

AI and the Quantity and Appeal of Creative Products: Have LLMs Boosted Creation of Valuable Books?*

Imke Reimers
Cornell University and ZEW

Joel Waldfogel
University of Minnesota, NBER, ZEW

September 15, 2026

Abstract

Monthly book releases tripled from 2022 to late 2025 as LLMs diffused, while the volume of non-AI titles remained constant. Average book appeal – ratings received, sales ranks, and average star ratings – worsened, but the surge in releases raised the number of books with modest appeal. AI-containing books are less used, and their growth, complementing a steady stream of human releases, explains the decline in average appeal. A nested logit calibration shows that the AI books raised consumer surplus only modestly. Previously active authors have increased their releases in the LLM era.

*We are grateful to seminar participants at the NBER Winter 2026 Digitization meetings, New York University, the University of Illinois, the University of South Carolina, the 2026 IIOC, Netflix, the University of Minnesota, the Workshop for Empirical Methods in Intellectual Property, the University of Vienna, the Econometric Society Interdisciplinary Frontier Economics and AI+ML Meeting, the Creative Industries Policy and Evidence Centre, the ZEW Mannheim ICT Conference, and the 9th AIEA-NBER Conference on Innovation and Entrepreneurship at Singapore Management University. We thank Shane Greenstein, Ginger Jin, Franziska Vogel, and Jiayu Yao for helpful comments, and we thank Pangram Labs for support with the use of their AI detection tool.

Introduction

Artificial intelligence, in the form of large language models (LLMs), promises to revolutionize many tasks, occupations, and industries, including – and perhaps especially – the creative industries. Because of LLMs’ ability to create text, book publishing may be particularly susceptible; and during the period 2022-2025, LLMs became capable of producing sustained text passages. LLMs reduce the cost of producing books; and reductions in entry costs raise the number of new products brought to market. Additional entry can be especially beneficial when product appeal is unpredictable, so that some new products have realized appeal in, or near, the right tail of the distribution.

The appearance of LLMs does not simply facilitate more “draws from the urn.” LLMs may also change the distribution of quality realizations, for example by attracting new authors of varying skill, while possibly altering the output of skilled incumbent writers. The net effects of additional author entry and changes in incumbent activity can bring new books of uncertain appeal. For example, if LLMs allow the creation of books with low average appeal and little variance, on top of what human creation alone delivers, the LLM era would bring no additional books in the right tail of the appeal distribution and little unpredictable benefit to consumers. By contrast, if the arrival of LLMs facilitated the creation of more books drawn from a distribution resembling the pre-LLM human distribution, then the number of high-appeal new releases would rise, with potentially substantial welfare benefits. Hence, the welfare effect of AI in the book market depends on what it does to both the number and appeal distribution of new, AI-enabled books, as well as to the continued supply of human books.

These considerations motivate this paper’s questions. First, how do LLMs affect both participation in authorship and the number of new books created? Second, how does the arrival of LLMs affect the appeal of books created, i.e., the distributions of usage, sales, and

the satisfaction that books deliver to their users? Third, how does the arrival of a large new set of LLM-enabled books affect consumer welfare? Finally, how does the availability of AI affect releases by incumbent authors?

Answering these questions requires data on the supply of, and demand for, new books.¹ Using data on new ebooks offered for sale at Amazon, we document that the number of new titles appearing each month tripled between 2022 and late 2025 and rose by a factor of nearly ten in some categories. Over 90 percent of the release increase following ChatGPT can be explained by a surge in new author entry, and these new authors' releases account for nearly all of the AI-containing books. The post-2022 increase in releases – releases less their pre-AI average – tracks the detected AI share closely, indicating a constant stream of human releases and providing our first piece of evidence that AI-containing books are incremental.

We document the impact of LLMs on the book usage distribution in two broad ways. First, we show how the availability of AI affects the overall usage distribution of new books based on an approach that does not rely on AI detection. We assemble a stratified random sample of over 330,000 titles that are representative of all ebooks released at Amazon between 2020 and 2025. We observe the number of ratings for each book as of April 2026 – our main usage measure – as well as each book's April-2026 sales rank and average star rating. We find that the LLM-affected vintages, in more heavily treated categories, have lower average usage and lower usage at particular percentiles of the distribution. Yet, we also find an increase in our usage measure conditional on absolute rank (e.g., among the top 500 book releases in a category and month), indicating growth in the number of moderately appealing books released.

Our second approach uses AI detection on nearly 50,000 randomly selected titles, and we compare the appeal of books that do, and do not, contain AI. First, AI books are less

¹These data are not readily available. Neither copyright registrations nor repositories of international standard book numbers (ISBNs) show the patterns we document in this study.

appealing to consumers. Most have very little usage, while a modest share attracts some usage, both when measured by the number of ratings and eventual sales ranks. AI books also are of lower quality in the sense of having lower star ratings. These AI differentials in usage, sales, and star ratings remain when we control for book prices and length, as well as for a measure of the books’ promotion by Amazon. Second, not only the volume but also the appeal of non-AI releases remained similar over time. Hence, the decline in appeal documented in the main approach equals the share of AI books times the difference in appeal between AI and non-AI books. The AI influx thus explains the reduction in average appeal.

Given the stable output of non-AI books during the LLM era, we treat the AI-containing books as incremental for a consumer surplus (CS) calculation. Using a simple calibrated nested logit demand model for comparing the value of all releases to just the human-written releases, we find that AI books raise consumer surplus by less than one percent in 2023, by just over three percent in 2024, and by about seven percent in 2025 when over half of new books employ AI and better authors use it. Because the appeal of AI books is low, the increase in CS due to AI is small in comparison with the rapid increase in new releases. Still, it is positive because AI’s arrival did not displace human creation.

Our data also allow us to address questions related to author careers and the human-AI gap. Using a separate dataset, a census of all 479,650 books released in eight subcategories between 2008 and 2025, along with an AI detection subsample, we find that author selection – in particular, the substantial entry of first-time authors – explains most of the human-AI gap. Inexperienced and less successful authors – whose pre-LLM books had lower average usage – are more likely to adopt AI; and just over a third of the raw human-AI gap remains when we identify the gap using within-author variation. Finally, the entry of new authors did not displace activity by authors active prior to the LLM influx, whose releases rose after 2023. Non-AI releases by incumbents remained roughly constant, and the growth in incumbents’ output reflects their AI adoption.

The paper proceeds in seven sections after the introduction. Section 1 discusses the background, including author use of LLMs, the evolution of the book market, and relevant academic literatures. Section 2 provides a simple theoretical framework showing how the welfare effect of an influx of new products depends on both the number of new products and their quality distribution. Section 3 describes the data used in the study, including information on the volume of books released, as well as book-level data on books’ success measures and, for a subset, book-level information on detected AI use. Section 4 presents results on the effects of LLMs on new book entry and on the distribution of book usage and quality, based on the whole sample without AI detection. Section 5 uses the AI detection subsample to directly compare the appeal differentials (in usage, sales ranks, and star ratings) of AI and non-AI books, as well as AI’s effect on consumer welfare. Section 6 explores the determinants and effects of author AI adoption, the sources of the human-AI quality gap, and the possibility of incumbent author displacement. A brief conclusion follows in Section 7.

1 Background

1.1 AI and authors

The public deployment of ChatGPT in late 2022 brought LLMs to prominence. According to Liang et al. (2025), “LLM usage surged following the release of ChatGPT in November 2022.” In book publishing, surveys indicate that LLMs are both attracting new writers and being used by existing writers. An April 2025 survey of 1,229 authors – most of whom were exclusively self-published and had started publishing in 2023 or earlier – found authors divided in their views. Nearly half (45 percent) were using AI to “assist with their work,” while 48 percent were neither using nor planning to use AI. The remaining 7 percent indicated they might use AI in the future.² Among those not using AI, more than three quarters

²See <https://insights.bookbub.com/how-authors-are-thinking-about-ai-survey/>.

thought its use “unethical.” Among authors using AI, ChatGPT was the most popular LLM, used by 85 percent of users.³

A cottage industry now provides LLM advice to would-be first-time authors. For example, `bookautoai.com` offers support so that “you can publish a fully formatted, human-level book in under 30 minutes.” Some authors have apparently achieved success using AI. Tim Boucher has written hundreds of books using AI and reports having earned thousands of dollars.⁴ At the same time, there is a great deal of low-quality work, which critics derisively term “AI slop.”⁵

Participants in the book publishing industry regard AI with concern and suspicion. The Authors Guild argues that publishers and sellers should be compelled to disclose the use of AI in book creation.⁶ An open letter from authors, including Colleen Hoover and Dennis Lehane, exhorted publishers to “pledge that they will never release books that were created by machines” and not to “replace their human staff with AI tools or degrade their positions into AI monitors.”⁷

Intellectual property protection for AI-enabled writing is limited. The US Copyright Office has determined that “the outputs of generative AI can be protected by copyright only where a human author has determined sufficient expressive elements.”⁸ Amazon also appears

³In another survey from November 2025, 61 percent of writers – and 42 percent of fiction authors – reported using AI tools. See <https://www.publishersweekly.com/pw/by-topic/digital/copyright/article/99019-new-report-examines-writers-attitudes-toward-ai.html>.

⁴See <https://www.newsweek.com/ai-books-art-money-artificial-intelligence-1799923>.

⁵Goodreads has many user-created lists of works identified as “AI slop” by Goodreads users. See <https://www.goodreads.com/list/tag/ai-slop>, <https://www.goodreads.com/list/tag/ai-book>, and https://www.goodreads.com/list/show/219041.Books_I_Have_No_Doubt_Are_Written_By_AI_.

⁶The Authors Guild argues for rules requiring “authors, publishers, platforms, and marketplaces to label AI-generated works and otherwise identify when a significant portion of a written work has been generated by AI.” See <https://authorsguild.org/advocacy/artificial-intelligence/faq/>.

⁷The initial list, from June 2025, included 70 authors in total. By April 16, 2026, the list included 448 authors and other professionals in the book publishing industry. See <https://lithub.com/against-ai-a-n-open-letter-from-writers-to-publishers/> and <https://docs.google.com/document/d/1TmtwFvRAQkR35lg4xqBNiU3TUIYNWN1xNxx9hrM14/edit?tab=t.0>.

