

NBER WORKING PAPER SERIES

AI AND THE QUANTITY AND QUALITY OF CREATIVE PRODUCTS:
HAVE LLMS BOOSTED CREATION OF VALUABLE BOOKS?

Imke Reimers
Joel Waldfogel

Working Paper 34777
<http://www.nber.org/papers/w34777>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

We received no funding for this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Imke Reimers and Joel Waldfogel. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Imke Reimers and Joel Waldfogel

NBER Working Paper No. 34777

January 2026

JEL No. L16, L82

ABSTRACT

With the diffusion of LLMs between 2022 and 2025, new book releases have tripled, raising a question of AI's impact on book quality. We develop a ratings-based usage measure that is comparable across book release vintages, and we find that the vintages from the AI influx period have lower average quality. Yet, the top 1,000 monthly releases per category - albeit not the top 100 - have higher quality than before; and the effect is larger in categories with faster growth in new titles. Authors entering since the LLM influx produce predominantly low-quality work; and the higher-quality output of pre-LLM authors entrants has risen. A nested logit calibration shows that LLM-enhanced book production could, in steady state, raise the surplus that consumers derive from book markets by a quarter to a half.

Imke Reimers

Cornell University

SC Johnson College of Business

imke.reimers@cornell.edu

Joel Waldfogel

University of Minnesota

Carlson School of Management

and NBER

jwaldfog@umn.edu

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 may reduce the cost of producing books; and reductions in entry costs can raise the number of new products brought to market. Additional entry can be especially beneficial when product quality is unpredictable at entry, so that some new products have realized quality in, or near, the right tail of the quality distribution ([Aguiar and Waldfogel, 2018](#)).

LLMs may well reduce the cost of creating books, but their effect does not simply facilitate more “draws from the urn;” the LLMs may change the distribution of quality realizations, for example by granting writing ability to those who lack it while complementing the ability of skilled writers. The effect of AI on the benefit that consumers derive from the products depends on what it does to both the *number* of new books and their *quality distribution*, and LLMs could have effects throughout the quality distribution. For example, if LLMs allow the creation of books with low average quality, and little quality variance, they will deliver no books in the right tail of the quality distribution and little benefit to consumers. On the other hand, if LLMs facilitate the creation of more books drawn from the pre-LLM quality distribution, then the number of high-quality new releases will rise (e.g. the n^{th} -best release in a month would be better than before), with potentially substantial welfare benefits.

We create two datasets for analyzing the effect of LLM availability on new book releases and quality. First, we assemble a stratified random sample of over 333,000 releases that are representative of the 10 million ebooks available at Amazon that were released between 2020 and 2025. We observe the number of persons who had rated each book as of late 2025 – our

main usage/quality measure – as well as each book’s average star rating. With population weighting, the stratified random sample allows us to draw inferences about the effect of the LLM influx on the distribution of new book qualities. Second, we collect data on 479,000 books released in eight subcategories between 2008 and 2025. The subcategory sample, while not a representative sample of releases generally, includes all releases in eight topic areas and allows us to follow author careers, for example to see whether incumbent authors become more productive, or are displaced, as LLMs diffuse.

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 nearly tripled between 2022 and late 2025 and rose by a factor of ten in some categories; and the increase in new titles 2022-2025 coincides with the diffusion of LLMs.

We then measure the impact of LLMs on the quality (appeal) of new books in three ways. First, using observations on books’ numbers of ratings at two points in time roughly a month apart, we translate each book’s number of ratings into a measure that is comparable across books of different vintages. We use this adjusted usage measure to document changes in quality across release vintages following the introduction of AI tools for writing. Second, we use a difference in differences approach to show that the effects we document are larger in faster-growing categories – presumably experiencing larger LLM impacts – relative to more slowly-growing categories. Third, we show how effects on quality vary with the continuous variation in new book entry across months and categories. Using each of these approaches, we document that the LLM-affected vintages have lower average quality and lower quality at particular percentiles of the distribution. At the same time, given the dispersion in realized quality, the large increase in releases raises the absolute number of moderately high-quality releases.

While the quality of the top 100 monthly releases per month and category is unaffected, the next thousand releases, by category and month, from LLM-heavy vintages are of higher

quality than before. This can raise welfare substantially: Using a simple calibrated nested logit demand model, we estimate that LLMs could raise consumer surplus by a quarter to a half in steady state, driven by the large increase in the number of new products.

The increase in book releases reflects both increases in new author entry and repeat entry by incumbent authors, which in part reflects an accelerated pace of new releases from authors who debuted before the LLM influx and who tend to produce high-quality works. Authors entering in the LLM era, by contrast, disproportionately produce low-quality work.

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, author use of LLMs, the evolution of the book market, and the related 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 the quality distribution. Section 3 describes the data used in the study. Section 4 describes effects of LLMs on new book entry and on the productivity of incumbent authors. Section 5 describes effects of the LLM influx on the distribution of entering book quality, as well as descriptive facts about the book and authors in the AI influx. Section 6 presents an explicit welfare analysis using a calibrated nested logit demand model. A brief conclusion follows.

1 Background

1.1 Large language models and author engagement

The development and diffusion of large language models in the past few years has held out the possibility of changing the book-writing process. 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.” This is reflected in Google search behavior. 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.

