AI and the quantity and quality of creative products: the case of books*

Joel Waldfogel
University of Minnesota, NBER, ZEW

Imke Reimers
Cornell University and ZEW

December 1, 2025

PRELIMINARY – DO NOT CITE

Abstract

The arrival of LLMs has revolutionized the creation of books. Using data on new books appearing for sale at Amazon, we document a doubling in the number of new titles appearing each month between 2023 and late 2025. The vast majority of the increase is new author entry, while the remainder is increased output of existing authors. Using data on the number of ratings books have received as a measure of "quality," we document that the post-2022 AI-heavy vintages have lower average quality; but the top 1,000 books per category have higher quality. The effect is larger in categories with faster growth in new titles. Despite low average quality, the AI influx is raising the quality of books near the top of the distribution.

*Author names are in random order.

Introduction

Artificial intelligence, in the form of large language models, promises to revolutionize many tasks, occupations, and industries, including – and perhaps especially – the creative industries. Because of LLMs' ability to create text, the book industry may be particularly susceptible; and during the period 2023-2025 LLMs became capable of producing sustained text passages. It is unclear, however, how LLMs might affect books, authors, and the supply of new books. How much do they act as complements to writing skills, making existing authors more productive? How much do they facilitate de novo entry? And, most importantly, what does the availability of AI tools do to the quality distribution for entering books and therefore the welfare of consumers?

Reductions in entry costs that raise the number of new products have traditionally benefited consumers through the predictability of product quality at entry. Cost reductions facilitate more "draws from the urn," and with complete unpredictability, additional products are as good, on average, as existing products. Increased entry delivers large welfare benefits (Aguiar and Waldfogel, 2018). While LLMs deliver cost reductions, their effect does not simply facilitate more draws; they may change the distribution of realizations by granting writing ability to those who lack it while complementing the ability of skilled writers. How does the use of LLMs in book writing affect the distribution of realized qualities? Will LLM-aided creation draw from the existing distribution of qualities, or will the new crop of works occupy a narrow range of realized quality?

Against this backdrop this paper asks three questions. First, how has the advent of LLMs affected the number of new books created? Second, which sorts of writers – new entrants or writers already working – are producing additional books using the new technology? Third, how do LLMs affect the distribution of the new works' quality? What happens to the average quality of works created, and what happens to the appeal of the choice set to consumers?

We assemble datasets on both the numbers of books published in each of 30 categories at Amazon each month, from 2006 to 2025, and – for subsets of the individual books – various measures of the books' eventual success and consumer appeal (as of November 2025), including the number of persons rating the books and the average star ratings.[1] We document that the advent of LLMs has had a large impact on the production of new books. The number of new titles appearing each month has doubled since late 2022 and has risen by a factor of ten in some categories; and the increase in new titles 2023-2025 coincides with the diffusion of LLMs. The increase in output comes largely from entry of new authors and, to a lesser extent, from increased productivity of incumbents.

We measure the impact of LLMs on the quality of new books in three ways. First, we construct a usage measure that is comparable across vintages, and we use it to show changes in "quality" – implied by our usage measure – following the introduction of AI tools for writing. Second, we show that the effects we document are larger in fast-growing categories (where new entry more than doubles between 2022 and 2025) relative to more slowly-growing categories. Third, we show how effects on quality vary with the continuous variation in entry growth across categories. Using each of these approaches, we document increases in the quality at particular rank positions in the distribution along with a decreased in average quality and a decrease in quality at particular percentiles of the distribution. Despite the low average quality of works in the AI influx, the quality of the top 1,000 books per category is as high or higher, with large potential benefit to consumers.

The paper proceeds in five sections after the introduction. Section 1 discusses the background, including the history of the development of LLMs and their features relevant to book authorship and the related academic literatures. Section 2 provides a simple theoretical framework showing how an influx of new products could differentially affect quality

[1] Our main measure of books' appeal is the number of star ratings they have received, which we show to be highly correlated with quantities sold.

at percentiles vs absolute rank positions in the distribution to guide the empirical analysis. Section 3 describes the data used in the study. Section 4 shows the increase in book entry as AI has diffused, as well as the parts of the increase coming from new vs existing authors. Section 5 presents our findings on the change in the distribution of quality with the increase in entry. A brief conclusion follows.

1 Background

1.1 Large language models

The development and diffusion of large language models in the past few years has held out the possibility of changing the book-writing process. While early LLMs were limited in their ability to produce reliable text, more recent products have stronger writing capabilities. Before 2019, neural models could not sustain narrative structure beyond a few paragraphs (Radford et al., 2019). GPT-2 demonstrated that transformers could generate stylistically plausible paragraphs but lacked memory and thematic control. The release of GPT-3 in 2020 marked a step-change: its 175 billion parameters, trained on extensive internet text, enabled extended continuity, stylistic mimicry, and few-shot learning. At this stage, professional and amateur writers began integrating the model into brainstorming workflows and prototyping dialogue or scene sketches.

The public deployment of ChatGPT in late 2022 brought LLMs to prominence. By removing technical barriers, OpenAI's chat interface allowed non-specialists to instruct a model conversationally, request revisions, and co-create long-form passages interactively (OpenAI, 2022). With GPT-4 in 2023, significant gains appeared in factual accuracy, contextual awareness, and stylistic consistency (Bang and et al., 2023). By 2025, advanced systems such as OpenAI o1, GPT-4.5, and open-source models (e.g., Mistral 7B, LLaMA 3) offered long-context reasoning over tens or hundreds of thousands of tokens, effectively allowing the

maintenance of characters, themes, and continuity across book-length manuscripts. The left panel of Figure 1 shows that the volume of Google searches for ChatGPT (and other LLMs) rises quickly beginning in late 2022 and especially after July of 2024.

1.2 Author engagement with LLMs

LLMs are not merely a theoretical possibility for writers; it is clear LLMs are both attracting new writers and being used by existing writers. Many existing authors are adopting AI. A 2025 survey of 1,229 authors found authors divided in their views. Nearly half (45 percent) were using AI to "assist with their work" while 48 percent were not 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 thought its use to be "unethical." ChatGPT was used by 85 percent of authors using AI. Just over a third of respondents had published through traditional publishers; the remainder had used self-publishing. Almost all (96 percent) had begun publishing in 2023 or earlier. Nearly half had begun in 2020 or earlier. In another survey, 61 percent of writers reported using AI tools.[3]

A cottage industry has arisen providing LLM advice to would-be de novo authors. For example, `bookautoai.com` offers support so that "you can publish a fully formatted, human-level book in under 30 minutes." They advise that AI is differently suited to different genres: For nonfiction (self-help, business), AI is suited for creating structures and outlines but "requires editing for expertise and tone." AI has "low" suitability for literary fiction, poetry, psychological/satire, academic/technical, and motivational/personal.[4] The best genres for AI, according to this source, are romance, romantasy, cross-genre stories (e.g. combining sci-fi, mystery, or horror), and thrillers and mysteries.[5]

[2] See `https://insights.bookbub.com/how-authors-are-thinking-about-ai-survey/`.

[3] See `https://www.publishersweekly.com/pw/by-topic/digital/copyright/article/99019-new-report-examines-writers-attitudes-toward-ai.html`.

[4] See `https://blog.bookautoai.com/ai-book-generator-genre-suitability/`.

[5] "The best genres for AI-generated books include romance, romantasy, thrillers, and cross-genre fiction,

Some authors have apparently achieved success using AI. Tim Boucher has written hundreds of books using AI.[6] There is a great deal of low-quality work, which critics derisively term "AI-slop." Goodreads has many user-created lists of slop.[7] The right panel of Figure 1 shows the number of monthly releases of works identified as "AI slop" by Goodreads users. The timing of the appearance – and identification – of these works provides additional evidence of AI's effect on authorship in the 2023-2025 period.

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 analysis of AI-driven content markets.

A first strand examines how digitization has affected the production and quality of creative goods. 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 suggest that technology shocks that reduce production and distribution costs can simultaneously raise both the volume and the realized quality of new creative outputs.

A second strand evaluates how generative AI affects cognitive and creative work. Exper-

as they have strong reader loyalty and predictable structures." See `https://blog.bookautoai.com/best-genres-ai-books-2025/`.

[6] See `https://www.newsweek.com/ai-books-art-money-artificial-intelligence-1799923`.

[7] See `https://www.goodreads.com/list/tag/ai-slop` and `https://www.goodreads.com/list/tag/ai-book`. "Wouldn't it be great if books written entirely by generative A.I. (ChatGPT or countless others) were prominently labelled as such? Unfortunately, in these wonderful times we live in, most people flooding online marketplaces with A.I.-generated books prefer to hide that fact, because of course no one would buy them if they knew. Fortunately, those of us who have some exposure to A.I. "writing" can usually tell. This is a list to call out books that you are 100% certain were written by A.I." See `https://www.goodreads.com/list/show/219041.Books_I_Have_No_Doubt_Are_Written_By_A_I_`.

iments show large productivity gains from access to large-language models. Noy and Zhang (2023) report that GPT-based writing assistance improves task completion times by roughly 40 percent and disproportionately benefits lower-skilled workers. Brynjolfsson et al. (2023) find that an AI-based assistant in a large call-center improves worker productivity by 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 or diversity (Doshi and Hauser, 2024). Together, these studies imply that AI may significantly increase the throughput of new book production but leave open whether quality rises or becomes homogenized.