⁸The Copyright Office further clarifies: “This can include situations where a human-authored work is perceptible in an AI output, or a human makes creative arrangements or modifications of the output, but not the mere provision of prompts.” See <https://www.copyright.gov/newsnet/2025/1060.html> and

to be cautious in its embrace of AI. Amazon requires authors to disclose AI-generated, but not AI-assisted, content, but it does not pass that information on to consumers.⁹ Moreover, Amazon limits the number of books that authors may upload per day.¹⁰

1.2 The evolution of the book market

Prior to digitization, books – and the book market – were easy to define. Books were physical products, initially released as hardbound editions and later as paperbacks, for sale at retail establishments. In sales terms, the market was dominated by a handful of New York publishing houses. Digitization changed the market substantially. With Amazon’s creation of the Kindle environment in 2007, including the Kindle device and the Kindle Direct Publishing platform, many books were released primarily as electronic products. Moreover, books came to market without the curation of the traditional gatekeepers. As [Waldfogel and Reimers \(2015\)](#) document, self-published works became a significant part of the market in the years after the release of the Kindle; and many works initially appearing as self-published books became best sellers ([Peukert and Reimers, 2022](#)). Self-publishing, the main vehicle for LLM-aided books post-2022, was already an established avenue to commercial success.

Measuring the number of new books is not easy in the digital era. Traditional publishers sought copyright registration and international standard book numbers (ISBNs). Individuals publishing outside of traditional avenues often do neither. Between 2020 and 2024, annual nondramatic literary copyright registrations show no evidence of an LLM-induced influx.¹¹

<https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf>. See also [Lutes et al. \(2025\)](#).

⁹“We define AI-generated content as text, images, or translations created by an AI-based tool. If you used an AI-based tool to create the actual content (whether text, images, or translations), it is considered “AI-generated,” even if you applied substantial edits afterwards.” See https://kdp.amazon.com/en_US/help/topic/G200672390.

¹⁰See <https://www.theguardian.com/books/2023/sep/20/amazon-restricts-authors-from-self-publishing-more-than-three-books-a-day-after-ai-concerns>.

¹¹New registrations went from 152 thousand to 156 thousand over this period, albeit with fluctuation between 129 and 189 thousand. See, for example, <https://www.copyright.gov/reports/annual/2020/ar2020.pdf>.

The number of ISBNs issued per year, from Bowker’s Books in Print, also shows no evidence of an LLM-induced increase in releases. Hence, the influx we document below is largely a digital – and an Amazon-centered – phenomenon. Although it is a single company, Amazon has over two thirds of the global ebook market.¹²

1.3 Relevant literatures

This paper connects three strands of prior research: (i) the economics of digitization and creative-goods supply, (ii) empirical and theoretical work on generative AI in creative and knowledge-intensive tasks, and (iii) recent analyses of AI-driven creative markets.

A first strand examines how digitization has affected the production and quality of creative goods. [Brynjolfsson et al. \(2003\)](#) document the benefit of access to a long tail of extant products. Other work explores supply impacts. In the book market, [Waldfogel and Reimers \(2015\)](#) show that digital distribution and self-publishing substantially increased the number of book titles released, while also lowering entry barriers for new authors. Related work in music finds similar effects: digitization expanded variety and may have increased realized quality by bringing forth many products of unpredictable quality, some of which turn out to be good ([Aguiar and Waldfogel, 2018](#)). These studies show that technological change that reduces production and distribution costs can simultaneously raise both the volume and the realized quality of the best new creative outputs.

A second strand evaluates how generative AI affects cognitive and creative work. Experiments show large productivity gains from access to large language models. [Noy and Zhang \(2023\)](#) report that GPT-based writing assistance reduces task completion times by roughly 40 percent and disproportionately benefits lower-skilled workers. [Brynjolfsson et al. \(2025\)](#) find that an AI-based assistant in a large call center improves worker productivity by

¹²“Amazon holds a dominant 68% market share in global e-book sales, which rises to 83% when Kindle Unlimited is included.” See <https://goodereader.com/blog/electronic-readers/amazon-is-the-dominant-player-in-ebook-sales>.

14 percent on average, with the largest gains for novice workers. In creative tasks, recent work documents improvements in fluency and coherence but mixed effects on novelty and diversity (Doshi and Hauser, 2024; Padmakumar et al., 2024). On the other hand, while time saved in some tasks may be reallocated to higher-level activities, the need to verify can offset gains when error costs are high or outputs are difficult to evaluate (Dell’Acqua et al., 2023; Vaithilingam et al., 2022). Relatedly, Kusumegi et al. (2025) document that academics adopting LLMs produce substantially more working papers. Together, these studies raise the possibility that AI may significantly increase new book production, but they leave open the exact nature of effects on quality. AI need not raise the quality of incumbent creators’ work: Chen et al. (2025) conduct an experiment in which incumbents facing competition from AI sacrifice quality.

A third strand considers the effect of generative AI on creative markets, thus far with a focus on images. Goldberg and Lam (2025) show that generative AI tools increase entry and product variety while sharply displacing output of non-AI images in a large stock image marketplace. Peukert et al. (2024) find that stock image creators reduce their output when they learn that their work is used for AI training. Finally, Kim et al. (2026) document that the launch of text-to-image generative AI depresses output among incumbent creators of anime- and manga-style artwork.

2 Theory

Potential authors create a book if the expected benefit of doing so exceeds its cost. AI tools have reduced the cost of book creation, and cost reduction stimulates additional creation. Moreover, the appeal of creative products is understood to be unpredictable at the time of entry (Caves, 2000). With unpredictability, general cost reduction can raise welfare by delivering additional products from a quality distribution similar to that for existing products, as

the first wave of digitization did with media products (Aguilar and Waldfogel, 2018). Yet, the arrival of AI is not simply a reduction in the cost of entry. AI may allow new author entry, bringing more books of uncertain quality, and it may allow existing authors to produce more books than they otherwise might. Hence, LLMs can bring changes throughout the realized distribution of product sales. These concerns motivate the paper’s main question: Given AI’s effect, what happens to the distribution of entering books’ appeal and their benefit to consumers?

What matters for consumer welfare is the distribution of realized products’ appeal and, in particular, the number of products attracting substantial usage. A simple model of the realized distributions illustrates how AI tools might affect the value of the choice set. Define δ_j as the mean utility of product j .¹³ Define μ and σ as the mean and standard deviation of realized mean utilities in either pre- (human-only) or post-LLM (combined human and LLM) regimes. Define n as the number of products introduced. With $\{\delta_j\}$ as the set of available products, per capita CS in the plain logit is proportional to $\ln(1 + \sum_{j=1}^n e^{\delta_j})$. Below we also allow for more flexible substitutability among products, which changes the CS formula slightly.

Some cases are helpful for organizing thinking. First, suppose that the number of releases n rose while the distribution from which mean utilities are drawn in the AI era remained constant so that AI and human books were drawn from the same distribution. Then the average value, as well as the value at any particular percentile of the distribution, would remain the same. For example, if the number of entering products doubled from 1,000 to 2,000, then the 90th-percentile products – the 100th-best before and the 200th-best after – would be of equal mean utility. Conversely, the 100th-best product would be better when 2,000 products entered than when 1,000 products entered; and there would now be more

¹³In the plain logit model, $\delta_j = \ln(q_j/M) - \ln((M - \Sigma q_j)/M)$, where q_j is the quantity sold for book j and M is market size. We return to this approach in Section 5.3.

products above the former 100th-best book threshold. The increase in entry would, of course, raise the value of the choice set.

For a second case, suppose that μ declined while n and σ remained constant. This could happen, for example, if some human-authored book production were displaced by AI books drawn from a lower-mean quality distribution. Then – comparing the new vs the old cohort – the average mean utility would fall, as would the value at any percentile of the distribution. Moreover, the k^{th} -best new product’s mean utility would also fall. A choice set produced in this way, rather than with the prior distribution, would have lower value.

Finally, consider a third case, where n rises while μ declines. This could happen, for example, with a large number of AI books entering with low average appeal alongside the new books with purely human authorship.¹⁴ Then effects would be slightly more nuanced: The average mean utility would naturally decline, as would the value at any particular percentile of the distribution; but the mean utility of the k^{th} -best product might rise or fall, depending on the quantity of entering books as well as their quality dispersion.

We illustrate the third case in Figure 1. Panel (a) compares the pre-LLM distribution (in blue) with the larger post-LLM distribution (in red), where the post-LLM (human plus AI) distribution is illustrated with five times as many products and a lower mean.¹⁵ Panel (b) compares the mean utilities of products at the same *percentile* of the pre- and post-LLM distributions, and the value at any percentile of the distribution is lower in the post-LLM distribution. Panel (c) compares products at the same *rank position* (e.g., the 200th-best product) in the pre- and post-LLM distributions, and the product at the k^{th} -best position is better under the post-LLM distribution. An increase in the number of new products in,

¹⁴If human production continued steadily, then the addition of AI books could only raise surplus, although the increase could be small – even with a large AI influx – if all AI books were of low quality. If human production fell, then surplus would rise only if some of the new AI books were better than the forgone human books.

¹⁵We characterize the pre-LLM period with 200,000 draws from a standard normal distribution, and we represent the LLM era with 1 million observations from a normal distribution with the same variance and a mean of -0.5 .

or near, the right tail of the distribution is important, as it can substantially raise the value of the choice set to consumers.

The framework in this section helps to focus our attention on this paper’s main objects of study: the number of new AI and non-AI products, the average appeal of entering books, the mean utility at particular percentiles of the distribution, and the values at absolute rank positions.