LLM usage is not merely a theoretical possibility for writers; surveys indicate that LLMs are both attracting new writers and being used by existing writers. 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 were planning to use, AI. The remaining 7 percent indicated they might use AI in the future.¹ Among those not using AI, more than three quarters thought its use to be “unethical.” Among authors using AI, ChatGPT was used by 85 percent. Just over a third of respondents had published through traditional publishers; the remainder had used only 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.²

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.³ 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.⁴ Some authors have apparently achieved success using AI. Tim Boucher has written hundreds of books using AI, and Time reports

¹See <https://insights.bookbub.com/how-authors-are-thinking-about-ai-survey/>.

²See <https://www.publishersweekly.com/pw/by-topic/digital/copyright/article/99019-new-report-examines-writers-attitudes-toward-ai.html>.

³See <https://blog.bookautoai.com/ai-book-generator-genre-suitability/>.

⁴“The best genres for AI-generated books include romance, romantasy, thrillers, and cross-genre fiction, as they have strong reader loyalty and predictable structures.” See <https://blog.bookautoai.com/best-genres-ai-books-2025/>.

thousands in his resulting earnings.⁵ At the same time, there is a great deal of low-quality work, which critics derisively term “AI-slop.” Goodreads has many user-created lists of slop.⁶ 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 2022-2025 period.

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 2008, including the Kindle device and eventually the Kindle Direct Publishing platform, many books came to be 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 Kindle. Many works initially appearing as self-published books became best sellers. Self-publishing, the main vehicle for LLM-aided books post-2022, had already been legitimized as an avenue to commercial success.

Measuring the number of new books is not easy in the digital era. Traditional publishers sought copyright registration and obtained international standard book numbers (ISBNs). Individuals publishing outside of the traditional publishers often do neither. Between 2020

⁵See <https://www.newsweek.com/ai-books-art-money-artificial-intelligence-1799923>.

⁶See <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_.

and 2024 the number of annual nondramatic literary registrations in the US copyright data went from 152 thousand to 156 thousand, albeit with fluctuation between 129 and 189, but showing no evidence of an LLM-induced influx.⁷ The number of ISBNs issued per year, from Bowker’s Books in Print, also shows no evidence of LLMs; it falls from 2.74 million new books in 2023 to 2.37 million in 2024. Hence, the influx we document below is largely a digital – and an Amazon-centered – phenomenon.

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, [Waldfoegel 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 Waldfoegel, 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. Experiments 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

⁷See, for example, <https://www.copyright.gov/reports/annual/2020/ar2020.pdf>.

documents improvements in fluency and coherence but mixed effects on novelty or 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). Together, these studies raise the possibility that AI may significantly increase new book production but leave open the exact nature of effects on quality. Kusumegi et al. (2025) document that academics adopting LLM produce substantially more working papers.

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 marketplace data, they show that generative-AI tools increase entry and product variety while sharply displacing output of non-AI images.

2 Theory

Creating a book is a product 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. AI tools have reduced the cost of book creation, and cost reduction would be expected to stimulate additional creation.

The quality 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 digitization did with media products (Aguilar and Waldfogel, 2018). The arrival of AI is not simply a reduction in the cost of entry, however. 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 quality but higher than what they could have created in its absence. Hence, LLMs can change not only the volume of creative entry but can also create changes throughout the realized product quality distribution. These concerns motivate the paper’s main question: What happens to the quality distribution – and aggregate value – of 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 and, in particular the number of products attracting substantial usage. A simple model of the realized distributions illustrates how AI tools might affect the distribution of product qualities and the value of the choice set. Define δ_j as the mean utility, or “quality” of product j .⁸ 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. Define N_k as the number of products introduced in each period. If $\{\delta_j\}$ is the set of available products, the consumer surplus associated with the choice set can be calculated. For example, per capita consumer surplus (CS) is proportional to $\log(1 + \sum e^{\delta_j})$ in the plain logit model.

Some cases are helpful for organizing thinking. First, suppose that N rises while the distribution from which qualities are drawn – μ and σ in the normal case – remain constant. Then the average quality would remain the same, and the quality of a product at any particular percentile (e.g. the 99th) would remain constant. 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 quality. As a result, the quality of the k^{th} best product would rise. That is, the 100th-best product would be better when 2,000

⁸In the plain logit model, $\delta_j = \log(q_j/M) - \log((M - \sum q_j)/M)$, where q_j is the quantity sold for book j and M is market size. We return to this approach in Section 6.

products entered than when 1,000 products entered. The increase in entry would, of course, raise the value of the choice set.

For a second case, suppose that μ declines while N and σ remain constant. Then – comparing the new vs the old cohort – the average quality would fall, as would the quality at any percentile of the distribution. Moreover, the k^{th} -best new product’s quality would also fall. A choice set produced in this way, rather than with the prior distribution, would have lower value. Finally, consider a third – and potentially more relevant case – when N rises while μ declines. Then effects are slightly more nuanced: The average quality would, of course, decline, as would the quality at any particular percentile of the distribution; but the quality of the k^{th} -best product might rise or fall. Depending on the magnitude of the increase in entry relative to the fall in average appeal, the value of the choice set could be higher or lower.

The third case 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 at any percentile of the distribution 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. An increase in the number of new products in, or near, the right tail of the quality distribution is important, as it can substantially raise the value of the choice set to consumers.⁹

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

⁹Of course, any increase in the number of available products raises the value of the choice set, but CS is little affected if all of the additional products have low quality.