A third strand considers the effect of generative-AI on an actual market for creative goods. Goldberg and Lam (2025) analyze a large image marketplace. Using theoretical modeling and marketplace data, they show that generative-AI tools increase entry and product variety.

2 Theory

Creating a book is an entry decision. Creators compare the cost of creation to the expected benefit, and they create a book if the expected benefit exceeds the cost. Benefits include revenue and whatever other gratifications come from producing a book. Because of unpredictability, the realized quality is a random variable (Aguiar and Waldfogel, 2018). AI tools have reduced the cost of book creation. People who lacked the ability to write a book may now be able to create them, and people who could already write books may now be able to do so more quickly. Hence, we expect the number of new books to rise.

Cost reduction gives rise to additional creation. With unpredictability, general cost reduction can raise welfare by delivering additional products from the same underlying quality distribution, some of which are valuable to consumers. But the arrival of AI is not simply a reduction in the cost of entry. AI may allow talented people to write better books, or more

books of similar quality. AI might also allow unskilled people to create works of uncertain but potentially higher quality than previously possible. These concerns motivate the paper's main question: What happens to the quality distribution for entering books when AI tools become available?

2.1 Quality distributions with an AI influx

Given that creation is a random process, what matters for consumer welfare is the distribution of realized product qualities. A simple model of the realized distributions illustrates how AI tools might affect the distribution of product qualities. Define μ_k and σ_k as the mean and standard deviations of normally distributed realized qualities in regime k, where k is either pre- or post-LLM. Use $f_k()$ to denote these distributions. Define N_k as the number of products introduced in each period. Define δ_j as the mean utility of product j; and the δs are generated from the distributions above.

Consider, for example, the plain logit model. Then the consumer surplus (CS) of a choice set is $(M/\alpha)\ln(1+\sum e^{\delta_j})$. If the mean μ falls, then CS will decline, all things equal. But if N rises, then CS will rise as well. Hence, if μ declines but N rises, then the effect on welfare is ambiguous.

Some cases are helpful for organizing thinking. First, suppose that N rises while μ and σ remain constant. Then the average quality will remain the same, the quality of a product at any particular percentile (e.g. the 99^{th}) will remain constant, and the quality of the k^{th} best product will rise. Second, suppose that μ declines while N and σ remain constant. Then the average quality will fall, as will the quality at any percentile of the distribution. Moreover, the k^{th} best product's quality will also fall. Third, if N rises while μ declines, then effects are slightly more nuanced. The average quality will, of course, decline, as will the quality at any particular percentile. But the quality of the k^{th} best product may rise or fall.

This is illustrated in Figure 2. Panel (a) compares the pre-LLM distribution (in blue)

with the larger post-LLM distribution (in red), where the post-LLM distribution is illustrated with five times more products and a lower mean. Panel (b) compares the quality of products at the same *percentiles* of the pre and post distributions, and the quality is lower in the post-LLM distribution. Panel (c) compares products at the same *rank position* (e.g. k^{th}-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.

As we discuss in the data section, we are unable to obtain the entire distributions of realized qualities. Still, we can take random samples for estimating the effect on the mean quality; and we can obtain samples at the various percentiles and points in the distribution. The framework above gives a way to interpret the findings and, in particular, what they mean for the impact of LLMs on the distribution of realized product quality.

3 Data

We are interested in characterizing volumes of new books entering over time – including times before and after the influx of AI tools that might affect authorship – as well as the distributions of those books' quality or success. Ideally, we would have a full list of all books published since, say, 2000, along with measures of those books' sales in each year (and perhaps other measures of quality such as star ratings). Moreover, we would also want to know the extent to which AI was used in the creation of each title. While ideal data are not available, we have ways to measure the phenomena of interest. We are able to observe the number of new books by category and month; and we can observe product-level information – including measures of attention such as ratings and the number of people rating each title – for percentiles and rank positions of the distribution, for two entire subcategories, and for random samples of the entire population of new books.

3.1 Sources and constraints

Our data are drawn from two sources, Amazon and Goodreads; and features of these sources determine what is feasible to observe. First, we use Amazon to get the numbers of books published by month and by category, using queries to Amazon's advanced book search function at `https://www.amazon.com/advanced-search/books`. We specify English-language books in new condition and in the Kindle format, and we specify their publication dates by month and by category (or subcategory). A query delivers a list whose first page indicates "1-16 of x results." When there are fewer than 1,000 results, x is exact. Up to 10,000, the result is reported to the nearest 1,000, From 10,000 on, the result is reported to the nearest 10,000. The search result also includes the number of search-result pages, with a maximum of 400, indicating 400×16, or 6,400 results. Hence, the number of books in search is exact for 1-1000, to the nearest 16 for 1,001-6,400, to the nearest 1,000 up to 10,000, and to the nearest 10,000 for higher numbers. Hence, we can produce monthly time series of new titles by category and subcategory.

We also use these search results 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 75 pages of search results and therefore only the first 1,200 ($= 75 \times 16$) books. There is an option to sort results by publication date within month. Using both ascending and descending sorts, it is possible to obtain up to 150 pages of monthly search results. The ability to see only 75 pages constrains the categories, or subcategories for which we can obtain book level data. In particular, the group of books chosen cannot include more than 2,400 titles published per month ($=150 \times 16$). In addition to being sorted by publication date, books can also be sorted by numbers of ratings received. This sort is only available descending, so one can see as many as the top 1,200 books by publication month can category.[8]

[8]Books may be sorted by "bestselling," and while this does deliver books with better sales ranks higher

We use Goodreads (`https://www.goodreads.com/`) to obtain author histories. Given an author name from an Amazon title-level dataset, we first obtain the authors' Goodreads page URL. We then obtain the information on each prior release. The information includes genre, publisher, publication date, and the number of Goodreads user ratings of 1-5 stars, respectively.

3.2 Additional datasets

Bookstat 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, if the title appears in the top few hundred thousand per day. For each of these titles, we have metadata, including publication date, genre, author name, and publisher. We use these data, from Bookstat, to show that the number of ratings a book receives is closely related to its sales quantity.

USA Today bestseller list Finally, we have the USA Today weekly bestseller list for July 2023-November 2025. We observe the top-150 sellers per week, along with author name, genre, publication date, and publisher. These data allow us to calculate share of each week's bestsellers that were published in, say, the last year.

3.3 Analysis datasets

We create four datasets. First, we have aggregate monthly time series measures of new works published, by month and by 30 Amazon categories, for 2020-2025.[9]

in the sort order, the order is not exactly according to sales rank.

[9]The 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.

Second, although we cannot obtain information on all titles, we can create an effectively random sample of books from all categories and publication months, 2020-2025. We obtain a random sample by sorting books by publication date. We then obtain the first two pages, delivering the 32 books per category and publication month that were published nearest to the beginning of the month. The random sample allows us to examine the effects of AI influx on the average characteristics of entering books. We refer to this sample as the "random sample."

In addition to the average quality of books, we would also like to characterize effects of the AI influx on the entire distribution, so we construct two datasets, for points and percentiles of the distribution. Our basic measure of book success is the number of ratings that a book has received. Ideally, we would have this measure for all books in each of the 30 categories for, say, 2020-2025. Collecting all of these data is infeasible given the limitations of Amazon's search feature. What we do, instead, is to collect a third dataset containing the titles on a few search result pages sorted by number of ratings, for each publication month and category. In particular, we collect the first page (the top 16 books for each category and publication month), the seventh page (books 97-112), and the thirteenth page (books ranked 193-208) and so on. This allows us to characterize how the influx of new books affects three points in the distribution, the top and at rank positions $100, \ldots, 600, 800,$ and $1,200$. We refer to this sample as the "rank position sample." We also collect analogous data at percentiles of the monthly distributions of ratings, by category. We do this for the 1^{st}, 5^{th} and 10^{th} percentiles of the distributions. We refer to this sample as the "percentile sample." The first three dataset are described in Table 1.

Fifth, we have a complete "census" of new titles in a few subcategories that are sufficiently small to allow their collection, including "romantasy" (a subcategory of Romance) economics (a subcategory of Business & Money), science fiction fantasy (Science Fiction & Fantasy), women sleuths (mystery), alternative history (Literature & Fiction), world history (History).

For each title we observe the author, the publication date, the number of persons rating the book as of November 2025, the average star rating for the book as of November 2005, and whether the book is rated. We collect all of the titles in these categories released between January 2008 and October 2025. These data are described in Table 2.

