobile phone short video use negatively impacts attention functions: an EEG study.” *Frontiers in Human Neuroscience*, 18: 1383913.

Zhang, Kuo, Peiyu Li, Ying Zhao, Mark D Griffiths, Jingxin Wang, and Meng Xuan Zhang. 2023. “Effect of Social Media Addiction on Executive Functioning Among Young Adults: The Mediating Roles of Emotional Disturbance and Sleep Quality.” *Psychology Research and Behavior Management*, Volume 16: 1911–1920.

## A Data Details

The OECD reports details of the cognitive test scores for each country. We report these results in Table . The survey also included Switzerland, Croatia, Latvia, and Portugal in its second round in 2022. But because they were not included in the first round and we cannot report changes in cognitive ability, we drop those countries from the sample.

Country	Literacy	Literacy Change	Numeracy	Numeracy Change
Finland	296	<b>15</b>	294	<b>17</b>
Japan	289	-6	290	-4
Sweden	284	5	285	6
Norway	281	4	285	<b>8</b>
Netherlands	280	3	284	6
Estonia	276	1	281	<b>9</b>
Flemish Region (BE)	273	<b>9</b>	279	<b>8</b>
Denmark	273	<b>9</b>	279	8
England (UK)	272	-1	268	<b>7</b>
Canada	271	-1	271	<b>7</b>
Germany	266	0	273	5
Ireland	263	-2	263	5
Czechia	260	<b>-9</b>	263	-3
New Zealand	260	<b>-21</b>	256	<b>-15</b>
United States	258	-12	249	<b>-7</b>
France	255	<b>-7</b>	252	-3
Singapore	255	-3	247	<b>-17</b>
Austria	254	<b>-12</b>	267	<b>-5</b>
Slovak Republic	254	<b>-20</b>	261	<b>-15</b>
Korea	249	<b>-23</b>	253	<b>-10</b>
Hungary	248	<b>-15</b>	246	<b>-17</b>
Spain	245	-10	243	-6
Italy	245	-8	239	-4
Israel	244	-10	246	-4
Lithuania	244	<b>-10</b>	238	<b>-22</b>
Poland*	236	<b>-31</b>	239	<b>-21</b>
Chile	218	<b>-2</b>	214	<b>-4</b>

TABLE 2: Adult Skills Survey OECD (2024). Statistically significant changes between 2012–2022 at the 5% level are highlighted in bold.