3 Data

We are interested in characterizing volumes of new books entering over time – including time periods before and after the influx of AI tools that might affect authorship – as well as the distributions of those books’ usage and therefore “quality.” 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 perceived quality such as star ratings). Moreover, we would also 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 usage (the cumulative number of ratings received) and satisfaction (average star ratings) for a representative sample of ebooks available at Amazon and released between 2020 and 2025, as well as the entire populations of ebooks released in eight smaller subcategories between 2008 and 2025.

3.1 Data sources and collection challenges

We rely on Amazon for 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 (ebook) 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 the exact number of releases in the category and release month. Up to 10,000, the number of search results is reported to the nearest thousand, 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 titles in the result. Hence, the number of books in a search is exact for 1-999, to

the nearest 16 for 1,000-6,400, to the nearest 1,000 up to 10,000, and to the nearest 10,000 for higher numbers. Subject to these limitations, 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 150 pages constrains the subcategories for which we can obtain a complete census of book-level data. In particular, the group of books chosen cannot include more than 2,400 titles published per month ($=150 \times 16$). Given the ways that we can search, we create four datasets.

Aggregate numbers: First, we have aggregate monthly time series on the numbers of new works published each month for 30 Amazon categories between 2008 and 2025.¹⁰ Seven of these 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 monthly subcategory totals by month to create the aggregate time series of new books released by category.

Random sample: The second dataset – and the main dataset for our analysis – is a random sample of books from all categories and publication months, 2020-2025. The aggregate

¹⁰The 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.

series show ten million releases between the start of 2020 and the end of 2025. We assemble a list of all possible pages using the information above on the number of titles released per category and month. We choose 10 random pages per month and category.

There is one complication. Because only the first 75 and last 75 pages for each monthly category-specific search are accessible to us, our sampling scheme – sorting by publication date – effectively allows us to access only books published near the beginning or the end of the month in large categories such as Literature & Fiction. Hence, we sample differently in the seven largest categories. Using the numbers of books released per month in each of their subcategories, we assemble lists of all search result pages in these subcategories. We then randomly sample 10 pages per month in each of these seven categories from the comprehensive subcategory page lists. This approach delivers books published throughout the month, even in the large categories.

Our random sampling approach obtains 160 random titles within each category and month. That is, our sample includes equal numbers of books in each category and month, even though the number of books released differs across categories and is changing over time. The aggregate category totals give rise to weights that allow calculation of statistics representative of the ten million books released 2020-2025. We use the random sample to explore the effects of the AI influx on the distribution of book qualities. The random sample allows us to explore not only effects on average quality but also effects at both percentiles and absolute rank positions in the distribution.

Table 1 summarizes the random category sample, weighted to reflect the population. The average number of ratings is 112.95, and the average star rating (among the books with ratings) is 4.44. The average price is \$9.70, although 51 percent of titles can be borrowed by Kindle Unlimited subscribers at a zero marginal price. The second panel of the table reports averages for rank position ranges 1-100, 101-200, and so forth. The top 100 have an average of 2,320.91 ratings and an average star rating of 4.49. By construction, the number of ratings

falls across these rank position ranges, to 830.79 for 101-200, to 407.73 for 201-300, and so forth.¹¹ The average star rating also declines across these sales rank ranges, until 701-800. Prices are higher on average for lower-ranked books; but the tendency to be included in Kindle Unlimited rises slightly. The bottom panel reports averages at percentile ranges of the monthly category-level distributions. The number of ratings averages 1,012.68 in the top decile (i.e.90-100), 75.30 in the second, 24.09 in the third, and so on. As with the rank position comparison, average star ratings fall from the top to the third decile.

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 third dataset, which includes a complete “census” of new titles in eight subcategories that are sufficiently small to allow their collection, including romantasy (a subcategory of Romance) economics (a subcategory of Business & Money), fantasy (Science Fiction & Fantasy), women sleuths (Mystery), alternative history (Literature & Fiction), world history (History), US travel (Travel), and sports biography (Sports & Outdoors). 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. We collect all of the titles in these categories released between January 2008 and December 2025. These data are described in Table 2. The census sample contains 479,648 books in total. The census sample differs from the representative category sample. These books have more ratings, lower average star ratings, higher prices, and a lower tendency to be included in Kindle Unlimited, though values vary.

Rank position data: We create a fourth dataset in which we sample specifically at rank positions. That is, we choose the pages, within a category and month, with 16 entries that cover the rank positions at 1, 100, 200, 300, . . . , 600, 800, and 1,200. We put this additional dataset to two uses. First, we use it to accurately measure the number of ratings near each

¹¹Different category-months have different numbers of releases, so monotonicity eventually breaks down.

rank position cutoff. Second, we use it as a supplementary dataset for measuring the change in book quality, conditional on rank position, as the number of releases rises. This dataset, which covers releases during 2020-2025 in all thirty categories, has 263,872 observations.

3.2 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. We would like to have 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 30 days apart. 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

$$r_{j,\tau+k} = r_{j,\tau} e^{g(\tau)k}, \tag{1}$$

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

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 60 months). We do this as follows: $r_{jT} = r_j e^{\sum_{i=\tau}^T g(i)}$. For notational simplicity we refer to our adjusted measure as \tilde{r}_j .

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. Panel (a) of Figure

[A.1](#) shows the average time pattern of ratings growth, among books with positive ratings, normalized to one at 60 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. Once a book has positive ratings, we adjust the ratings forward as above. We refer to the adjusted number of ratings as \tilde{r}_j .