3.4 Validity of our usage measure

We do not observe book sales. Instead, we have the number of ratings that each book has ever received as of data collection in November 2025. Define r_{jt} as the number of cumulative ratings received by title j which was published at time t. Using the Bookstat estimates of Amazon sales, along with information on the number of Amazon ratings ebooks have received, we explore how well ratings reflect sales. Starting with the ebooks published in the January of 2018, we calculate the estimated cumulative sales of each title by month until mid-2022. The left panel of Figure 3 shows the end-of-sample correlations of the logs of sales quantities and the number of ratings, by book category. These correlations are generally between 0.6 and 0.8. The right panel of Figure 3 shows how the cross-book correlations evolve with time since publication. By 30 days after publication it reaches nearly 0.4, and it rises to 0.6 after about seven months.

4 The growth in new book entry

This section describes our empirical on volumes of book entry over time. Section 4.1 describes overall entry patterns before, and after, the advent if AI tools. Section 4.2 describes book entry by new and already-existing authors.

4.1 Aggregate patterns

Figure 4 shows the number of new Kindle-available titles by month and for the 30 broad Amazon categories, for January 2008-October 2025. The number of new titles in each category rises substantially from 2008, when Amazon introduced the Kindle device and store, to about 2014 (Waldfogel and Reimers, 2015). The number of new titles is, for the most part, roughly stable between 2014 and the early 2020s, then rises sharply in 2024 and especially 2025 in all categories. These categories encompass both fiction (e.g. literature & fiction, romance) as well as nonfiction (e.g. law, business & money).

Figure 6 reports the numbers of new titles by month for the two categories for which we have the full census. Both Romantasy and Economics experience sharp increases in the numbers of new titles during 2025. The large increases in the numbers of books appearing in these subcategories, as well as the full categories, mirror the growth in attention to ChatGPT in Figure 1.

The growth in the rate of new publications is large across most categories. As Figure 5 shows, the monthly number during 2025 is between 1.5 and 12.5 times as high as the average during 2020-2022. Categories with particularly large proportionate growth – in some cases, more than a factor of five – include travel, sports & outdoors, computers & technology, self-help, and teen & young adult.

Entry with a Kindle edition, but without a hardbound edition, has a much lower fixed cost than entering with a hardbound edition as well. And, as Figure 6 shows, hardbound entry is rare for new titles in romance, science fiction, mystery, and literature. Moreover, the tendency for books to have a hardbound edition rises very little with the high entry growth following 2022. This suggests that the growing number of titles released following 2022 were not expected, at release, to have returns justifying the larger investment associated with a hardbound book.[10]

[10]They are also likely not distributed in physical stores.

4.2 New production by entrants and incumbents

AI's effect on new products might operate through the extensive or intensive margins. We might see growth in works released by incumbent authors, or we might see entry of new authors. We can explore this directly using the census datasets on the six subcategories, in Figure 7. We divide the total number of books appearing each month into those by authors first appearing in the month, additional books by authors first appearing before 2023, and additional books by authors first appearing since 2023. The growth in entry during the LLM era is primarily driven by entry in authors' first month of appearance and, to a lesser extent, by repeat appearances from authors entering since 2023. These six subcategories show no growth in entry by authors who were active prior to 2023. In short, the growth in entry is driven entirely by authors entering since the existence of LLMs. If the author surveys cited above are to be believed, then effects of AI on authors entering earlier would have to operate through something other than quantity.

5 Entry and quality

Below we address our main question. Section 5.1 describes our approaches to measuring the effect of the AI influx on the distribution of book qualities. Section 5.2 presents our main evidence on quality effects. Section 5.3 discusses additional evidence on whether the AI-influx is valuable based on shares of 2025 sales to recent books and recently-entering authors.

5.1 Empirical strategies

We have three broad empirical strategies available to us, including a purely intertemporal comparison and a cross-genre intertemporal comparison. We describe these – and their advantages and disadvantages – below.

5.1.1 Intertemporal approach

LLMs become available over time, so one important source of variation is temporal. As LLMs diffused to authors, was there a growth in new works created, and what happened to their quality? One challenge to the purely intertemporal approach is that out main outcome measure, the cumulative #of ratings received since release, grows over time within title. Hence, older titles appear to be "better," simply because they are older.

We address this by adjusting our basic measure for time of exposure. Using data on all of the ebooks released in April 2017, and the number of ratings they receive in the subsequent 56 months, we develop measures of the numbers of ratings received by month since release and genre. This allows to scale up our measure for recent books, and to scale-down the measure for older books, for a normalized measure based on 12 months since release, which we term "adjusted attention," or $\tilde{r}_j$.

We then explore how adjusted attention evolves over time. We are interested in the evolution of adjusted qualities at various points in the distribution as LLMs diffuse. We explore this via regressions of the form

$$\tilde{r}_j^k = \mu_c + \mu^k + \psi_t^k + \varepsilon_j$$

where k indicates the point in the distribution, and c indexes category. The coefficients ψ_t^k show how book quality evolves for group k. We consider the following groups k: a) the average across the entire distribution, b) percentiles of the distribution (e.g. the books near the 5^{th} percentile), and c) absolute points in the distribution. We implement this approach directly using both the census samples and as well as the random, percentile, and rank position samples.

5.1.2 Cross-category intertemporal approach

Different sorts of books may be amenable, to different extents, to LLM effects. We also exploit this to check whether categories that are "more treated" experience bigger effects on quality. Because this approach relies on comparisons between more and less treated groups of books over time, we can include time effects to deal with the changing meaning of r_j over time.

A simple version of this approach divides book categories into two groups, those highly exposed to AI and those less exposed. We then ask how various outcome measures change over time for the highly exposed groups, relative to the less exposed groups.

That is, we run regressions of the form

$$r_{j(c)t} = \mu_t + \mu_c + \lambda_t \times \delta_j^h + \varepsilon_{j(c)t}$$

where δ_j^h indicates whether book j is in the highly exposed group. The coefficients λ show how the outcome measures for the highly exposed group evolve, relative to their pattern for the less exposed group. We term this the "high-low" approach.

5.1.3 Quality and the volume of category entry

A second version of this approach uses the number of books appearing in a category per month (n_{ct}) as a measure of the extent of the category's treatment.

We run the regression

$$r_{j(c)t}^k = \mu_c + \mu_t + \mu^k + \phi^k n_{ct} + \varepsilon_{j(c)t}.$$

The coefficient ϕ^k shows the extent to which our quality measure (for distributional group k) changes more in groups experiencing a greater "LLM treatment," i.e. and increase in the number of books appearing.

This approach requires cross-group variation in the number of books appearing, so we implement it using the random, percentile, and rank position samples covering all categories of books. We term this the continuous approach.

5.2 Quality and success of new work

We study effects on quality using two kinds of data, the full censuses for six subcategories, as well as the category samples describing random books, those at particular rank positions, and those at two percentiles of the distribution. We measure effects in the three broad ways described above.

5.2.1 Intertemporal results

Figures 8–10 show the results of the intertemporal approach on a) random books, b) absolute ranges in the distribution (1-25, . . . ,201-225), and c) percentile deciles. Figure **??**, with a vertical line at October 2022, shows that average quality declines in 2024 and 2025. Figure 9 shows that the quality measure rises after 2022 at any of the top 200 rank positions. Figure 10 shows that quality declines in the top 5 deciles, is constant in the next two, and rises in the bottom two deciles. These results Figures 9 and **??** indicate that the influx of books with AI reduces average quality while raising the quality at absolute points of the distribution (to 225). The percentile figure shows that, for high percentiles, quality falls while it rises for low percentiles.

Figures 11–13 repeat the intertemporal exercises for the category samples. Figure 11 shows that the average quality of books holds steady, which stands in contrast with the results from the census data. Figure 12 shows that quality rises after 2022, particular for rank positions 200-1,200. Figure 13 shows that the average quality in the first and fifth percentiles of the distributions falls, particularly in 2024 and 2025.

5.2.2 Comparison of high and low-AI diffusion categories

Figure 14 shows the evolution of various quality measures for the high-growth categories relative to the low-growth categories. These figures are based on the category random, rank position, and percentile samples. Panel A shows that the quality of random books in the high-growth groups falls relative to the low-growth groups. Panel B – which pools rank positions and includes a rank position fixed effects – shows that quality rises for the high growth categories relative the low-growth categories after 2022. Panel C shows that quality falls in the first and fifth percentile regions for high relative to low-growth categories.

Figure 15 disaggregates the rank position results in panel B of Figure 14 by rank position, showing that quality is flat at the top of the distribution (rank position 1 - 100) and rises at positions 200-1,200. While the AI influx seems to raise the quality near the top of the distribution, it appears not to raise quality at the very top of the distribution.

Similarly, Figure 16 disaggregates the percentile results of panel C and shows declines in quality, for high-growth relative to low-growth categories, at both the 1^{st} and 5^{th} percentile ranges.