3 Data

We are interested in characterizing volumes of new books entering before and after the influx of AI tools, as well as the distributions of those books’ appeal to consumers. Ideally, we would have a full list of books published over time, along with measures of those books’ sales in each time period. While ideal usage data are not available, we have two promising ways to measure the phenomena of interest. First, we observe the number of new books released by category and month, 2020-2025, as well as product-level information on cumulative usage at two points in early 2026. We use these data for a measurement approach in which we compare usage – and implied appeal – distributions across categories and over time. We employ this approach to estimate the effect of the presence of AI on the appeal of new books created. Second, we use AI detection on a random sample of these books to compare usage measures, sales ranks, and star ratings for books that do, and do not, contain detected AI.

3.1 Aggregate book release data

We obtain the number of books published over time using queries to Amazon’s advanced book search function at <https://www.amazon.com/advanced-search/books>. We request new, English-language ebooks, and we query the search tool by release month and category

for all 30 Amazon categories between 2020 and 2025.¹⁶ A query delivers a list whose first page indicates “1-16 of x results.” When there are fewer than 1,000 results, x is the exact number of releases in the category and release month. Up to 10,000, the number of search results is reported to the next-lowest thousand (e.g., “1-16 of over 2,000 results”). From 10,000 on, the result is reported to the next-lowest 10,000. The search result also includes the number of search-result pages, with a maximum of 400, indicating 400×16 , or 6,400 titles in the result. Hence, the number of books in a search is exact for 1-999; rounded up to the nearest 16 for searches returning 1,000-6,400 results; rounded down to the nearest 1,000 (or reported as 6,400) up to 10,000; and rounded down to the nearest 10,000 thereafter.¹⁷

Amazon lists some books in multiple categories, in which case the sum of the numbers of books across categories would overstate the total number of books released each month. We have two ways of estimating the extent of overstatement, a direct empirical approach based on a comparison of subcategory totals to category totals (for which the totals can be measured accurately) and a probabilistic approach based on the number of repeated titles across categories in the random sample we discuss below. Using the subcategory approach, the unduplicated count averages 82 percent of the observed, while the probabilistic approach indicates an average of 57 percent. Importantly, neither changes systematically over time. See Appendix A.1 for a discussion of these adjustments.

¹⁶The categories are arts & photography; biographies & memoirs; business & money; children’s books; comics & graphic novels; computers & technology; cookbooks, food & wine; crafts, hobbies & home; education & teaching; engineering & transportation; health, fitness & dieting; history; humor & entertainment; law; LGBTQ+ books; literature & fiction; medical books; mystery, thriller & suspense; parenting & relationships; politics & social sciences; reference; religion & spirituality; romance; science & math; science fiction & fantasy; self-help; sports & outdoors; teen & young adult; test preparation; and travel.

¹⁷Seven of the 30 categories (children’s books; health, fitness & dieting; literature & fiction; religion & spirituality; romance; science fiction & fantasy; and self-help) become sufficiently large during the sample period that the monthly category totals are reported only to the nearest 10,000. For those categories, we obtain their 148 subcategory totals by month to create the aggregate time series of new books released by category.

3.2 Main book-level analysis dataset

We use the search results pages to obtain book-level information. Although Amazon reports the number of pages in a search result up to 400, it is only possible to directly access the first and last 75 pages of search results and therefore only 2,400 ($= 150 \times 16$) books. Sorting by publication date, we sample books' Amazon Standard Identification Numbers (ASINs) from these pages, and we obtain updated information on the numbers of ratings, sales ranks, and average star ratings per book as of February 6 and April 5, 2026, from ASIN Data API (<https://app.asindataapi.com>).

Our main dataset is a random sample of books from all categories and publication months, 2020-2025. The aggregate series shows ten million title listings (including repeats) between the start of 2020 and the end of 2025. We assemble a list of all possible 16-title search pages using the information above on the number of titles released per category and month. We choose ten random pages per month and category.¹⁸ Our random sampling approach obtains 160 titles within each category and month, or a total of 331,360 observations. We use this stratified random sample, suitably weighted, to explore the effects of the AI influx on the distribution of books' appeal measures. The top panel of Table 1 summarizes the sample. Using the full stratified sample, in row 1, the average number of ratings is 83.38, and the average star rating (among the books with ratings) is 4.40. The average price is \$14.85, and the average publication year is 2022.5.

AI Detection subsample: We have a direct measure of AI use for a random sample of 48,193 books among those in the main dataset. This subsample includes 9,923 random titles from 2020-2022 and 38,270 from 2023-2025.¹⁹ We obtained 1,000-word text previews of

¹⁸For the largest seven categories, we choose ten random pages from their subcategories to avoid oversampling of books published near the beginning or end of the month.

¹⁹We sampled more heavily from 2023 to 2025.

books at Amazon and ran them through the Pangram Labs AI detector.²⁰ The AI detector classifies a text passage as fully human or fully AI, or returns an AI share. We classify books that are not fully human as AI.²¹ The quality of AI detection varies across tools, and Pangram’s tool has been shown to perform well.²² The false positive and false negative rates for prose evaluation at Pangram translate to error rates of at most a few percent for our AI/non-AI usage comparison. See Appendix A.2 for details.

Table 1 shows that books in the AI Detection subsample have an average of 59.65 ratings, an average star rating of 4.42, and an average price of \$13.90. The human-created books, which, on average, have had more time to accumulate ratings, have much higher usage, 82.53 vs 5.96 average ratings. The human-created books have higher star ratings as well, 4.46 vs 4.22. Finally, the AI books have lower average prices, \$7.94 vs \$16.46. The last column reports log sales ranks as of April 5, 2026, which average 13.82 in the AI Detection subsample. Log sales ranks are also worse – 14.09 vs 13.72 – for AI vs human-created books, even though AI books tend to be more recent.

Additional controls: Some of our analyses below explore the role of platform promotion in the relative success of AI and non-AI books. The advanced book search utility allows books to be sorted according to how much the book is “featured” by Amazon. The associated featured ranking depends on the book’s sales history and “whether [Amazon] think[s] the

²⁰The text previews are obtained from [https://read.amazon.com/sample/\[ASIN\]](https://read.amazon.com/sample/[ASIN]). Pangram Labs AI detection is available at <https://www.pangram.com/>. Our main sample uses the first 1,000 words of the previews; but Amazon text samples often include more words. We obtained up to nine additional 1,000-word blocks of text for 1,000 randomly chosen books in our sample. Within books, there is no statistically significant relationship between the block position and the tendency for the text to contain detected AI ($F = 0.63$, p -value = 0.77).

²¹The vast majority of text samples are classified as 0 or 100 percent AI. Just 4.94 percent of the classified books are partially AI; of the partial AI books, the median AI share is 0.65.

²²Jabarian and Imas (2025) provide a thorough comparison of AI detectors and conclude that “a clear ranking in detectors emerges. Pangram maintains near-perfect accuracy across long and medium length texts. It achieves very low error rates even on shorter passages and ‘stubs.’ These results are robust to exogenous changes in the threshold and the use of ‘humanizers’ such as StealthGPT.”

item will be of interest.”²³ With up to 16 results per page and a maximum of 75 pages, we can observe featured ranks from 1-1,200. We searched each category and month (from January 2023 through the end of 2025) in early July of 2026. We then merged these featured ranks with the AI Detection subsample, for a total of 22,759 matched observations. The distributions of featured ranks for AI and non-AI books have substantial overlap. The interquartile ranges run from 311 to 859 for AI books and from 241 to 799 for non-AI books.

3.3 Census of subcategories

For some questions, such as whether incumbent authors change their productivity in the face of AI, it is useful to have all of the titles released in a genre. Hence, we collect a complete “census” of new titles in eight subcategories that are sufficiently small to allow their full collection each month, from January 2008 to December 2025.²⁴ For each title, we observe the author, the publication date, the number of persons rating the book, and the average star rating, as of early January 2026. These data are described in the bottom panel of Table 1. The sample contains 479,650 books in total.

AI Detection author subsample: In order to study AI use over time within authors, we create an LLM-era subsample that contains all of the books (within these subcategories) by two subsets of authors. We search for text samples from all of the LLM-era works by 4,957 randomly selected authors active prior to 2023 and all of the works by 1,436 randomly chosen authors with no works before 2023.²⁵ The resulting AI Detection subsample contains 9,120 books released between 2023 and 2025, 7,269 by incumbents and 1,851 by authors debuting in 2023 or later. As in the main AI Detection sample, AI-containing works have

²³See <https://www.amazon.com/gp/help/customer/display.html?nodeId=GSUNWNFT2ALMPR3L>.

²⁴These subcategories are romantasy (a subcategory of romance), economics (business & money), fantasy (science fiction & fantasy), women sleuths (mystery), alternative history (literature & fiction), world history (history), US travel (travel), and sports biography (sports & outdoors).

²⁵We excluded authors with more than ten books total.

less usage, lower average star ratings, and lower prices. However, the share of AI books in this sample – which we oversample to include more books by incumbent authors – is lower than in the main AI Detection sample.

3.4 Usage measure validity

Our strategy for inferring appeal from our main usage measure has three potential shortcomings that we address in detail in the Appendix. First, we use information on the number of ratings that books have received rather than actual sales. We show in Appendix A.3 that the number of ratings is highly correlated with sales in a large sample of Amazon ebooks from Bookstat. We also undertake analyses using the April 2026 Amazon sales rank as a usage measure to circumvent the possibility that the propensity to leave a rating varies between human and AI books.

Second, the number of ratings that books have received depends not only on consumers’ level of interest but also on how long the books have been available. While most of our approaches rely on within-release-date comparisons and hence circumvent usage timing issues, some of our measurement approaches depend on comparability of the usage measure across release months. We develop an age adjustment using information on books’ cumulative numbers of ratings at two points in time, February 6 and April 5, 2026. These data show how the number of ratings received – and the probability that books have any ratings – grow over time for books of different ages. We use the resulting growth rates to calculate the adjusted number of ratings as of 72 months after release, which we term $\tilde{r}_j$. We use the February-to-April transitions to adjust zero ratings to the unconditional averages (the share changing to positive $\times$ the average ratings among now-positive), so that all adjusted ratings are positive. See Appendix A.4.