In effect, relatively few books observed with no ratings shortly after publication will still have zero ratings at month 60, while books observed with zero ratings longer after publication are more likely to have no ratings at month 60. See Panel (b) of [Figure A.1](#).

3.3 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_{jt}^τ as the number of cumulative ratings as of data collection date τ 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 April 2017, and in the Januaries of 2018-2020, we calculate the estimated cumulative sales of each title by month (τ) until mid-2022. [Figure A.2](#) 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. This indicates that our usage measure is good for books published prior to a month before the end of the sample period.

we also use these data to describe the relationship between cumulative sales and ratings for each title. Using all of the dates for each title, a regression of the log of cumulative sales on the log of ratings yields a constant of 3.448 ($se = 0.001$) and a ratings coefficient of 0.667 (0.0002). We use this relationship below at Section 6.

4 LLMs and book entry

This section documents new book entry and the productivity of incumbent authors. Section 4.1 discusses examples of prolific authors, whose output appears to have been increased by LLM use. Section 4.2 describes overall entry patterns before and after the advent of LLMs, by new and existing authors.

4.1 LLMs and prolific authors

The volume of new works by the most prolific authors suggests LLM involvement. Figure 4 shows the annual output of the most prolific authors in our census sample over time. The solid line shows the annual releases, among books in our census sample, of the tenth-most prolific author for that year. The tenth most prolific author (in each year) produced about 20 works per year between 2015 and 2022. The annual output rate then grew past 100 for 2025. We see similar, if less dramatic, increases for authors at the 25th, 50th, and 100th positions. The jump in output at these points in the productivity distribution are suggestive of AI use.

The census dataset contains 36 authors whose first work appeared in 2024 or 2025 and who had 50 or more book releases in the sample. For example, one author has 456 titles in

our sample; another author has 172 sample titles during 2025 alone. Quantity of releases and effects on usage – and therefore welfare – are not the same. The works by the author with 456 sample titles had collectively been rated 37 times, and the works by the other prolific authors had collectively amassed no ratings as of early 2026.¹²

4.2 Aggregate book release patterns

The left panel of Figure 3 shows the new titles by month for the 30 broad Amazon categories between January 2020 and December 2025. The figure shows the total book releases, based on the sum of the category monthly totals (where the totals for the seven largest categories are aggregated from the subcategory totals). The number of new titles is roughly stable at 100,000 per month between 2020 and mid-2022 and rises during 2023 and 2024 and then rises more sharply in 2025, peaking at over 300,000 per month in late 2025. Over ten million titles are released as ebooks during the six-year period.

The growth in the rate of new publications is large across most categories. As the right panel of Figure 3 shows, the highest monthly total during 2025 is as much as 9.1 times as high as the highest monthly figure for 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.

The time pattern of entry in the subcategory sample is similar to the aggregate data in Figure 3. The large increases in the numbers of books appearing in these subcategories, as well as books overall, coincide with the growth in attention to ChatGPT in Figure 1.

¹²Google searches of the 456-title author’s name and “AI slop” reveal Amazon reviews surmising the work to be AI.

5 Quality effects

It is clear from the foregoing section that LLMs can make documents that are being offered for sale as books. The more interesting question, explored in this section, is whether the presence of LLMs has facilitated the creation of valuable books. What does the presence of LLMs do to the distribution of book qualities produced?

We employ three broad empirical strategies for measuring the impact of the AI influx on the quality distribution. These include a purely intertemporal comparison and two cross-genre intertemporal comparisons. We describe these – and their advantages and disadvantages – below. Section 5.1 presents a simple temporal approach. Section 5.2 presents a difference in differences comparison of quality in high entry growth vs low-growth categories. Section 5.3 presents a continuous approach based on the relationship between the quality distribution and the number of monthly releases in the category.

5.1 Temporal approach

LLMs became available over time, so one important source of variation is temporal. As LLMs diffused to authors – and the number of new works rose – what happened to the quality distributions of the new works created? While this approach – unlike the two that follow – lacks a control group, the abrupt growth in new titles makes the raw time pattern worthy of attention. Another challenge with the purely temporal approach is that our main outcome measure, the cumulative number of ratings received since release, grows over time within title, so older titles have higher cumulative usage measures even if they are not better. We address this with our adjusted cumulative rating measure \tilde{r}_j discussed above.

We are interested in how the distribution of adjusted qualities evolves over time. We begin by exploring the effect on average quality via the following regression, estimated on the random sample. Because the adjusted ratings measure contains zeroes (22 percent of obser-

vations), we use the inverse hyperbolic sine of adjusted ratings as our regression dependent variable in:

$$\operatorname{arcsinh}(\tilde{r}_j) = \mu_c + \psi_t + \varepsilon_j. \quad (2)$$

where μ_c is a category fixed effect, ψ_t is a time effect, 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. Figure 5 shows the resulting estimate of time effects ψ_t . Panel (a) shows that between 2020 and early 2022 average quality is stable, then it falls noticeably after late 2022.

We measure effects on quality at particular percentiles of the distribution by augmenting Equation (2) to include fixed effects for within-category percentile deciles. Panel (b) of Figure 5 shows the within-percentile results. Like the overall average, the quality conditional rank percentiles is stable prior to late 2022, then falls.