5.2.3 Quality and the number of entering books

We explore the relationship between outcome measures and the by-category growth in titles directly using regressions, and we report results not only for our main quality measure (the # of ratings) but also whether books are rated and the average star ratings. Table 3 reports the coefficient on $\log(n_{ct})$ from the regression for random (average) books. Categories with larger growth have a lower propensity to receive any ratings ($I(rated) > 0$), fewer ratings, and lower star ratings. Table 4 reports results of analogous regressions, with separate coefficients for each of the rank positions. Faster-growing categories experience increases in the number of ratings at all points starting with 100, while the tendency to be rated rises for rank positions 400 and above. Average star ratings rise for positions 100-300. Table 5 does the analogous

exercise for the first and fifth percentiles of the distribution. The tendency to be rated is unaffected, but the number of ratings falls substantially. Average star ratings are unaffected.

The results above paint a very consistent picture. AI-induced growth in book entry reduces the average quality, while increasing quality near the top of the distribution and decreasing quality at particular percentile ranges. The framework in Section 2 is useful for guiding interpretation of the results. The contrast between the positive effects at particular rank positions and negative effects at particular percentiles (including the average) is consistent with the large observed increase in book entry, along with a decrease in the mean of the quality distribution.

5.3 The value of the new AI products

We have seen an enormous increase in the number of products entering in 2024 and 2025, along with evidence that AI is increasing the quality/appeal of products at particular points in the rank distribution.

If this is true, then we should also see the AI-infused vintages attracting high shares of interest and revenue. We have two ways to explore this. The USA Today data allow us to calculate the weekly shares of bestsellers that come from recent vintages. The left panel of Figure 17 shows the share of bestsellers whose authors first appeared in the sample less that a year before the sales date. This share oscillates around about 15 percent between mid-2023, then rises above 30 percent during 2025. The share of bestsellers attributable to books published in the last year, in the right panel, rises similarly.

6 Conclusion

LLMs, which have diffused rapidly between 2023 and 2025, have become increasingly capable of writing long passages of text. Accordingly, the have seen rapid adoption by existing, and

especially new, book authors. We document a doubling in the number of new books coming to market between the start of 2023 and late 2025 that appears strongly to have been driven by use of AI tools. Large increases in the number of products can have large effects on welfare when the quality of the new works is drawn from the same distribution as existing works (so that the new works will include some high-value products). We find that the new AI-driven vintages of have low average quality; but given the large increase in the number of products – along with variance in realized quality – these vintages raise the quality of the $100^{\text{th}} - 1,200^{\text{th}}$ entrant by category. As a result, AI – like other cost-reducing technological changes in the cultural industries – has raised the value of the consumer choice set.

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.

Bang, Y. and et al. (2023): "A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning and Factuality," *arXiv preprint*.

Brynjolfsson, E., D. Li, and L. Raymond (2023): "Generative AI at Work," *Quarterly Journal of Economics*, forthcoming.

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.

Noy, S. and W. Zhang (2023): "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, 381, 187–192.

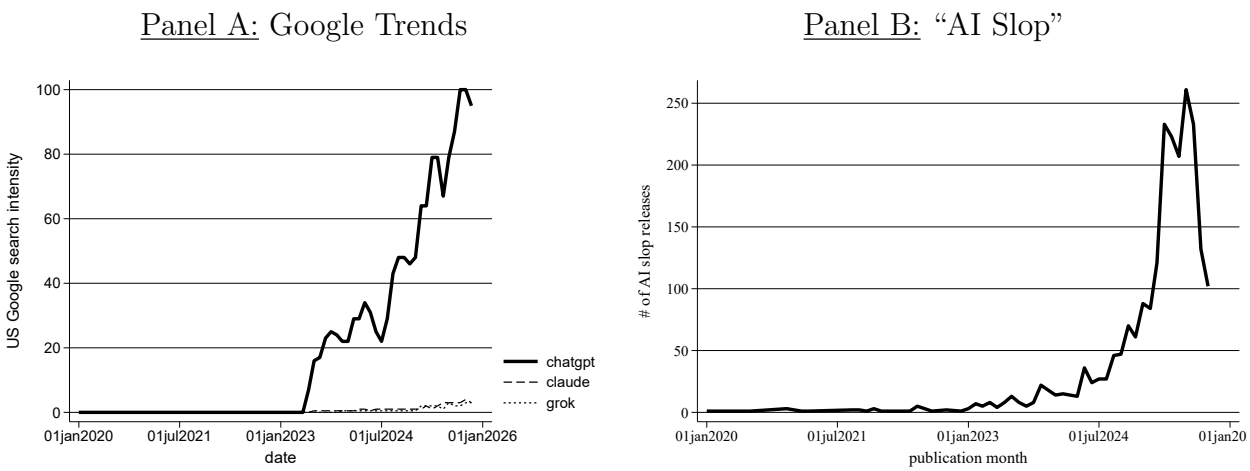
OpenAI (2022): "Introducing ChatGPT," `https://openai.com/blog/chatgpt`, accessed 2025-11-08.

Radford, A., J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever (2019): "Language Models are Unsupervised Multitask Learners," *OpenAI Technical Report*.

Waldfogel, J. and I. Reimers (2015): "Storming the gatekeepers: Digital disintermediation in the market for books," *Information economics and policy*, 31, 47–58.

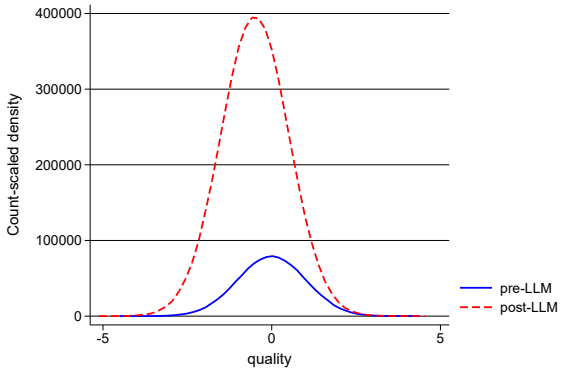
7 Figures and tables

Figure 1: The arrival of AI in books

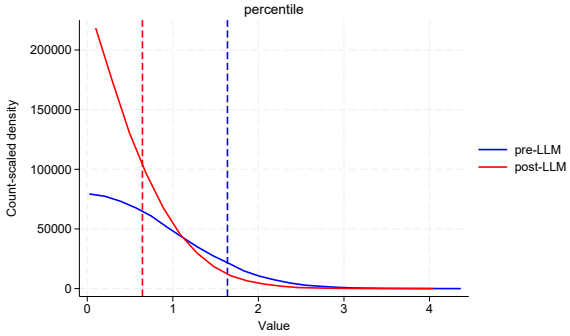


Notes: The left panel shows Google Trends data on search intensity for ChatGPT, Claude, and Grok. The right panel shows "AI slop" identified by Goodreads users.

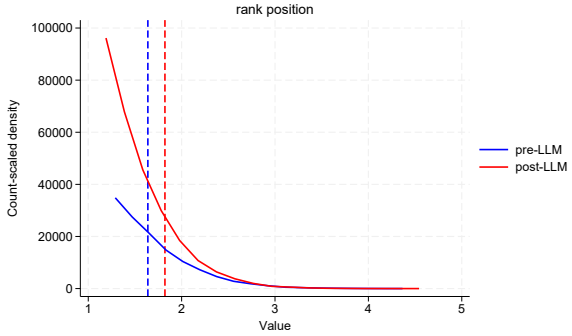
Figure 2: LLM influx and effects on rank positions vs percentiles