Finally, it is possible that the increase in the number of books divides rating-leaving users’ attention among titles, reducing the tendency for books of similar appeal to receive

ratings. This could make average appeal seem to fall even if it remained steady. We address this in two ways. First, we compare the usage measure, as well as sales ranks and star ratings, for books from particular release dates with and without detected AI. Same-release-date comparisons are unaffected by possible changes in consumption and rating behavior over time. Second, if dilution causes usage measures to understate appeal, the increases we document at absolute rank positions would understate actual increases.

4 Entry and usage distribution effects

This section explores two broad questions. Section 4.1 documents new book entry and the evolving AI-detected share. Section 4.2 uses our main dataset to explore how the distribution of books’ appeal, as measured by usage, evolves as the aggregate detected AI share rises from 2023 to 2025. We provide a complementary estimate and explore the mechanism for the change in book appeal with a direct comparison of AI and non-AI books in Section 5.

4.1 Aggregate book release patterns and AI

The solid line in Panel (a) of Figure 2 shows a monthly index of aggregate new titles across all Amazon categories between January 2020 and December 2025, normalized to 100 based on average monthly releases in the first ten months of 2022. The index is roughly stable between 2020 and late 2022. It rises during 2023 and 2024 and then rises more sharply in 2025, peaking at over 300 in late 2025.²⁶

The timing of the growth in releases corresponds to the growth in detected AI. Panel (b) of Figure 2 shows the weighted share of books in the AI Detection subsample detected to contain AI. Its growth coincides with the growth in releases in Panel (a). Using the AI

²⁶The number of distinct releases over the time period depends on the extent of titles listed in multiple categories. The sum of the category totals exceeds ten million. Using the mean (median) overlap calculations in Appendix A.1, the total is between 5.7 (6.6) and 8.2 (8.0) million.

detected share, we calculate the time pattern of human book production, which is the dashed line in Panel (a) of Figure 2. The resulting human book release index is flat, indicating that human book output continued at its pre-ChatGPT level during the AI era. The pattern also suggests that AI-containing books are incremental to what was produced prior to AI and would have continued to be produced without AI’s appearance.

We get additional evidence about the role of AI in book creation from book releases by author vintage, which the census of subcategories includes directly. The solid line in Panel (c) of Figure 2 shows overall releases in the subcategory census, and it mirrors the overall pattern in Panel (a). Using information on each author’s debut year, we calculate that the vast majority of the increase in releases between 2022 and 2025 comes from higher-than-normal releases by authors who debuted after 2022. The excess of those authors’ releases over typical n^{th} -year output accounts for 93 percent of the overall increase in Panel (c).²⁷ Removing these new-author excess releases from the total yields the dashed line in Panel (c) of Figure 2. This reinforces the implication from Panel (a) that if not for the influx of new, AI-using authors, releases would have been flat.

Panel (a) of Figure 2 shows that AI accounts for the whole increase in book releases, and Panel (c) shows that excess entry by post-ChatGPT entrants accounts for 93 percent of the increase. Hence, it is not surprising that the vast majority (94.8 percent) of AI-containing books are by authors who entered in the LLM era.²⁸ That the rise in releases is attributable to an increased number of newly entering authors – and that nearly all AI books are by the new authors – indicates that AI-containing works are incremental, over and above the works that humans continue to create.²⁹

²⁷Between 2015 and 2021, first-year entrants in the eight subcategories release an average of 17,183 works each year, while entrants in the second and third years release an average of 1,874 and 1,268 respectively. We treat the post-2021 entrants’ releases less these averages as excess releases.

²⁸Of all books in the census of subcategories released between 2023 and 2025, 82.1 percent are by newly entering authors, 40.4 percent of whose releases contain AI. The remaining 17.9 percent of releases are by pre-existing authors; 10.1 percent of these releases contain AI. Hence, 94.8 percent of AI books are by authors entering in the LLM era.

²⁹Section 6.2 shows that incumbent authors increase their output in the AI era by using AI, providing an

Returning to the full sample, the growth rate of new releases is large across most categories. The left panel of Figure 3 shows the ratios of the highest monthly totals during 2025 relative to 2022, by category. Categories with particularly large proportionate growth are mostly nonfiction and include travel (9.1 times higher), sports & outdoors (6.9), self-help (5.3), and computers & technology (4.5). The right panel of Figure 3 shows the category share of titles published between 2023 and 2025 in which AI is detected. The share is lowest in comics & graphic novels and highest in travel and computers & technology. Categories with fast growth tend to have higher AI shares for 2023-2025.

4.2 Effects of LLMs on the usage distribution

We use the data in our main sample to employ two approaches to measuring the impact of the LLM influx on the overall distribution of product appeal. These include a “high-low” approach comparing categories with more rapid and slower release growth, and a continuous approach estimating outcomes as functions of the number of new releases in each category and month.

High-low approach: Our first approach asks whether success measures evolve differently after the LLM-induced influx in categories with rapid release growth ($\mathbb{1}_j^h = 1$) relative to categories with slower growth. We divide the 30 book categories into two groups, those in the top quartile of weighted release growth and those less exposed.³⁰ We then ask how our usage measure changes across release dates for the highly exposed categories relative to the less exposed categories via regressions of the form:

$$\ln(\tilde{r}_j) = \mu_c + \mu_t + \lambda_t \mathbb{1}_j^h + \varepsilon_j, \quad (1)$$

explanation of the remainder of the 2022-2025 increase in releases.

³⁰We obtain similar results dividing at the weighted median.

where μ_c and μ_t are category and publication month fixed effects, and ε_j is an error term. We weight observations in a category and publication month to reflect the total number of books, and we cluster standard errors by category $\times$ release month. The coefficients λ_t show how the outcome measures for books in the highly exposed groups evolve over time, relative to the time pattern for the less exposed groups. Panel (a) of Figure 4 reports results on the average usage. Between 2020 and early 2022, the difference in average usage measures is stable, and then it falls noticeably after 2022. The overall post-LLM differential, from a regression on a post-2022 indicator interacted with $\mathbb{1}_j^h$, yields a coefficient of -0.422 (se = 0.053).³¹

We measure effects on usage at particular percentiles of the distribution by augmenting Equation (1) to include fixed effects for 20 within-category-and-month rank percentile quantiles. Panel (b) of Figure 4 shows the within-percentile results. Like the overall average, the usage conditional on rank percentiles is stable prior to late 2022, then falls.

To measure effects on usage at particular rank positions, we augment Equation (1) to include fixed effects for within-category-and-month rank ranges, 1-100, 101-200, etc. Panel (c) shows the resulting time fixed effects. In contrast to the usage at the mean and at percentiles, the usage at particular rank positions rises. A decline in average usage, and usage at percentiles of the distribution, may reflect a decline in users’ propensity to rate individual books driven by the larger volume of releases. The increase in measured usage conditional on rank positions reflects an increase that overcomes the possible dilution and therefore indicates an actual increase. Together, these results indicate a shift in distributions – more products with lower average appeal – as in the third theoretical scenario above.³²

³¹The fast-growing categories grow by 160 percent, while the low-growth categories grow by 70 percent. Hence, the coefficient shows the effect of a 90 percent differential. Because the average category release growth rate is 90 percent, this also provides an estimate of the overall LLM-era effect.

³²We also estimate a model with a purely temporal approach, comparing the period before and after the advent of LLMs with and without fixed effects for percentiles and rank positions. Results are very similar to the high-low results. See Appendix Figure A.3.

Continuous approach: We also ask how the usage measure evolves with the number of releases by category using a regression of log ratings on release month and category fixed effects as well as the log of the number of releases in that category and month:

$$\ln(\tilde{r}_j) = \mu_c + \mu_t + \phi \ln(n_{ct}) + \varepsilon_j, \quad (2)$$

where n_{ct} is the number of books entering in category c in month t and μ_c and μ_t are category and publication-month fixed effects. The regressions are weighted as above, with standard errors clustered at the category $\times$ release month level; and we employ both the adjusted and raw usage measures. We also add rank percentile and rank range fixed effects to examine effects at different points in the distribution.

Table 2 shows the results, first – in columns (1) and (2) – for average usage. Using both adjusted and unadjusted usage measures, categories with faster growth in entry experience larger decreases in average usage. The raw and adjusted samples differ in that books with zero raw ratings are excluded from the raw ratings regressions but not the adjusted ratings regressions.³³ Columns (3) and (4) add rank percentile range dummies to Equation (2) to explore impacts within percentile ranges. The negative coefficients for both usage measures indicate that categories with more growth in releases experience larger reductions in usage, conditional on percentile. Columns (5) and (6) explore impacts within rank position ranges. Using both usage measures, usage rises more – within rank ranges – in categories with faster growth in entry.

While we have constrained ϕ to be constant across rank positions in Table 2, we relax this in Figure 5. The figure reports the coefficients on interactions of $\ln(n_{ct})$ with rank-position ranges from a regression of $\ln(\tilde{r}_j)$ that includes month, category, and rank-position fixed effects. The coefficient on $\ln(n_{ct})$ is not statistically distinguishable from zero for books ranked 1-200. The coefficient then grows as ranks worsen, reaching 1.2 to 1.35 for ranks

³³We obtain virtually identical results using both measures estimated on the raw subsample.

beyond 600. Thus, the large, LLM-aided growth in the number of books does not raise the usage of the top 200 books released per category-month; but it does raise the usage of books outside of the top 200.

The results above paint a consistent picture of the relationship between LLM-induced entry and the distribution of appeal, and the framework in Section 2 is useful for guiding interpretation of the results. The number of books released rises substantially, and despite a decline in average appeal, the increase in our book usage measure at particular rank positions (for all but the top 200 titles released per category and month) indicates a potential for modest welfare gains.

5 AI appeal differentials and welfare effects

This section uses the AI Detection subsample of the main dataset for a different approach to measuring the impact of AI on book appeal. Section 5.1 measures the usage, sales, and quality differentials between LLM- and human-generated books; and we show how the usage gap relates to the overall effect of AI that we documented in Section 4. Section 5.2 explores alternative explanations of AI books’ low usage measures. Section 5.3 presents estimates of the welfare impact of the LLM book influx.