We measure effects on quality at particular rank positions by augmenting Equation (2) to include fixed effects for within-category rank ranges 1-100, 101-200, etc. Finally, Panel (c) shows the evolution of quality conditional on rank positions for the top 2000 by category and month. In contrast to the quality at the mean and at percentiles, the quality at particular rank positions rises.¹³

Using the temporal approach, the average quality of entering books falls with with the arrival of LLMs. Hence, the books at any percentile of the distribution are also worse. However, the large increase in the number of books released raises the quality of books at rank positions in the right tail of the distribution.

¹³We obtain very similar results using the rank position sample. See Panel (a) of Figure A.3.

5.2 Cross-category intertemporal approach

Different categories of books may be amenable, to different extents, to the effects of LLM availability. We exploit this, and a difference in difference approach, to ask 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 exposure to ratings across publication dates. As a result, this approach can be implemented with either adjusted or unadjusted usage measures (\tilde{r}_j or r_j). Moreover, we can compare results based on adjusted and unadjusted measures to explore the role of usage adjustment in our study.

We divide the 30 book categories into two groups, those highly exposed to AI and those less exposed. We then ask how the usage/quality measures change across release dates for the highly exposed groups, relative to the less exposed groups. That is, we run regressions of the form

$$y_j = \mu_t + \mu_c + \lambda_t \times \delta_j^h + \varepsilon_j, \quad (3)$$

where y_j is the inverse hyperbolic sine of one of the usage measures, μ_c and μ_t are category and publication date fixed effects, and δ_j^h indicates whether book j is in the highly exposed group. As above, we weight for frequencies to reflect total releases; and we cluster standard errors by category \times release month. The coefficients λ_t show how the outcome measures for the highly exposed group evolve over time, relative to the time pattern for the less exposed groups. We define $\delta_j^h = 1$ if the category’s growth in entry is in the top quartile. We term this the “high-low” approach.

The upper left panel of Figure 6 shows that average quality falls more sharply for the highly treated categories than the others, indicating that a stronger LLM treatment reduces quality more. Panel (b) compares estimates of the λ_t coefficients using both the adjusted

and raw usage measures. Both measures decline following 2022; and the results are very similar. This provides some assurance that our results are not driven by the adjustment procedure for \tilde{r}_j .

Panel (c) reports the time effects from a version of Equation (3) augmented to include rank percentile fixed effects. Conditional on percentiles, quality falls. Panel (d) shows the month effects for the highly treated categories, conditional on rank position fixed effects. As with the temporal approach, quality rises conditional on rank position.¹⁴

5.3 Quality and the volume of category entry

We also make a cross-category temporal comparison with the number of books appearing in a category per month (n_{ct}) as a continuous measure of the extent of the category’s LLM treatment. We first explore the impact of additional entry on average quality using the regression

$$y_j = \mu_c + \mu_t + \phi \ln(n_{ct}) + \varepsilon_j, \quad (4)$$

where y_j is an outcome measure for book j , 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.

Table 3 shows the usage results. The first two columns show the results for average quality. Using both adjusted and unadjusted usage measures, categories with faster growth in entry experience larger decreases in average quality; and the coefficients are similar. Columns (3) and (4) explore impacts within percentile ranges. The regressions are based on Equation (4) augmented with indicators for within-category percentile ranges 0-5, 6-10, and so on. The negative coefficients for both usage measures indicate that categories with more growth

¹⁴We obtain very similar results using the rank position sample. See Panel (b) of Figure A.3.

in releases experience larger reductions in quality, conditional on percentile. Columns (5) and (6) explore impacts within rank position ranges. Here, the regressions from Equation (4) are augmented with fixed effects for rank ranges 1-100, 101-200, and so on. Using both usage measures, quality rises more – within rank ranges – in categories with faster growth in entry.

While constant in Table 3, the coefficient on $\ln(n_{ct})$ could vary across percentiles and rank positions. The left panel of Figure 7 plots rank percentile-specific coefficients from regressions that include month, category, and rank percentile fixed effects. The coefficient on $\ln(n_{ct})$ is smaller, albeit negative, for the top percentiles and grows more negative for worse ranking percentiles. The right panel of Figure 7 reports the rank-position-specific coefficients from regressions with month, category, and rank-position fixed effects. The coefficient on $\ln(n_{ct})$ is insignificantly different from zero for books ranked 1-100 among the titles released in the category and month. The coefficient then grows as ranks worsen, reaching roughly 1 for ranks beyond 800. Thus, the large, LLM-aided growth in books does not raise the quality of the top 100 books released per category-month; but it does raise the quality of books outside of the top 100.

The results above paint a very consistent picture of the relationship between LLM-induced entry and the distribution of quality, and the framework in Section 2 is useful for guiding interpretation of the results. There is an enormous increase in the number of books released. Despite a decline in average quality – which also produces reductions in quality within percentile ranges – there is an increase in the quality of books at particular rank positions near the right tail of the distribution, for all but the top 100 titles released per category and month.

5.4 Incumbent vs new-author contributions

Figure 8 provides a visual summary of the paper’s main results, as well as a preview of their welfare effects. The figure compares the numbers – and adjusted ratings distributions – of

books from before and after the AI influx, among the books with an adjusted rating above a small threshold that excludes the books with zero raw ratings. The light-shaded region, for books released in the LLM era (2023-2025), is far larger than the darker region, for the pre-LLM era (2020-2022), reflecting the large increase in the number of books released. The AI influx has brought a moderate increase in the absolute number of moderately-valuable books – near the right tail of the distribution – and an enormous increase in the number of books in the left tail of the distribution.