(a) compare distributions

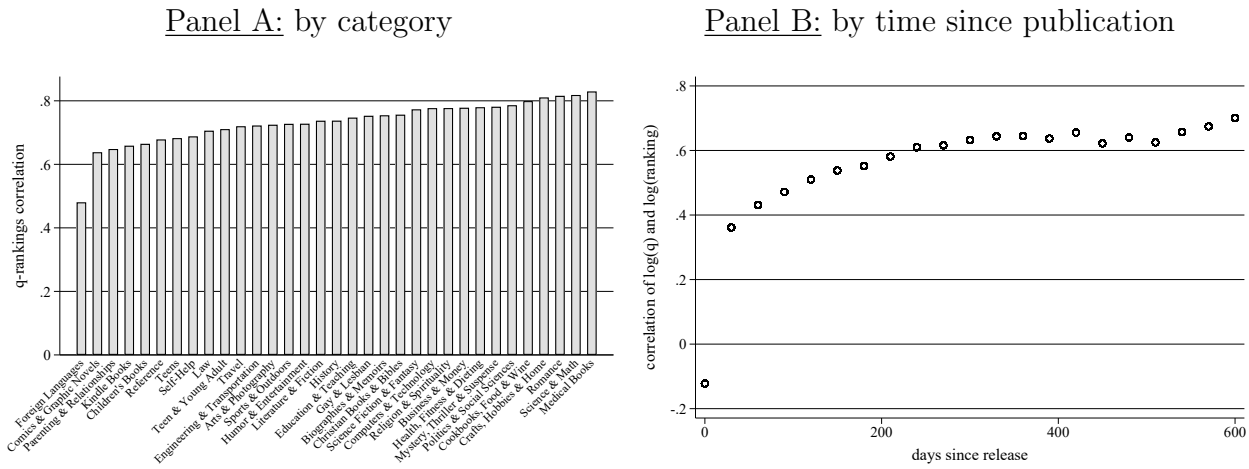


(b) percentile



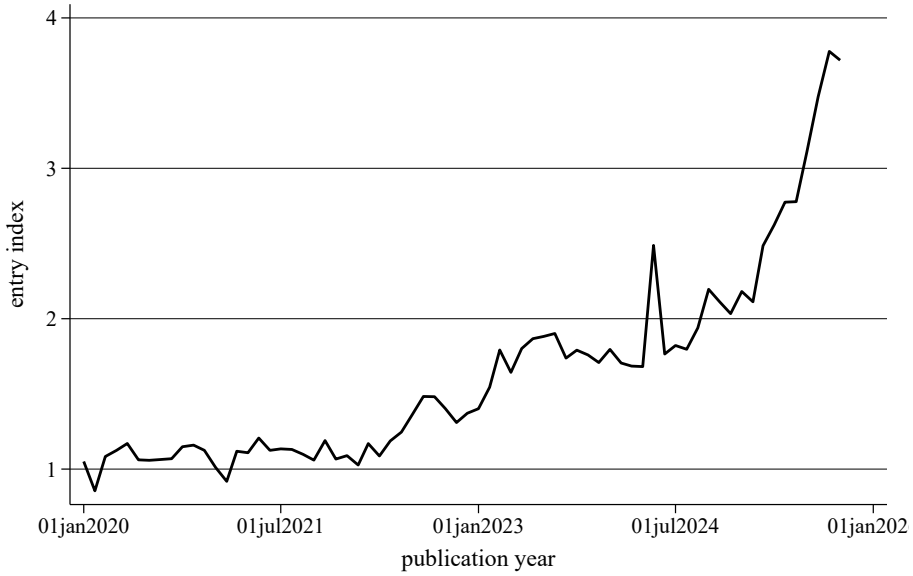
(c) rank position

Figure 3: Quantities and the number of ratings



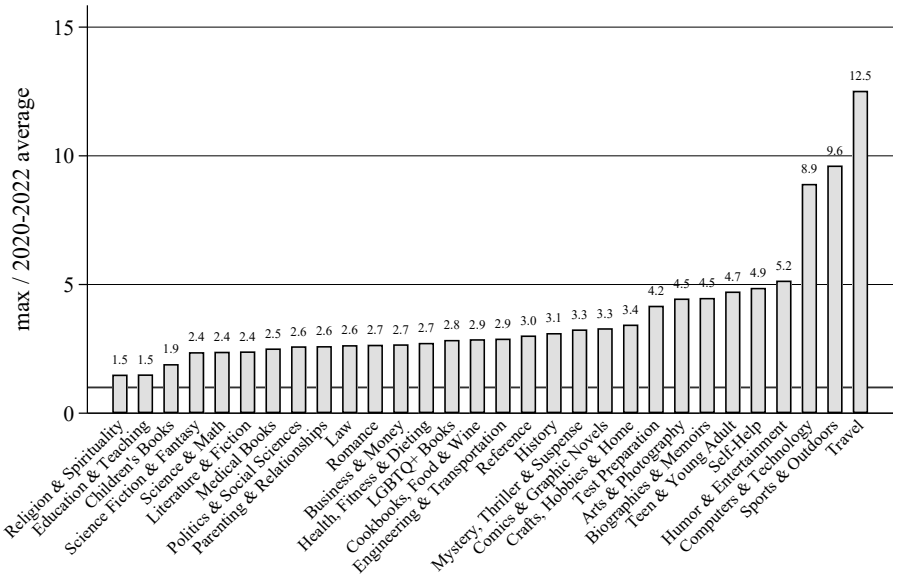
Notes: The left panel shows the correlation between log sales and the log of rankings for each book category. Included ebooks were published in January, 2018. The right panel shows the correlation as of 0, 30,...,600 days after publication.

Figure 4: Overall entry index



Note: These are exponentiated publication month fixed effects from a regression of log monthly entry, by category, on category and publication months.

Figure 5: Growth by category



Note: The bars show the ratio to the maximum monthly new publications if 2025 to the average monthly publications during 2020-2022.

Figure 6: Book entry by subcategory, 2020-2025

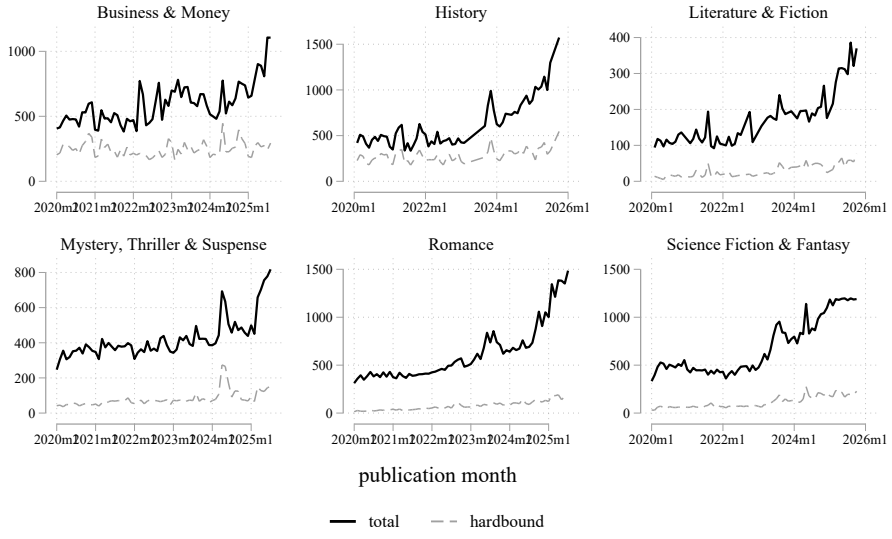


Figure 7: Book entry by subcategory, new vs existing authors, 2020-2025

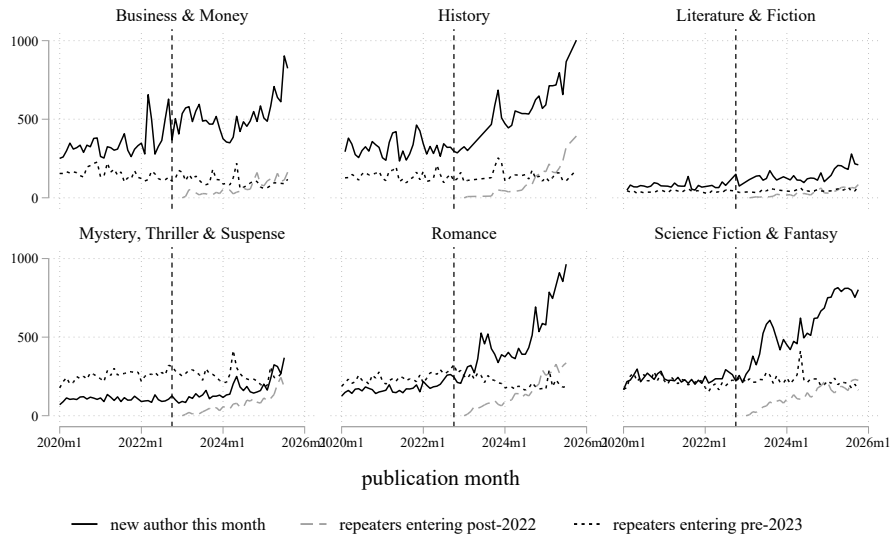
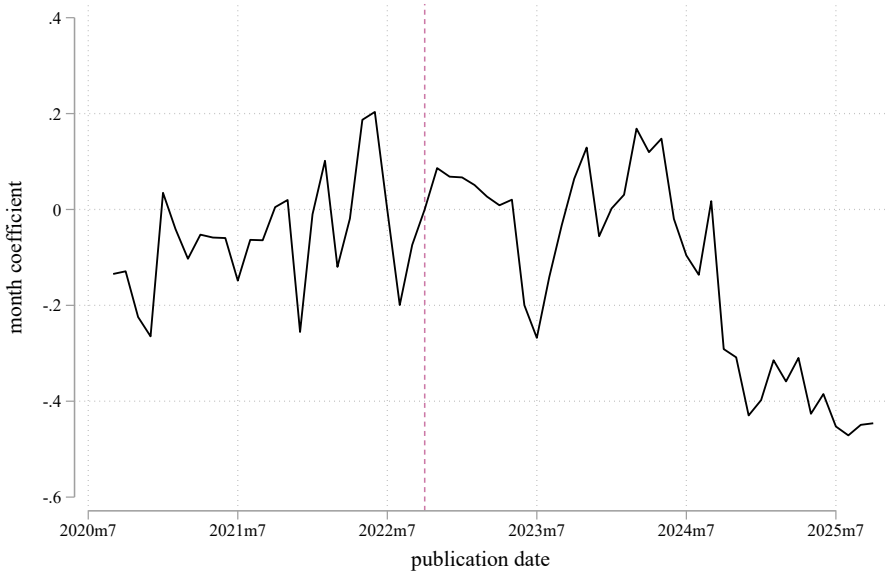
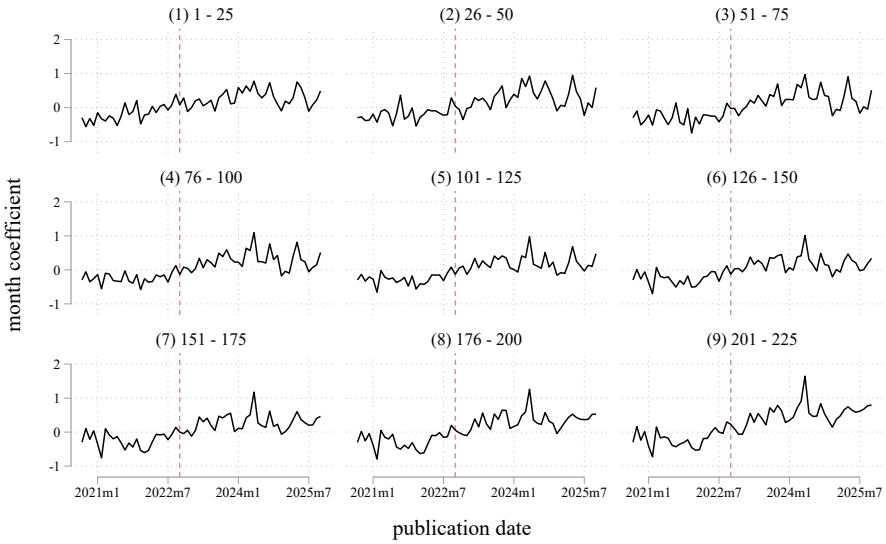