5.1 AI and appeal differentials

If the arrival of AI delivered additional books while leaving the production of purely human books as it was, then the effect of AI on average book appeal could be ascertained by arithmetic rather than econometrics. Let the usage measures for AI and non-AI books be $\tilde{r}_j^{AI}$ and $\tilde{r}_j^{\text{non-AI}}$, respectively, and let π^{AI} denote the AI share. Then $\pi^{AI}\tilde{r}_j^{AI} + (1 - \pi^{AI})\tilde{r}_j^{\text{non-AI}}$ is the average usage measure under AI. If, additionally, the mean of $\tilde{r}_j^{\text{non-AI}}$ were both the pre-AI and LLM-era non-AI average, then the effect of AI on the overall distribution would

be simply the human vs AI appeal gap times the AI share, or $\pi^{AI} \times (\tilde{r}_j^{\text{non-AI}} - \tilde{r}_j^{AI})$.

We have already seen, in Figure 2, that the number of human books is roughly constant in the LLM era; but what happened to their average usage? The solid line in Panel (a) of Figure 6 shows the average usage of human books 2020-2025 (using all books from our main sample before 2023 and the non-AI books from the AI Detection subsample for 2023-2025), and it is roughly constant over time. We see a similar constancy for star ratings in Panel (b). Hence, we can use the arithmetic approach, along with an estimate of the AI differential, to provide an additional estimate of the effect of AI on average appeal.

We measure the difference in appeal between AI- and human-generated books in three ways, with the adjusted log number of ratings (our preferred measure), with a sales-rank measure, and with average star ratings. The first two columns of Table 3 present weighted regressions of adjusted ratings on the book-level AI indicator.³⁴ Column (1) includes no fixed effects and shows that the human-AI differential is -1.147 ($se = 0.053$). Accounting for category and month fixed effects, in column (2), the gap remains large, at -0.945 (0.042). The books created with AI during 2023-2025 attract substantially less usage than their human-generated counterparts.

The AI share times the human-AI gap estimated here ($39.6 \text{ percent} \times -0.945 = -0.374$) is close to the temporal decline in average usage from the high-low approach estimated earlier (-0.422). The difference is not statistically significant, with a p -value of 0.198 . Hence, the human-AI gap provides a rather direct explanation for the change in the book appeal distribution after the adoption of LLMs.

Robustness to alternative usage measures: We have interpreted the AI gap as a reflection of AI books’ lower appeal, but it is possible that the tendency to leave ratings differs between AI and non-AI books, which would challenge our basic usage measure. We

³⁴The weights here inflate the AI Detection subsample to the population.

address this with an alternative usage measure that does not depend on whether users leave ratings, the sales rank as of April 5, 2026. Despite its advantage, this measure has the shortcoming that it is the sales rank at a date with differing delays after publication, and ranks weight recent sales more heavily. The measure has the further characteristic that books without recent sales have an undefined sales rank, even if they have sold substantially in the past and have large numbers of ratings.

We regress the log sales rank for book j on an indicator for whether it contains AI, along with category and publication month fixed effects. We estimate this by least squares using only the observations with an observed sales rank, and we weight and cluster as above. We also estimate a censored normal model in which books without observed sales ranks are assumed to have a rank worse (higher) than the worst observed sales rank. Columns (3) to (5) of Table 3 report the coefficients from these regressions. Column (3) reports a specification without fixed effects, yielding a coefficient of 0.321 (se = 0.037), while column (4) adds the fixed effects, giving a coefficient of 0.303 (se = 0.040). The coefficient in the censored model (column (5)) is 0.773 (se = 0.047). Hence, this approach corroborates the main approach in showing that books containing AI are substantially less appealing to consumers, and the difference is statistically significant.

Columns (6) and (7) of Table 3 report regressions of books' average Amazon star ratings on the AI indicator, with and without category and month fixed effects. The AI coefficients are negative and significant, indicating that AI books are also worse in the sense that users rate them lower.

5.2 Alternative mechanisms

AI books may receive little usage because of their prices or lengths, or because they are not promoted, rather than because they are unappealing. We explore this, in Table 4, by augmenting the regressions in Table 3 with the books' prices, lengths (log pages), and – in

even-numbered columns – a measure of Amazon promotion as well, for the subset of books from our AI Detection subsample that we could match to the featured ranks.

Columns (1) and (2) use the adjusted log ratings as the dependent variable. As column (1) shows, books with higher prices have less usage, and longer books receive more usage. The AI indicator continues to have a negative and significant coefficient. Column (2) adds the books’ Amazon featured rank, and its coefficient has the expected sign: books with higher (less favorable) featured ranks have lower usage. The AI coefficient remains negative and significant, indicating that AI books’ lower usage does not simply reflect lower promotion. The remainder of the table repeats the exercises of columns (1) and (2) using the censored log sales ranks and the average Amazon star rating as dependent variables. While promotion and book characteristics sometimes matter in these regressions, our conclusion that AI books receive less usage, sell less, and are less appreciated by consumers remains when accounting for these other potential explanations.

5.3 Consumer welfare effects

Because the quality and quantity of human releases remained constant over time, the effect of AI on the choice set can be described by the addition of the AI books. This facilitates a simple calculation of welfare effects. We start by looking at where in the quality distribution these new books appear. To this end, Figure 7 presents histograms showing the number of AI and non-AI releases, by year, in each adjusted ratings interval.³⁵ The bar at each interval shows the number of releases at that usage level; and the dark components of the bars – starting in 2023 – show the AI books. In 2023, AI books are concentrated in the lower intervals and absent from the highest intervals. The dark areas become larger in 2024 but remain concentrated at the lower intervals. In 2025, a larger number of AI books appear

³⁵We exclude books whose ln adjusted ratings are below 2, as these books have little effect on total usage or welfare. Changes in CS calculated excluding these books are nearly identical to welfare effects reported below.

in the higher intervals. AI books add to the choice set in all years; by 2025 the AI influx delivers books with higher usage than in prior years.

The area of the light portions of the bars in Figure 7 is similar across years, providing additional evidence that the volume and usage distribution for human-authored books remain similar as LLMs appear and find wide use. Hence, the growing dark portions of the bars in Figure 7 foreshadow the effects of AI on welfare, and the AI Detection sample allows a simple calculation of the welfare effects of AI. For each year from 2023 to 2025 we compare the CS from the entire choice set – including both human and AI books – to the CS from the subset that would arise if only the human books from that year had been available.

We measure the magnitude of the change in CS using a nested logit calibration in the spirit of Berry (1994). Consumer i chooses among books $j \in J$ and the outside good according to utility given by

$$u_{ij} = \delta_j + \zeta_{ig} + (1 - \sigma)\epsilon_{ij},$$

where ϵ_{ij} follows an extreme value distribution, and the idiosyncratic inside-good preference ζ_{ig} follows a distribution such that $\zeta_{ig} + (1 - \sigma)\epsilon_{ij}$ also follows an extreme value distribution. The outside good has utility zero, or $u_{i0} = 0$.

The mean utility δ_j of product j is given by the observed market shares and a calibrated taste correlation parameter σ with $\delta_j = \ln(s_j) - \ln(s_0) - \sigma \ln(s_j/(1-s_0))$, where $s_j = q_j/M$ and s_0 is the market share of the outside good. The resulting consumer surplus CS associated with a choice set is proportional to

$$\ln \left(1 + \left(\sum_{j \in J} e^{\delta_j/(1-\sigma)} \right)^{1-\sigma} \right).$$

We treat our adjusted usage measures $\tilde{r}_j$ as quantities, and we define the market size M for

each year as the US population times ten.³⁶ We use the σ estimate from [Reimers and Wald-fogel \(2021\)](#) of 0.443 for the baseline nested logit implementation, and we explore robustness to a wide range of σ values. We focus on the percent change in CS across environments.

Table 5 shows the ensuing welfare calculations. The column headed “ $\sigma = 0.443$ ” reports baseline estimates. Using the 2023 releases in our AI Detection subsample weighted to the population of releases, all releases (including those with AI) deliver 0.44 percent more CS than just the non-AI releases. For 2024, the AI books raise CS by 3.26 percent. Finally, the AI books in 2025 – when total releases reach 240 percent of their 2022 level – also raise CS by only 7.23 percent. The last two columns of Table 5 report welfare estimates for alternative substitution parameters. If $\sigma = 0.25$, then the increase in CS ranges from 0.6 percent in 2023 to 9.8 percent in 2025. If $\sigma = 0.75$, the analogous percentages are 0.2 and 3.2, respectively.³⁷ We conclude that the influx of AI-assisted books creates positive consumer benefits, although these benefits are modest relative to the increase in the number of books.

6 Author AI adoption, quality, and displacement

This section explores two questions related to AI adoption and possible incumbent author displacement. Section 6.1 discusses how author selection and the appeal of AI-supported writing determine the gap between AI and non-AI book appeal. Section 6.2 explores whether the influx of AI displaces incumbent authors. Because this section studies author careers, we use the subcategory census dataset, which allows us to observe all books by certain authors. We also use the AI Detection author subsample, which includes books by authors who switch from human to AI-assisted authorship.

³⁶Our choice of the market size M does not drive our results. The estimated CS effects from adding AI books are nearly identical for ranges of M between $0.1 \times$ population and $100 \times$ population.

³⁷AI books tend to have lower prices than non-AI books, so the measured CS benefit reflects not only more books of modest quality but more books of modest quality that also have low prices.

6.1 Why are AI books less appealing?

The appeal of AI and non-AI books might differ for two reasons: the selection of authors using AI and the effect of AI on the writing in books whose authors have adopted it. First, nearly all (94.8 percent) of the AI books are from newly entering authors, so the quality of AI books is highly dependent on whether the new entrants, when using AI, produce appealing books. The quality of the remaining 5 percent of AI-containing books depends on both the quality of those authors and the effect of AI use on the quality of their books.