What sorts of books and authors are delivering the changing choice set in the LLM era? For example, how well do the books in the 100-1000 rank range sell? The relationship between ratings and sales estimated above in Section 3.3 allows us to translate numbers of ratings into sales estimates. Averaging across categories and months, books at rank position 100 (within category \times month) sell an average (median) of 2,291 (686) copies by 60 months; and books at rank position 1,000 sell an average (median) of 397 (57).¹⁵ The LLM-affected choice set includes many additional books that are used by many people.

The increase in book releases reflects both increased new author entry and repeat entry by incumbent authors. Panel (a) of Figure 9 shows the numbers of debut and repeat-author releases in the census sample by month. Both rise considerably in the LLM era, but the repeat authors of Panel (a) include additional releases by authors who first entered in the LLM era. We can separately document the productivity of pre-LLM authors with a simple cohort approach that accounts 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 v . Productivity of a cohort evolves over time as it ages, where “age” $a = t - v$. To identify the productivity of pre-LLM authors over time from these age effects, we estimate the following

¹⁵In reality, some of these “sales” are instances of borrowing through Kindle Unlimited. Still, the additional books of the LLM eras are consequential books in the sense that they each reach thousands of consumers.

model on author cohorts entering before 2023:

$$\ln(n_{tv}) = \omega_t + \mu_a + \varepsilon_{tv}, \tag{5}$$

where ω_t and μ_a are time and age fixed effects, and ε_{tv} is an idiosyncratic error.

Panel (b) of Figure 9 shows the time pattern of the ω_t coefficients. They are stable between 2015 and 2021 and then dip in 2021, prior to the widespread diffusion of LLMs; and they remain below their earlier level during 2023. The coefficients rise to their earlier levels in 2024, then rise beyond pre-LLM levels in 2025. This shows that LLM availability on balance raises the output of incumbent authors.¹⁶

To what extent do LLM-era debut authors, compared with those entering earlier, create the successful books among the large LLM book influx? We explore this with an alternative cohort approach. We aggregate the data to the number of entering books by author entry vintage v and author age a (months between release and debut months). In particular, we aggregate to the number of books among the top 500 most-rated books per release month in the census sample and the remainder. That is, we run regressions of the form

$$\ln(n_{va}^k) = \mu_a + \gamma_v + \varepsilon_{va}$$

where k is either top 500 or remainder. Panel (c) of Figure 9 shows author debut vintage γ_v effects for the more successful books. The resulting γ_v coefficients are lower for the LLM-era author entry vintages, indicating that recent entrants are less likely to produce books in the top 500. Panel (d) of Figure 9 reports results for the less successful books (outside the top 500), and LLM-era author vintages have a sharply elevated tendency to produce these

¹⁶These results contrast with [Goldberg and Lam \(2025\)](#), who find that the entry of AI-assisted images causes “extensive” crowd out of traditional production. Their analysis is at the product, rather than the artist level, so their finding might reflect incumbent artists adopting AI tools. The survey evidence cited above indicates that roughly half of incumbent authors use LLMs. The implied effect on adopters is thus twice what’s shown in Figure 9.

less-successful books. Hence, much of benefit from the LLM-era book influx comes from pre-LLM-era authors.

6 Welfare effects

In addition to welfare effects operating through additional products, changes in pricing and user satisfaction (as reflected in star ratings) could also affect welfare. Section 6.1 discusses effects of entry on prices and star ratings, and we measure the change in CS resulting from the changes in the number, and quality, of products in Section 6.2.

6.1 Prices and user satisfaction

We explore the relationships between entry and both log prices and average star ratings using regressions analogous to those in Table 3. None of the six resulting coefficients (reported in Table A.1) are statistically significant. Additional entry does not appear to change these channels of influence on consumer well-being.

While the additional entry does not change prices, it is important to note that many of books are available not only for an a la carte price but also through Amazon’s Kindle Unlimited subscription service. For \$12 per month, users have access to a large number of books at Amazon at no additional cost.¹⁷ In 2025 the Kindle Unlimited catalog reportedly included “more than 5 million digital books, audiobooks, comics, manga, and magazines.”¹⁸ This is consistent with our data: Roughly half of the books in the category sample representing ten million underlying books were included in Kindle Unlimited as of the start of 2026.

¹⁷At launch in 2014, the service included 600,000 titles. See <https://time.com/3004637/amazon-kindle-unlimited/>. By 2018 the catalog had grown to 1 million titles. See <https://www.theatlantic.com/technology/archive/2018/07/amazon-kindle-unlimited-self-publishing/565664/>.

¹⁸See <https://www.aboutamazon.com/news/devices/what-is-kindle-unlimited>.

6.2 Consumer surplus calibration

LLMs have delivered large increases in the numbers of books created, along with increases in the quality of products at various rank positions near the right tail of the distribution. This by itself indicates a welfare benefit to consumers. This section presents a calibration of a simple structural model to quantify the welfare benefit to consumers in the book market.

One could measure the welfare effect in various ways. One approach would be to ask how the recent, unusually large crop of books delivered by AI has added to the CS associated with a choice set that consists of every book ever written. Our interest, however, is in comparing the benefit of LLM-accelerated creation vs pre-LLM creation. Hence, we seek to compare environments with and without LLMs. Accordingly, we compare hypothetical environments in which all extant books were created by the processes delivering the 2020-2022 vs the 2023-2025 books.