Figure 8: Number of ratings – census of 6 subcategories



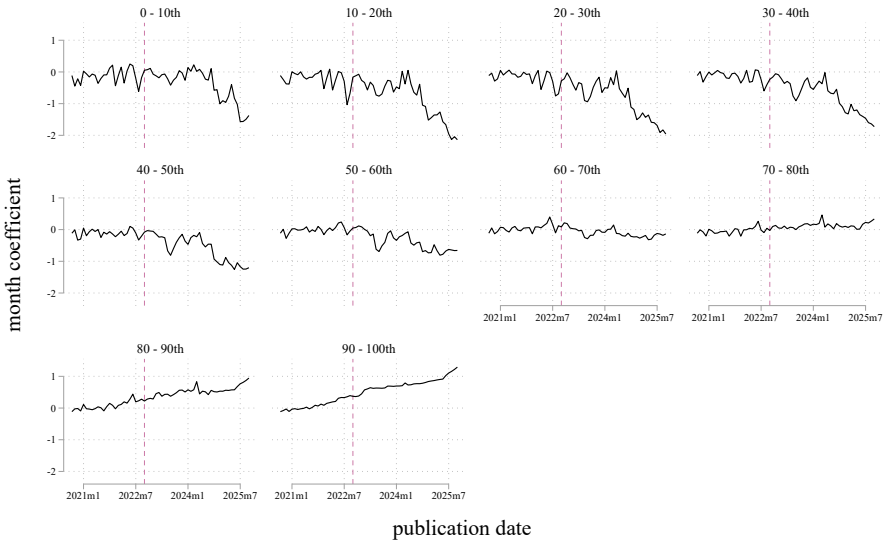
Note: These are the coefficients on month fixed effects from a regression of the inverse hyperbolic sine of the number of consumer ratings – adjusted for the time since publication – on month and category fixed effects. The underlying sample includes all books published between September 2020 and October 2025, for the six subcategories for which we observe the census of all published Kindle books.

Figure 9: Number of ratings at varying rank positions – census of 6 subcategories



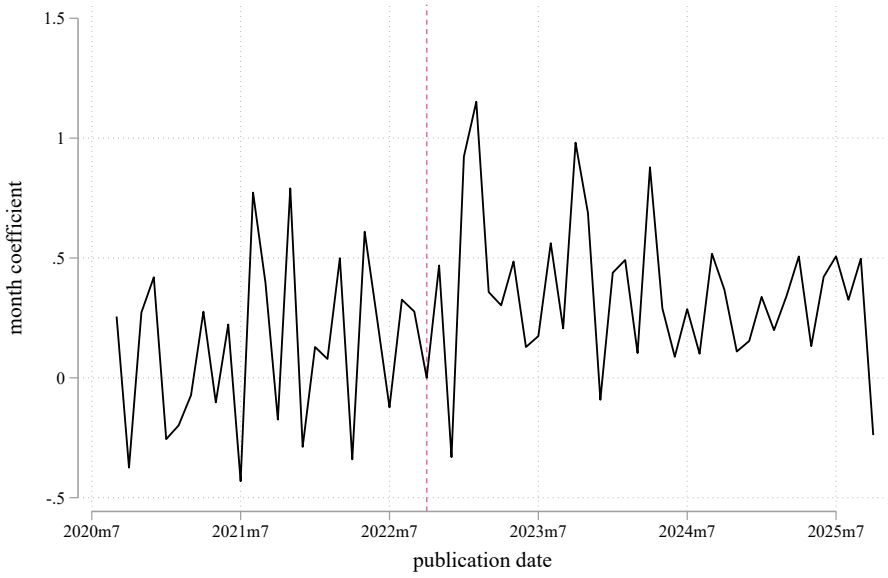
Note: These are publication month fixed effects from regressions of the inverse hyperbolic sine of the number of consumer ratings – adjusted for the time since publication – on month fixed effects along with category fixed effects and an indicator variable for whether the author is experienced, i.e. has published at least one book more than a year prior. The regressions were run separately for different rank position ranges according to the total number of reviews within a category and month. For example, Panel 1 shows estimates from regressions using only books at positions 1 – 25.

Figure 10: Number of ratings at varying distribution deciles – census of 6 subcategories



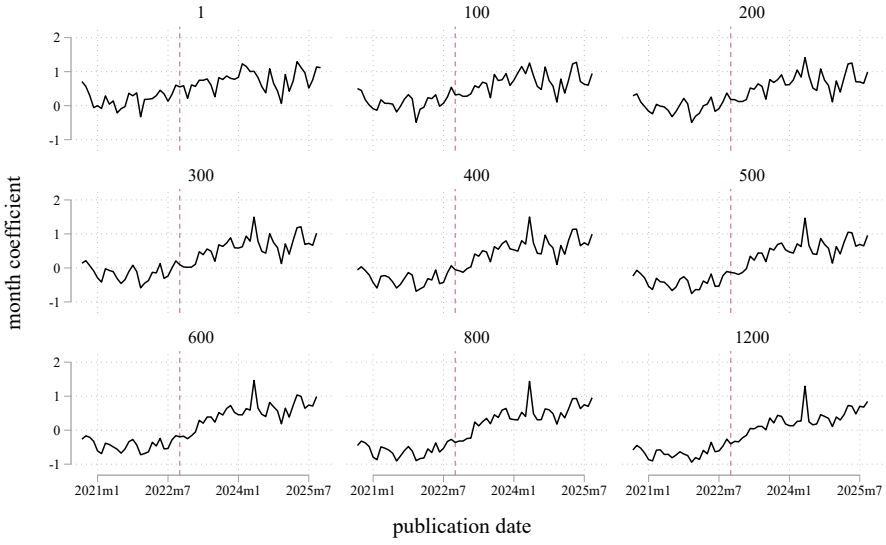
Note: These are publication month fixed effects from regressions of the inverse hyperbolic sine of the number of consumer ratings – adjusted for the time since publication – on month fixed effects along with category fixed effects and an indicator variable for whether the author is experienced, i.e. has published at least one book more than a year prior. The regressions were run separately for different percentiles of the category-month-specific ratings number distributions. For example, Panel 1 shows estimates from regressions using only books in the top 10 percentile of the respective ratings distribution.