To decompose the AI usage differential, we estimate variants of

$$\ln(\tilde{r}_{ij}) = \lambda_t \mathbb{1}_{ij}^{AI} + \mu_i + \mu_t + \mu_{t-v} + \varepsilon_{ij} \quad (3)$$

on the AI Detection author sample, where the dependent variable is the log of adjusted ratings for book j by author i , $\mathbb{1}_{ij}^{AI}$ is an indicator for whether book j by author i contains AI, μ_i is an author fixed effect, μ_t is a publication calendar-month fixed effect, μ_{t-v} is an author age dummy (calendar time t less v , the month in which the author's first book appeared), and ε_{ij} is an error term.

Column (1) of Table 6 reports a variant of Equation (3) that excludes the author fixed effects and includes a single AI differential for the entire period. As in the main sample, AI books in this sample have lower usage (-1.74 , $se = 0.08$). Column (2) reports year-specific AI differentials, again without an author fixed effect, and the AI usage differential declines in absolute value from -2.20 in 2023 to -1.50 in 2025.

Column (3) adds author and author age fixed effects, so that the AI coefficients show the within-author differentials for authors with more than one sample book; and this sample overwhelmingly consists of incumbent authors. Even among these authors, the differential remains significant: When a given author uses AI, the resulting book is less appealing. However, the coefficients shrink substantially in column (3) relative to column (2): The

within-author component of the differentials is between 35 and 40 percent of total yearly differentials.³⁸ Hence, the majority of the AI usage differential stems from selection: Low-quality – and almost exclusively newly-entering – authors use AI.

The decline in the column (3) AI coefficients over time shows that the within-author AI penalty falls, which could reflect author learning or improving technology. Column (4) explores this by adding a rolling count of authors’ prior AI-containing books. The count variable has a positive but insignificant coefficient. The generally declining within-author AI differentials in columns (3) and (4) therefore appear to reflect improvement in the technology rather than author learning.

6.2 Has AI displaced incumbent authors?

The appearance of a large number of new books by entering – and in many cases AI-using – authors might displace continued activity by incumbent authors, those who were active prior to LLMs. We explore the productivity of pre-LLM authors with a cohort approach using the full census of subcategories and accounting for author activity patterns relative to the time of their first release. Define n_{tv} as the number of books released in time t by authors debuting at time (vintage) v . Productivity of a cohort evolves over time as it ages. We estimate the productivity of pre-LLM authors over time from these age effects. In particular, we estimate the following model on author cohorts entering between 2010 and 2022:

$$\ln(n_{tv}) = \omega_t + \mu_{t-v} + \varepsilon_{tv}, \quad (4)$$

where ω_t and μ_{t-v} are time and age fixed effects, and ε_{tv} is an idiosyncratic error. We use robust standard errors.

Panel (a) of Figure 8 shows the pattern of the time coefficients ω_t . They are stable

³⁸For example, in year 2023, the ratio between columns (2) and (3) is $(e^{-0.438}-1)/(e^{-2.197}-1) = 0.399$. These ratios remain similar if we estimate the column (2) regression on the column (3) sample.

between 2020 and the middle of 2023. The coefficients then rise gradually beyond pre-LLM levels starting in 2024. Conditional on incumbent author age, output rises by roughly ten percent in 2024 and 25 percent in 2025, relative to 2022. This shows that incumbent authors also increase their output during the LLM era. As documented above, incumbent authors account for only about 5 percent of AI books. Still, these AI-containing works account for nearly all of the growth in incumbent output in Figure 8, providing additional evidence that the non-AI component of incumbent output is stable over time.³⁹ These results, which are robust to the inclusion of different pre-LLM time spans, contrast with existing literature that finds substantial crowding out of human-generated creative work by AI in image markets (Goldberg and Lam, 2025; Kim et al., 2026; Peukert et al., 2024).

Finally, the growth in incumbent output can be decomposed into extensive (number of authors, au_{tv}) and intensive (number of books per author, np_{tv}) parts. That is, $n_{tv} = au_{tv} \times np_{tv}$. We run analogues of Equation (4) using these component variables as dependent variables and report them in Panels (b) and (c) of Figure 8. Both components rise during the AI era, and roughly three quarters of the increase occurs on the number-of-participating-authors margin. AI appears to encourage authors who otherwise might have exited the market to remain active.

7 Conclusion

LLMs, which diffused rapidly between 2022 and 2025, have become increasingly capable of writing long passages of text. Accordingly, they have seen rapid adoption by book authors. We document a tripling in the number of new books coming to market between late 2022 and late 2025, while fully human output remains roughly constant. The influx is explained

³⁹AI-containing work in the AI Detection author sample accounts for 8.5 percent of the books by incumbents in 2024 and 20.4 percent in 2025. Multiplying the human share (1 minus the AI share) by the total level increase in Panel (a) of Figure 8, incumbent non-AI book output rises by 0.7 percent in 2024 (0.915×1.10) and falls by 0.5 percent in 2025 (0.796×1.25), relative to 2022.

almost entirely by a surge in releases by post-ChatGPT authors, who also account for nearly all of the AI-containing releases.

The average usage and star ratings of new books have fallen with the LLM-induced influx while the quantity and appeal of new human releases have been roughly constant. Because human output has been stable, the effect of AI is also reflected, arithmetically, by the AI share times the difference in appeal. The effect of the AI influx on consumer welfare depends on the appeal of the additional books. The additional books, containing detected AI, are substantially less popular than human-authored books, so that much of the new work is of little value to consumers. Still, the LLM influx has delivered some books in the middle range of the usage distribution, and the book releases of the LLM era have delivered modest increases in consumer surplus.

Much of the human-AI quality gap stems from new – and low-quality – authors’ adoption of AI, but the human-AI gap shrinks between 2023 and 2025. Moreover, the arrival of LLMs does not appear to have displaced activity by incumbent authors, in part because AI adoption appears to increase incumbent authors’ productivity as well. Despite the controversy surrounding LLMs, the experience of book consumers in their wake – as with other cost-reducing technological changes in the cultural industries – seems to have been positive. Yet, because the new books mostly have low appeal, the effects are modest. AI does not appear to have the substantial positive welfare effect of the earlier round of digitization documented elsewhere for books and other cultural products.

While AI facilitates the creation of content that is somewhat valuable to consumers, there is a separate question of whether AI fosters the development of content that would be well-regarded by critics or other cultural elites. Commercial success and cultural importance are not the same thing, and the absence of an effect at the top suggests a limit to how much LLMs can help with elite cultural production. The answers to these questions may change, depending on hard-to-predict developments in the capabilities of LLMs.

References


- AGUIAR, L. AND J. WALDFOGEL (2018): “Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music,” *Journal of Political Economy*, 126, 492–524.
- AIGNER, D. J. (1973): “Regression with a Binary Independent Variable Subject to Errors of Observation,” *Journal of Econometrics*, 1, 49–59.
- BERRY, S. T. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *The RAND Journal of Economics*, 25, 242–262.
- BRYNJOLFSSON, E., Y. HU, AND M. D. SMITH (2003): “Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers,” *Management Science*, 49, 1580–1596.
- BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025): “Generative AI at Work,” *The Quarterly Journal of Economics*, 140, 889–942.
- CAVES, R. E. (2000): *Creative Industries: Contracts Between Art and Commerce*, Harvard University Press.
- CHEN, Y. J., J. GONG, J. LI, AND Z. ZHAO (2025): “Better Technology, Worse Motivation: GenAI’s Mediocrity Trap,” *Available at SSRN 5208163*.
- DELL’ACQUA, F., S. RAJENDRAN, E. I. MCFOWLAND, E. MOLICK, H. LIFSHITZ-ASSAF, L. KRAYER, F. CANDELON, K. R. LAKHANI, AND K. C. KELLOGG (2023): “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” Tech. Rep. 24-013, Harvard Business School Working Paper.
- DOSHI, A. R. AND O. P. HAUSER (2024): “Generative AI enhances individual creativity but reduces the collective diversity of novel content,” *Science Advances*, 10, eadn5290.
- GOLDBERG, S. AND H. T. LAM (2025): “Generative AI in Equilibrium: Evidence from a Creative Goods Marketplace,” *Working Paper*, available at SSRN.
- JABARIAN, B. AND A. IMAS (2025): “Artificial Writing and Automated Detection,” Tech. rep., National Bureau of Economic Research.
- KIM, S., G. Z. JIN, AND E. LEE (2026): “Does Generative AI Crowd Out Human Creators? Evidence from Pixiv,” Tech. rep., National Bureau of Economic Research.
- KUSUMEGI, K., X. YANG, P. GINSPIRG, M. DE VAAN, T. STUART, AND Y. YIN (2025): “Scientific Production in the Era of Large Language Models,” *Science*, 390, 1240–1243.

- LIANG, W., Y. ZHANG, M. CODREANU, J. WANG, H. CAO, AND J. ZOU (2025): “The Widespread Adoption of Large Language Model-Assisted Writing Across Society,” *arXiv preprint arXiv:2502.09747*.
- LUTES, B., J. GANS, S. GREENSTEIN, A. B. JAFFE, A. NAGARAJ, I. REIMERS, M. D. SMITH, R. TELANG, C. E. TUCKER, AND J. WALDFOGEL (2025): “Identifying the Economic Implications of Artificial Intelligence for Copyright Policy,” *First published by the US Copyright Office*.
- NOY, S. AND W. ZHANG (2023): “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science*, 381, 187–192.
- PADMAKUMAR, V., B. DHINGRA, W. AMMAR, F. METZE, D. DOWNEY, H. HAJISHIRZI, E. HOVY, AND N. A. SMITH (2024): “Does Writing with Language Models Reduce Content Diversity?” in *International Conference on Learning Representations (ICLR)*.
- PEUKERT, C., F. ABEILLON, J. HAESE, F. KAISER, AND A. STAUB (2024): “AI and the Dynamic Supply of Training Data,” *arXiv preprint arXiv:2404.18445*.
- PEUKERT, C. AND I. REIMERS (2022): “Digitization, Prediction, and Market Efficiency: Evidence from Book Publishing Deals,” *Management Science*, 68, 6907–6924.
- REIMERS, I. AND J. WALDFOGEL (2021): “Digitization and Pre-Purchase Information: The Causal and Welfare Impacts of Reviews and Crowd Ratings,” *American Economic Review*, 111, 1944–1971.
- VAITHILINGAM, P., T. ZHANG, AND E. L. GLASSMAN (2022): “Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models,” *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.
- WALDFOGEL, J. AND I. REIMERS (2015): “Storming the Gatekeepers: Digital Disintermediation in the Market for Books,” *Information Economics and Policy*, 31, 47–58.