To this end, we develop a nested logit calibration in the spirit of [Berry \(1994\)](#). Consumer i chooses between books j and the outside good according to

$$u_{ij} = x_{ij}\beta - \alpha p_j + \zeta_{ig} + (1 - \sigma)\epsilon_{ij},$$

where ϵ_{ij} is extreme value, α is the marginal utility of a dollar, 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 of product j is given by $\delta_j = \ln(s_j) - \ln(s_0) - \sigma \ln(s_j/(1 - s_0))$. Then the consumer surplus associated with a choice set is

$$CS = \frac{M}{\alpha} \ln \left(1 + \left(\sum e^{\delta/(1-\sigma)} \right)^{1-\sigma} \right).$$

We treat our adjusted usage measures \tilde{r}_j as quantities, and we estimate M for each year as

the US population times ten. We use an estimate of σ from [Reimers and Waldfogel \(2021\)](#) of 0.443 for the nested logit implementation. We lack an estimate of α and so will focus on the percent change in CS across environments, which does not depend on α .

Estimating the welfare effects of LLMs on the book market requires two elements from each of the pre-LLM and LLM eras. First, we need the numbers of books that would be produced per period. Second, we need the quality distributions from which new books are drawn in each period.

As Panel (a) of [Figure 3](#) shows, book releases were steady during 2020-2022 at roughly 10,000 per month. Releases rose to 15,000 per month during 2023, to 20,000 by the end of 2024, and reached a peak of over 30,000 per month in late 2025. We treat 10,000 per month as the pre-LLM output rate. The level of post-2022 output that reflects the full availability of LLMs is less clear, as it is unclear which – if any – of the output levels since 2023 will be sustained. Accordingly, we consider various possibilities, including the actual number produced 2023-2025, as well as a replication, a doubling and a tripling of the pre-LLM quantity.

We take as given the distribution of product qualities for the pre-LLM era. We create LLM-era choices in two ways. We use the actual entering products for the “observed” comparison. For the simulations involving multiples of the pre-LLM number of products, we draw randomly – with replacement – from the observed LLM-era distribution. This gives LLM-era counterfactual choice sets that are one, two, and three times the size of the pre-LLM choice set.

[Table 4](#) shows the simulation results. If the period 2023-2025 had the same number of releases as the 2020-2022 period, then – given the lower average quality of the LLM era books – CS would have fallen by 13 percent. The actual increase in books, from 3.6 million released 2020-2022 to 6.7 million released 2023-2025, delivers 21 percent more CS in the second period. A counterfactual doubling in the number of books, with the new books

drawn from the LLM era distribution, would raise CS by 24 percent. Finally, a tripling – reflecting the maximum LLM-era release rate – would raise CS by 53 percent. While it is hard to know the appropriate counterfactual entry expansion to entertain, the effect of LLMs on welfare appears to be substantial.

7 Conclusion

LLMs, which have diffused rapidly between 2022 and 2025, have become increasingly capable of writing long passages of text. Accordingly, they have seen rapid adoption by existing, and especially new, book authors. We document a tripling in the number of new books coming to market between the start of late 2022 and late 2025 that coincides with growing use of AI tools. The effects of this influx on welfare depend on the quality of the additional books. While the average quality of new books has fallen with the LLM-induced influx – and while much of the new work appears to be of little value to consumers – the LLM influx has delivered a large number of works near the top of the usage/quality distribution. The LLM-era entry process delivers a quarter to a half more consumer surplus from books than the pre-LLM process. Moreover, the arrival of LLMs does not appear to have displaced activity by incumbent authors. Despite the controversy surrounding LLMs, their effect on book consumers – like other cost-reducing technological changes in the cultural industries – appears to be positive and large.

While AI has improved the welfare of consumers, participants in the book publishing industry have regarded AI with concern and suspicion. The Authors Guild has argued that publishers and sellers should be compelled to disclose the use of AI in book creation.¹⁹ An open letter from authors, including Colleen Hoover and Dennis Lehane, exhorted publishers

¹⁹The 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/>.

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.”²⁰

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.”²¹ Even Amazon appears to be cautious in its embrace of AI. Amazon requires authors to disclose AI-generated, but not AI-assisted, content to Amazon; but Amazon does not disclose the information to consumers.²² Moreover, Amazon limits the number of books that authors may upload per day.²³

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

²⁰On January 25, 2026, the list included 70 authors in total. See <https://lithub.com/against-ai-a-n-open-letter-from-writers-to-publishers/>.

²¹The 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. 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 <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf>. See also Lutes et al. (2025).

²²“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.

²³See <https://www.theguardian.com/books/2023/sep/20/amazon-restricts-authors-from-self-publishing-more-than-three-books-a-day-after-ai-concerns>.