Figure 11: Number of ratings for random books – all categories



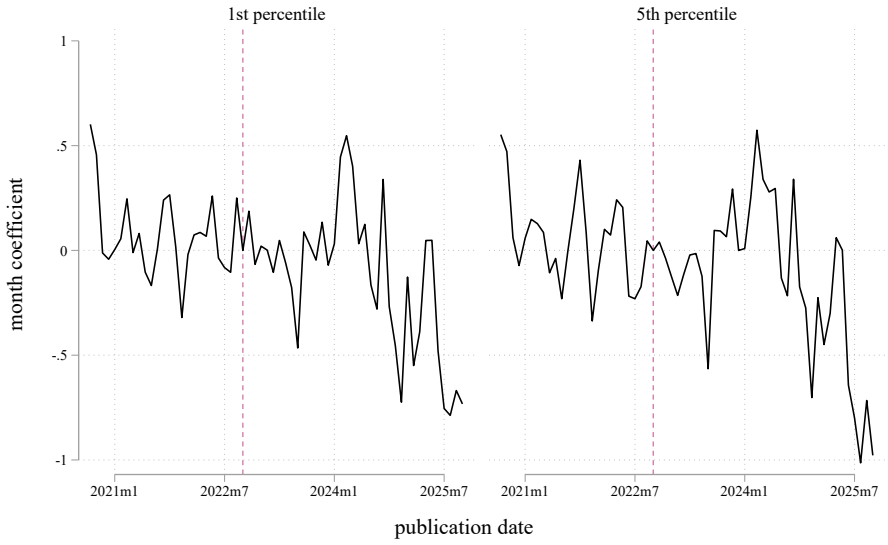
Note: These are the coefficients on month fixed effects from a regression of the inverse hyperbolic sine of the number of consumer ratings – adjusted for the time since publication – on month and category fixed effects.

Figure 12: Number of ratings at varying rank positions – all categories



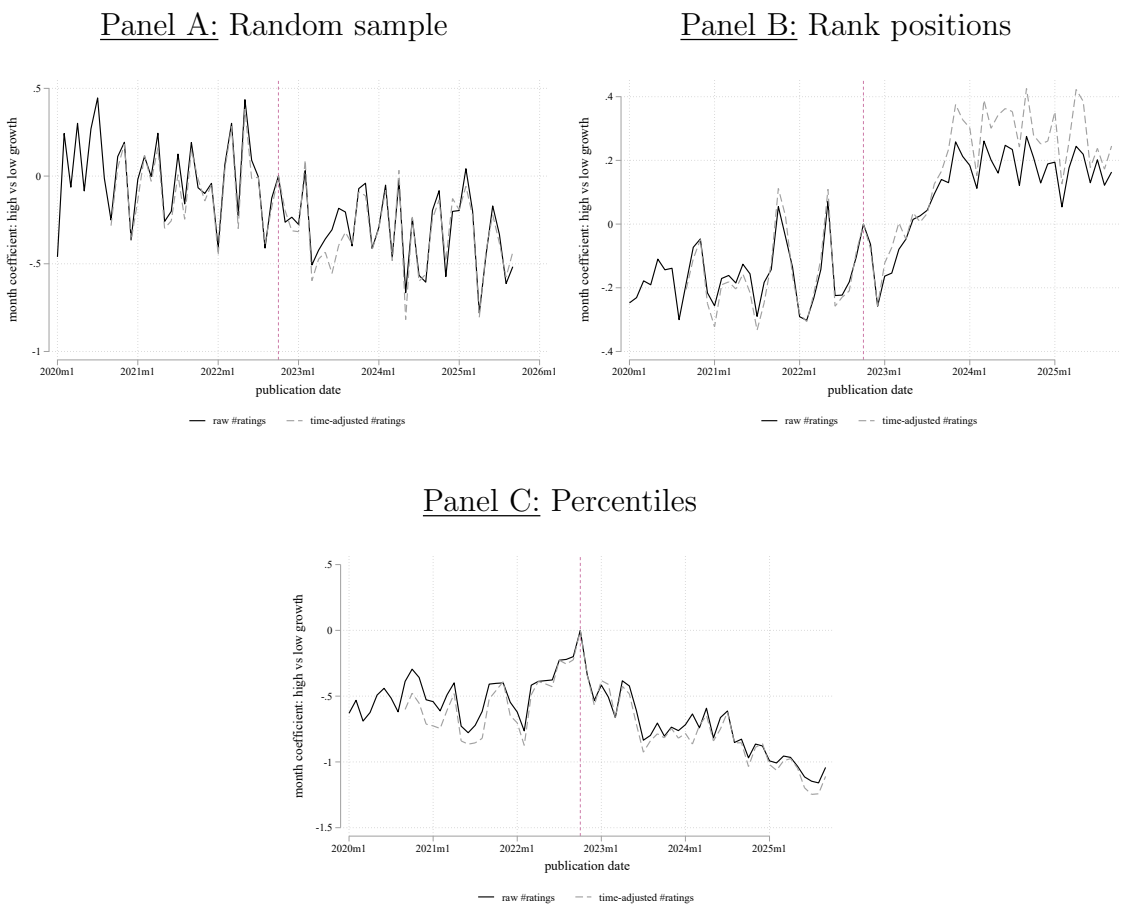
Note: These are estimates from a regression of the inverse hyperbolic sine of the number of consumer ratings – adjusted for the time since publication – on month fixed effects interacted with nine indicators for different search pages, along with category and search page fixed effects. Each panel shows the coefficients on the interactions of the month fixed effects with search page dummies. The search results were sorted in descending order of the book's number of ratings. Accordingly, the top left panel ("1") captures the 16 books of a category and published in a month that have garnered the most star ratings; the second panel ("100") includes the 97^{th} to 112^{th} most-rated books from that category and month, and so on.

Figure 13: Number of ratings at two distribution percentiles – all categories



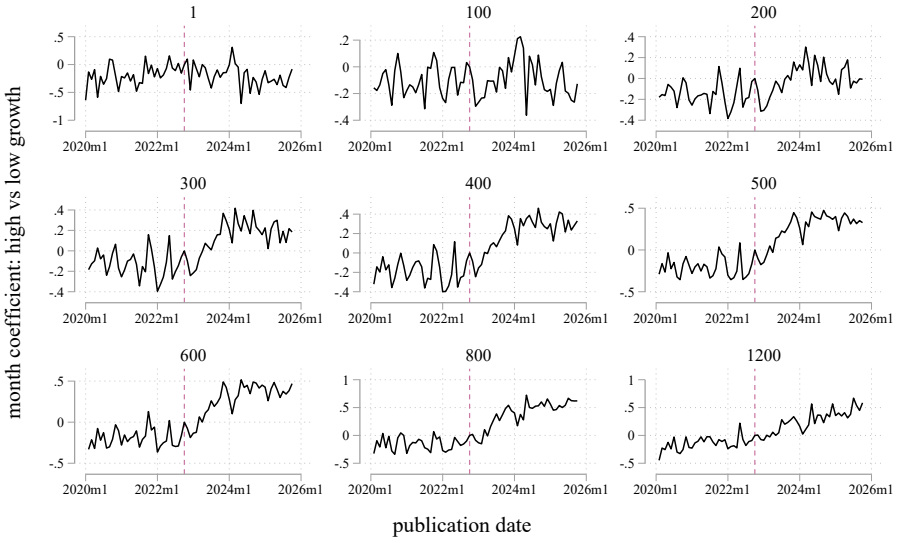
Note: These are estimates from a regression of the inverse hyperbolic sine of the number of consumer ratings – adjusted for the time since publication – on month fixed effects interacted with indicators for different percentiles of the category-month specific distributions of the number of star ratings, along with category and percentile fixed effects. Each panel shows the coefficients on the interactions of the month fixed effects with the percentile dummies. The search results were sorted in descending order of the book's number of ratings. Accordingly, the left panel ("1st percentile") captures the 16 books at the search page at the top 1^{st} percentile of the category-month specific ratings distribution, and the right panel includes those 16 books on the search page at the top 5^{th} percentile.