Figures and tables

Figure 1: LLM influx and effects on rank positions vs percentiles


(a) Full distributions



(b) Value at 95th percentile




(c) Value at 200th rank position



Notes: This figure compares hypothetical pre-LLM (blue) and LLM-era (red) quality distributions, where the LLM-era distribution has a lower mean and more products. Panel (a) shows the overall distributions. Using these distributions, Panel (b) shows that the 95th (top 5) percentile product has lower quality in the LLM era, while Panel (c) shows that a product at the 200th rank position has higher quality in the LLM era.


Figure 2: Book entry and detected AI usage




Notes: Panel (a) shows an index of aggregate book releases across all categories (average of January-October 2022=100), along with an index of the number of human releases implied by the detected AI share. May 2024 is excluded because search queries deliver oddly large results in many categories for that month. Panel (b) shows the weighted share of books with AI detected, based on 9,923 titles for 2020-2022 and 38,270 for 2023-2025. Panel (c) shows the number of monthly releases in the census of subcategories. The solid black line shows the total number of releases by all authors. The dashed gray line shows the total releases, less the excess releases of the post-ChatGPT-entering authors (above the usual numbers of releases from 1st- or 2nd-year authors).

Figure 3: Entry and AI use by category

(a) Growth by category




(b) AI share by category




Notes: The left panel shows the ratios of the maximum monthly new publications in 2025 relative to the maximum monthly publications in 2022, for each category. The right panel shows the weighted AI share of the books in the AI Detection subsample, 2023-2025.

Figure 4: Usage distributions across publication dates


(a) Overall average



(b) Conditional on rank percentiles




(c) Conditional on rank positions




Notes: The figures show coefficients of the interactions between month dummies and an indicator for high-growth categories from the high-low approach in Equation (1). Panel (a) is the overall average. Panel (b) includes fixed effects for 20 rank-percentile quantiles. Panel (c) includes fixed effects for rank position ranges within categories by 100, up to 2,000.

Figure 5: Coefficients on $\ln(n)$ by rank position range




Notes: This figure shows rank-position-specific coefficients from regressions of our adjusted usage measure on the log of monthly releases per category. Standard errors are clustered on category $\times$ release month. The regression includes month, category, and rank group fixed effects.

Figure 6: AI and non-AI usage and quality time series



Notes: This figure shows weighted time series of usage ($\ln$ adjusted number of ratings), in Panel (a), and average star ratings, in Panel (b), for books of different types. The solid lines show time series for human-authored books (all books released before 2023 and those detected not to contain AI in 2023 and later). For 2023-2025, we also show the average usage and star ratings for AI books (blue, dotted lines) and for all (AI and non-AI) books (gray, dashed lines).


Figure 7: Annual AI and non-AI usage distributions




Notes: The figures show distributions of AI and non-AI books by log adjusted rating bins and release years. We report the distribution for $\ln(\tilde{r}) \geq 2$ and group all books with log adjusted ratings greater than 9.

Figure 8: Productivity of incumbent authors


(a) Total works



(b) Number of authors



(c) Works per author



Notes: This figure shows publication month fixed effects from regressions of author vintage productivity measures on publication month and author age fixed effects. We use the census of subcategories, and we include authors with debuts between 2010 and 2022. The regressions use robust standard errors. Panel (a) uses the log of overall author output (n_{tv}) as its dependent variable, Panel (b) uses the log of the number of authors, and Panel (c) uses the log of works per author.

Table 1: Summary statistics across samples

	# books	# ratings	avg stars	price	pub year	ln rank
<u>Main dataset:</u>						
Full sample	331,360	83.38	4.40	14.85	2022.52	
AI Detection sample	48,193	59.65	4.42	13.90	2023.38	13.82
human-created	33,792	82.53	4.46	16.46	2023.02	13.72
AI-containing	14,401	5.96	4.22	7.94	2024.21	14.09
<u>Census of subcategories:</u>						
Full sample	479,650	166.58	4.33	14.60	2019.47	
AI Detection sample	9,120	178.65	4.44	11.23	2024.04	13.59
human-created	7,614	208.63	4.45	12.01	2023.95	13.57
AI-containing	1,506	27.14	4.25	7.31	2024.47	13.77

Notes: The table shows average values for variables of interest in the main datasets. The main dataset, described in the top panel, is a stratified random sample of books released between 2020 and 2025. The top row shows raw averages for the full sample. We run 1,000-word samples from 48,193 of these books through Pangram Labs AI detection to create the AI Detection sample. The census of subcategories, in the bottom panel, contains all 479,650 books released in eight small subcategories between 2008 and 2025. We run 9,120 of these books through AI detection.

Table 2: Continuous approach usage regressions

	average		percentile		rank position	
	(1)	(2)	(3)	(4)	(5)	(6)
	adjusted	raw	adjusted	raw	adjusted	raw
ln monthly releases	-0.621*** (0.0671)	-0.459*** (0.0609)	-0.614*** (0.0677)	-1.188*** (0.0936)	1.580*** (0.0726)	0.672*** (0.0921)
N	326,975	157,373	326,975	157,372	326,954	157,361

Notes: This table shows weighted regressions of adjusted and raw ln ratings measures on the natural log of the number of monthly releases in the category, using the main dataset. All regressions include category and month fixed effects. The percentile regressions also include fixed effects for percentiles of adjusted usage within category-month (0-5, 6-10, etc.). The rank position regressions also include fixed effects for adjusted usage rank positions 1-100, 101-200, etc., within category-month. Standard errors are clustered at the category $\times$ month level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: AI book usage and quality differentials

	log adjusted ratings		log sales rank			avg. star rating	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
AI-containing	-1.147*** (0.0529)	-0.945*** (0.0422)	0.321*** (0.0368)	0.303*** (0.0401)	0.773*** (0.0470)	-0.146*** (0.0254)	-0.200*** (0.0244)
Fixed effects	None	cat. + mo	None	cat. + mo	cat. + mo	None	cat. + mo
N	38,270	38,270	23,628	23,628	38,270	15,277	15,277

Notes: This table shows regressions of usage measures on an indicator for whether the book contains AI. Columns (1) and (2) use $\ln(\text{adjusted ratings})$ as the dependent variable. Column (2) includes category and month fixed effects. Columns (3)-(5) use $\ln(\text{sales ranks})$ as the dependent variable. Columns (3) and (4) include only observations with enough recent sales to have ranks. Column (5) is estimated as a censored normal regression, and books without ranks are assumed to have worse ranks than the worst observed rank. Columns (6) and (7) use average star ratings as the dependent variable, for books that have received at least one rating. All models are estimated on the AI Detection subsample of the main dataset (2023-2025). Standard errors are clustered at the category $\times$ month level, and regressions are weighted to reflect the numbers of entering books by category and release month. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Product characteristics, featured ranks, and AI differentials

	log adjusted ratings		log sales rank		avg. star rating	
	(1)	(2)	(3)	(4)	(5)	(6)
AI-containing	-0.884*** (0.0762)	-0.688*** (0.0666)	0.298*** (0.0554)	0.150*** (0.0518)	-0.134*** (0.0255)	-0.135*** (0.0255)
price	-0.0115*** (0.000781)	-0.00641*** (0.000713)	-0.000377 (0.000815)	-0.00440*** (0.000846)	-0.0000540 (0.000674)	-0.0000891 (0.000700)
log pages	0.683*** (0.0352)	0.512*** (0.0354)	-0.334*** (0.0207)	-0.207*** (0.0201)	0.0716*** (0.0130)	0.0723*** (0.0137)
featured rank		-0.00231*** (0.0000859)		0.00170*** (0.0000709)		0.00000991 (0.0000228)
N	22,759	22,759	22,759	22,759	13,230	13,230

Notes: This table shows regressions of usage and quality measures on an indicator for whether the book contains AI as well as category and month fixed effects, on the subset of the main dataset's AI Detection subsample matched to featured ranks. The even columns additionally control for the book's featured rank. Columns (1) and (2) use $\ln(\text{adjusted ratings})$ as the dependent variable. Columns (3) and (4) use $\ln(\text{sales ranks})$ as the dependent variable, in a censored normal regression. Columns (5) and (6) use the average star rating as the dependent variable. Standard errors are clustered at the category $\times$ month level, and regressions are weighted to reflect the numbers of entering books by category and release month. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Consumer surplus effects of LLMs by year

year	AI share	%Δ CS		
		$\sigma = 0.443$	$\sigma = 0.25$	$\sigma = 0.75$
2023	23.12%	+0.44%	+0.59%	+0.20%
2024	33.57%	+3.26%	+4.38%	+1.45%
2025	52.70%	+7.23%	+9.82%	+3.19%

Notes: This table shows consumer surplus effects from adding AI books in each year. The “AI share” column shows the weighted average annual AI-detected share. The remaining columns show the percent increase in consumer surplus from the full sample (both human and AI books) in each year, relative to just the human-authored books of the year, based on our preferred nested logit substitution parameter $\sigma = 0.443$ and two alternative substitution parameters.

Table 6: AI books and usage over time

	(1)	(2)	(3)	(4)
AI-containing	-1.740*** (0.0855)			
AI $\times$ year 2023		-2.197*** (0.222)	-0.438* (0.237)	-0.426* (0.238)
AI $\times$ year 2024		-1.863*** (0.133)	-0.354* (0.186)	-0.364* (0.187)
AI $\times$ year 2025		-1.503*** (0.111)	-0.341** (0.138)	-0.391*** (0.149)
# prior AI books				0.0559 (0.0462)
Add'l fixed effects	None	None	author + age	author + age
N	9,026	9,026	4,629	4,629

Notes: Regressions of $\ln(\text{adjusted ratings})$ on an indicator for whether the book contains AI using the AI Detection author subsample (2023-2025). Columns (1) and (2) include neither author nor age fixed effects. Column (3) includes author and author age fixed effects, and column (4) adds a variable for the author's number of prior AI-containing books. All columns include release-month fixed effects. Standard errors are clustered on category $\times$ release month. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A Appendix

A.1 Adjustments for over-counting

Books may be indexed in more than one Amazon category, so the sums of the category totals may overstate total releases. We develop two adjustments to support estimates of total release volume.

Direct empirical approach using subcategories: Each Amazon category is made up of subcategories. For example, the Romance category has 38 subcategories. We obtain one estimate of over-counting by comparing sums of subcategory totals to category totals.⁴⁰ In particular, we calculate this ratio for each of 30 categories in each January, for all category-months for which the total does not exceed 6,400 so that our release estimates are accurate up to the next 16. In all, we calculate this ratio for 154 category-months. The mean (median) ratio is 0.823 (0.80). A regression of these ratios on year indicators shows no time trend, and we cannot reject the hypothesis that the coefficients on the year indicators are jointly zero.

Probabilistic approach: If we observed the entire population of books listed in all categories, we would see any over-count directly. Suppose there were two symmetric categories with 1 million total listings, half of which are duplicated across the two categories. Each category would have 0.25 million unique titles and 0.25 million in common with the other. Overall, there would be 0.75 million distinct books, and the correct unique title count would be 0.75 times the observed number of listings.

Suppose that instead of seeing all million listings, we saw half as many, chosen at random. We would see 0.5 million listings, with 0.25 million listings for each category. How much overlap would we see? The probability that one of the 0.25 million actually-duplicated titles in the population would be duplicated in the sample would then be 0.5^2 , or 0.25. Hence, 62.5 thousand titles would appear in both of the 0.25 million category samples; and the number of distinct titles among the 0.5 million sample listings would be 0.4375 million ($= 2 \times (0.25 - 0.0625) + 0.0625$). The observed overlap in the sample understates the population overlap: Of the 0.5 million sample listings, 87.5 percent would represent distinct listings, rather than the 75 percent in the full population. This logic allows us to infer the population overlap from the overlap in a random sample.

⁴⁰In principle one could do this for the sum of category totals in relation to the overall total, but the overall totals are always too high to return a precise, and therefore comparable, count.

More formally, we estimate the extent of this over-counting using the independently drawn samples within each category and month. Let N_{ct} denote the true number of listings in category c in month t , and let n_{ct} denote the number sampled. The category-specific sampling probability is therefore $p_{ct} = n_{ct}/N_{ct}$. A product that belongs to categories c and c' is observed in both category samples with probability $p_{ct} p_{c't}$. Accordingly, each observed ASIN-category pair is weighted by $1/(p_{ct} p_{c't})$ to estimate the underlying pairwise category intersection. Products observed in three categories analogously receive a triple-intersection weight of $1/(p_{ct} p_{c't} p_{c''t})$.

We then define $L_t = \sum_c N_{ct}$ as the summed category count, $\hat{P}_t$ = the estimated sum of pairwise intersections (the observed intersections weighted by their probabilities), and $\hat{T}_t$ = the estimated sum of triple intersections. The estimated number of distinct products in month t is

$$\hat{D}_t = L_t - \hat{P}_t + \hat{T}_t.$$

Thus, $\hat{s}_t = \hat{D}_t/L_t$ is the estimated monthly ratio of distinct to total listings in the population. The procedure assumes independent random sampling across categories. Because books have the same publication date regardless of their category, overlap can occur only within each month. At the same time, small samples can make month-specific estimates noisy. When we calculate $\hat{s}$ separately by month, we do not find systematic time variation in overlap. A regression of the monthly estimates on a time trend yields a coefficient of -3.63×10^{-6} (se = 0.0022); and we cannot reject the hypothesis of common coefficients in a regression of the monthly ratios on year effects ($F = 0.35$, p -value = 0.88). We calculate a single mean (median) probabilistic adjustment factor of 0.57 (0.66).

A.2 AI detection errors

We can use what we know about the AI detection error rates to bound the inaccuracy of our estimates. Books produced before the advent of LLMs cannot contain AI. Hence, we measure Pangram’s false positive rate as the AI share among sample books released prior to November 2022. The false positive rate is 0.97 percent if we classify a book containing any AI as AI-created. With a threshold of 0.9, the false positive rate falls to 0.59 percent. We do not observe the false negative rate among AI-containing works after October 2022, but [Jabarian and Imas \(2025\)](#) have undertaken highly relevant tests, running 1,000-word human- and AI-written novel excerpts through various AI detectors. They report that Pangram produced a

false negative rate between 0.45 and 1.61 percent for ChatGPT-produced excerpts, between 1.20 and 2.86 percent for Claude Opus 4, and between 1.41 and 3.82 percent for Gemini 2.0 Flash. In their tests, Pangram’s false positive rate was 0.10 percent at a threshold of 10 percent AI and zero for higher thresholds.

Given the information above about false positive and false negative rates, we can adjust our estimates (Aigner, 1973). If q = the detected AI rate, f = the false positive rate, and g = the false negative rate, then the true AI rate is $\pi = (q-f)/(1-f-g)$. Moreover, if $\tilde{\beta}$ = the estimated coefficient on a measured AI indicator, the true β is governed by

$$\beta = \tilde{\beta}/\lambda,$$

where $\lambda = \frac{q(1-q)}{\pi(1-\pi)(1-f-g)}$. Taking the most conservative (highest) false positive and false negative rates from our context ($f = 0.0097$, $g = 0.0382$), if detected AI (q) were 0.5, then the true AI share would be very similar ($\pi = (0.5-0.0097)/(1-0.0097-0.0382) = 0.515$). Similarly, the estimated coefficient from a regression of, say, a usage measure on the error-containing measured AI indicator would be misstated by only about 5 percent. We conclude that our results are not driven by inaccurate AI detection.

A.3 Validity of our usage measure

Given our reliance on the number of ratings a book has received (r_j) as a usage measure, it is important to validate our quantity/usage proxy. For this, we use an additional dataset. We have daily estimates of sales, average star ratings, and the number of ratings received, for ebooks sold at Amazon between mid-2017 and mid-2022, from Bookstat. The data include titles appearing among the top few hundred thousand per day in sales. For each of these titles, we have book metadata, including publication date, genre, author name, and publisher. For the validation, define $r_{j,\tau}$ as the cumulative number of ratings received by book j as of τ periods since its release. Using the Bookstat estimates of Amazon sales, along with information on the number of Amazon ratings the ebook has received by that time, we explore how well ratings reflect sales. Starting with the ebooks published in April 2017 and in January of each year from 2018 through 2020, we calculate the estimated cumulative sales (q) of each title by month (τ) until mid-2022. Figure A.1 shows how the cross-book correlation of $\ln(q)$ and $\ln(r)$ evolves with time since publication. By 30 days after publication, the correlation reaches nearly 0.4 and rises to 0.6 after about seven months. This indicates that our usage measure is informative about sales for books published through the end of

the sample period. Because the main usage measure is obtained in April of 2026, even the December 2025 sample books have four months to accumulate ratings.

A.4 Adjusting the number of ratings

The number of ratings a book has received as of the time of data collection depends in part on the amount of time that has elapsed since publication. Some of our estimation approaches require a measure of the number of ratings received that is comparable across books, and that requires a way to adjust the number of ratings for the time elapsed between publication and observation.

We create an adjustment using information on the evolution of books' ratings across time. In particular, we observe the number of ratings for each book at two points in time, roughly 60 days apart.⁴¹ That is, we observe r_j at a time τ (τ periods since j 's publication) and at time $\tau + k$. Because our books have a range of ages at each observation time, we can characterize how r_j varies over time with the age of the book. We assume that the number of ratings for a book grows proportionally according to

$$r_{j,\tau+k} = r_{j,\tau} e^{g(\tau)k}, \quad (\text{A.1})$$

where $g(\tau)$ is the daily ratings accumulation rate at time τ since publication. We calculate $g(\tau)$ for each 30-day age interval.

We want to normalize the ratings measure to the number of ratings that would have accumulated by an age of T , i.e., a time that is T after publication (we use 72 months). We do this as follows: $r_{jT} = r_j e^{\sum_{i=\tau}^T g(i)}$, where i runs across 30-day intervals.

We calculate the daily ratings growth rate from age τ as $\ln(r_{\tau+k}/r_\tau)/k$. Ratings grow quickly immediately after publication and more slowly as books age. Figure A.2 shows the average time pattern of ratings growth, among books with positive ratings, normalized to one at 72 months.


We handle books with zero ratings at the time of observation differently. Using our repeat observation data, we see the share of zero-rating books transitioning to having positive numbers of ratings at each age, as well as the average ratings for those books at time $\tau + k$. We use these empirical transition shares to assign the books obtaining ratings to the average number of ratings that books getting positive ratings obtain, times the probability of transitioning. Adjusted ratings are all then positive, and we adjust the non-zero ratings

⁴¹We collected data on February 6 and, later, on April 5, 2026.

forward as above. We refer to the adjusted number of ratings as $\tilde{r}_j$.


A.5 Appendix figures and tables

Figure A.1: Quantities and the number of ratings



Notes: The figure shows the correlation between log cumulative sales and the log of ratings as of various numbers of days after publication. Included ebooks were published in April 2017 and in January of each year from 2018 through 2020.


Figure A.2: Ratings adjustments by book age at collection




Notes: The figure shows the proportional continuous adjustment factors for ratings. We divide raw ratings by the adjustment to create our 72-month “adjusted rating” ($\tilde{r}_j$).

Figure A.3: Usage distributions across publication dates – temporal approach


(a) Overall average



(b) Conditional on rank percentiles



(c) Conditional on rank positions



Notes: The figures show month effects from regressions of $\ln(\tilde{r}_j)$ on month effects and category dummies, weighted to reflect total releases in the categories. Panel (a) is the overall average. Panel (b) includes fixed effects for 20 rank-percentile quantiles. Panel (c) includes fixed effects for rank position ranges within categories by 100, up to 2000.