Figure 14: Number of ratings for high- vs low-growth categories



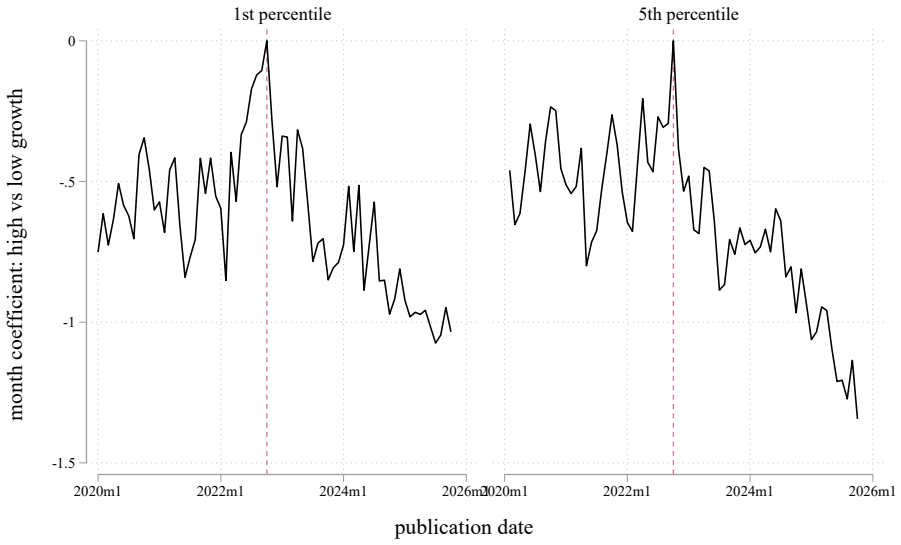
Notes: These are estimates from regressions of the inverse hyperbolic sine of the number of consumer ratings on month fixed effects, month fixed effects interacted with an indicator for book categories seeing growth in the number of books from 2022 to 2025 above the median category, and category dummies. Each panel shows the point estimates of the coefficients on the interaction terms. The solid lines show results when the ratings are not adjusted for time since publication, and the dashed, gray lines depict results when ratings are adjusted. The panels differ in the underlying sample. For example, Panel A shows that among the "random sample," the average inverse hyperbolic sine of the number of reviews decreased slightly for high-growth categories relative to low-growth categories after the introduction of ChatGPT; whereas it increased substantially for books in the "search positions sample."

Figure 15: Number of ratings for high- vs low-growth categories by search page



Note: These are estimates from a regression of the inverse hyperbolic sine of the number of consumer ratings on month fixed effects interacted with an indicator for book categories seeing growth in the number of books from 2022 to 2025 above the median category as well as a triple interaction of these terms with each observed search page. Additional controls include publication month and category-search page fixed effects. Each panel shows the coefficients on the triple interactions. The search results were sorted in descending order of the book's number of ratings. Accordingly, the top left panel ("1") captures the 16 books of a category and published in a month that have garnered the most star ratings; the second panel ("100") includes the 97^{th} to 112^{th} most-rated books from that category and month, and so on.

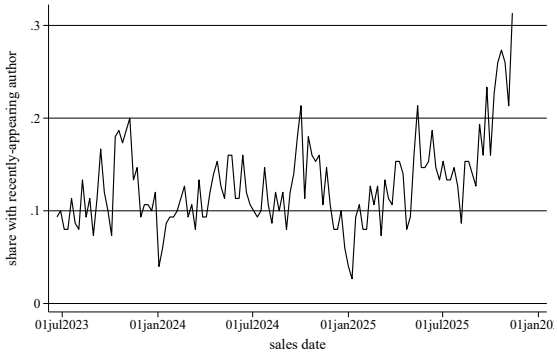
Figure 16: Number of ratings for high- vs low-growth categories by percentile



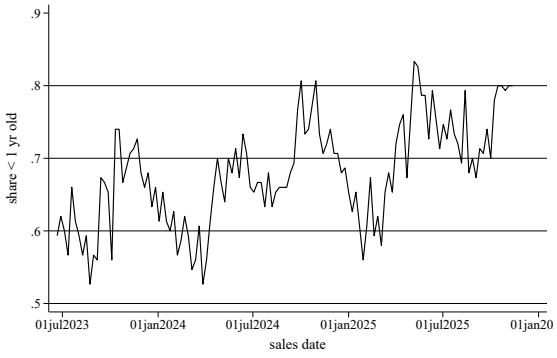
Note: These are estimates from a regression of the inverse hyperbolic sine of the number of consumer ratings on month fixed effects interacted with an indicator for book categories seeing growth in the number of books from 2022 to 2025 above the median category as well as a triple interaction of these terms with the two popularity percentiles. Additional controls include publication month and category-percentile fixed effects. Each panel shows the coefficients on the triple interactions. The search results were sorted in descending order of the book's number of ratings. Accordingly, the left panel ("1st percentile") captures the 16 books at the search page at the top 1^{st} percentile of the category-month specific ratings distribution, and the right panel includes those 16 books on the search page at the top 5^{th} percentile.

Figure 17: Success of recent books and authors

Panel A: recently-appearing authors



Panel B: recently-appearing books



Notes:

The graphs show the shares of weekly USA Today bestsellers from authors (left panel) and books (right) that appeared less than a year before the sales date.

Table 1: Summary statistics – categories data

	# books	# ratings	avg rating
Random	51,903	108.92	4.43
Rank position: 1	22,191	5,985.49	4.52
Rank position: 100	30,722	554.12	4.48
Rank position: 200	31,237	300.78	4.45
Rank position: 300	31,597	192.72	4.42
Rank position: 400	32,027	135.34	4.41
Rank position: 500	32,113	107.44	4.40
Rank position: 600	32,826	115.71	4.39
Rank position: 800	30,295	63.54	4.37
Rank position: 1200	26,066	41.31	4.45
Percentile: 1^{st}	27,846	1,599.73	4.52
Percentile: 5^{th}	30,316	218.33	4.49

Note: Descriptions of the categories samples. The "Random" panel describes the random samples consisting of two pages (32 observations) of Amazon search results for each category and publication month. The "Rank position" panel describes the sample with one page (16 books) of search results for each category and publication month. We obtain the page with 16 search results nearest to, or surrounding, the rank position (according to the number of ratings a book has received). The "Percentile" panel describes the samples based on one page in the first percentile of results and a page in the 5th percentile of search results (according to the number of ratings).

Table 2: Summary statistics – census of six subcategories

	# books	# ratings	avg rating
Business & Money	94,954	38.66	4.33
History	92,591	71.72	4.37
Literature & Fiction	19,841	83.82	4.29
Mystery, Thriller & Suspense	56,105	537.34	4.35
Romance	73,197	315.91	4.36
Science Fiction & Fantasy	90,172	156.16	4.36

Note: The table describes the "census" data, the complete list of books published in selected subcategories of the listed categories, Jan 2008 - October 2025. The particular subcategories are economics, history (world), alternative history, mystery (women sleuths), romantasy, and science fiction adventure. Subcategories are chosen to be small enough to make data collection feasible. See the text.

Table 3: Book entry and average quality measures

	(1) I(rated)	(2) asinh(R)	(3) stars
log N	-0.0676***	-0.298***	-0.140**
	(0.0191)	(0.0903)	(0.0572)
Observations	2100	2100	2098

Note: Dependent variables are the shares of books receiving ratings, in Column (1), the inverse hyperbolic sine of $\#ratings$ in (2), and the average star rating in (3). The explanatory variable of interest is the number of books published in that category and month. All specifications include time and category fixed effects. Monthly data, January 2020-October 2025, for 30 Amazon book categories. The underlying dataset is the "random sample."

Table 4: Book entry and search rank specific quality measures

	(1) I(rated)	(2) asinh(R)	(3) stars
rank pos=1 × log N	-0.00754	0.0453	-0.0121
	(0.00527)	(0.113)	(0.0173)
rank pos=100 × log N	-0.00207	0.261***	0.0542**
	(0.00238)	(0.0633)	(0.0249)
rank pos=200 × log N	0.00458	0.411***	0.0477*
	(0.00295)	(0.0715)	(0.0269)
rank pos=300 × log N	0.0431	0.559***	0.0584*
	(0.0290)	(0.0919)	(0.0316)
rank pos=400 × log N	0.145**	0.732***	-0.0313
	(0.0595)	(0.134)	(0.0371)
rank pos=500 × log N	0.318***	0.832***	-0.0341
	(0.101)	(0.131)	(0.0262)
rank pos=600 × log N	0.350***	0.773***	-0.0485
	(0.106)	(0.116)	(0.0370)
rank pos=800 × log N	0.429***	0.846***	-0.0515
	(0.0865)	(0.150)	(0.0626)
rank pos=1200 × log N	0.286**	0.702***	0.0434
	(0.108)	(0.167)	(0.0405)
Observations	18273	18273	16293

Note: Dependent variables are the shares of books receiving ratings, in Column (1), the inverse hyperbolic sine of $\#ratings$ in (2), and the average star rating in (3). The explanatory variables of interest are the number of books published in that category and month, interacted with the search page. All specifications include rank-position-specific time and category fixed effects. Monthly data, January 2020-October 2025, for 30 Amazon book categories. The underlying dataset is the "search positions sample."

Table 5: Book entry and average and percentile quality measures

	(1) I(rated)	(2) asinh(R)	(3) stars
percentile=1 × log N	-0.000552	-1.017***	-0.0169
	(0.00105)	(0.0998)	(0.0202)
percentile=5 × log N	-0.00112	-1.181***	-0.0289
	(0.000776)	(0.108)	(0.0250)
Observations	4108	4108	4108

Note: Dependent variables are the shares of books receiving ratings, in Column (1), the inverse hyperbolic sine of *#ratings* in (2), and the average star rating in (3). The explanatory variables of interest are the number of books published in that category and month, interacted with the percentile. All specifications include percentile-specific time and category effects. Monthly data, January 2020-October 2025, for 30 Amazon book categories. The underlying dataset is the "percentile sample."