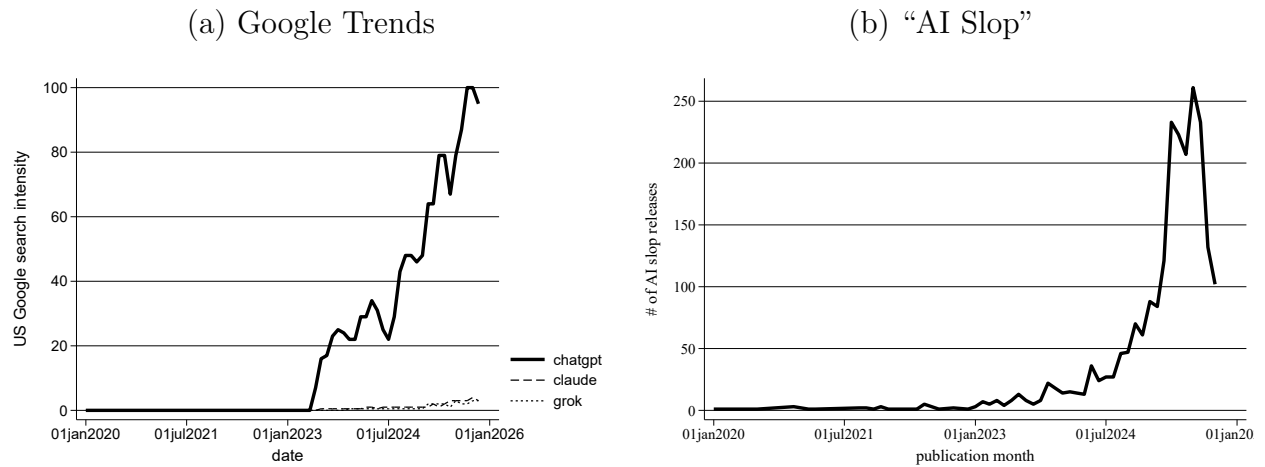
References

- 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.
- BERRY, S. T. (1994): “Estimating discrete-choice models of product differentiation,” *The RAND Journal of Economics*, 242–262.
- BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2023): “Generative AI at Work,” *Quarterly Journal of Economics*, forthcoming.
- CAVES, R. E. (2000): *Creative industries: Contracts between art and commerce*, 20, Harvard university press.
- DELL’ACQUA, F., S. RAJENDRAN, E. I. MCFOWLAND, E. MOLLIK, 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.
- 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)*.

- 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.

8 Figures and tables


Figure 1: The arrival of AI in books



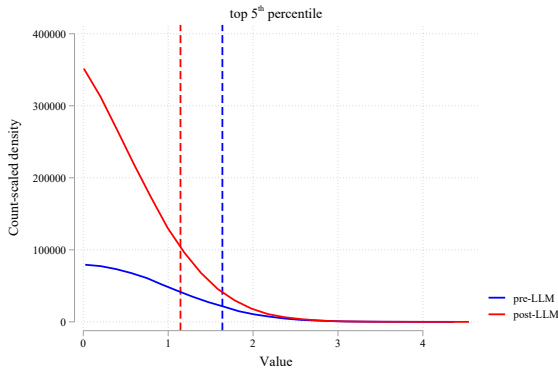
Notes: The left panel shows Google Trends data on search intensity for ChatGPT, Claude, and Grok. The right panel shows "AI slop" identified by Goodreads users.

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

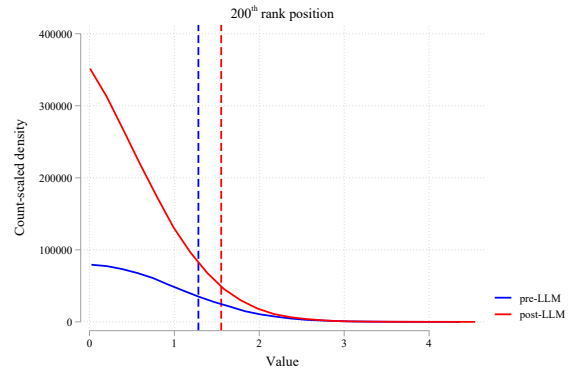
(a) Full distributions



(b) Value of top 5th percentile

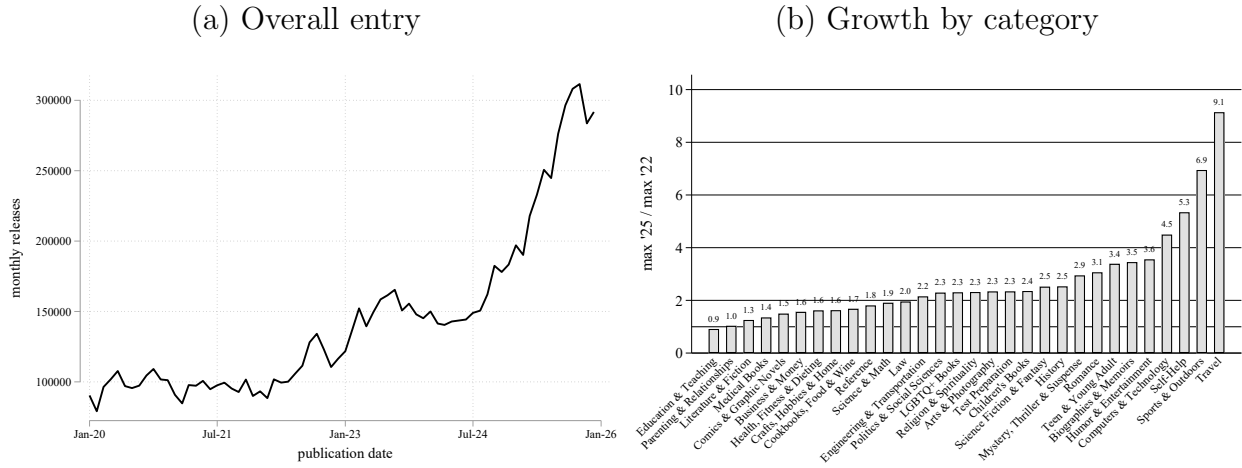


(c) Value at 200th rank position



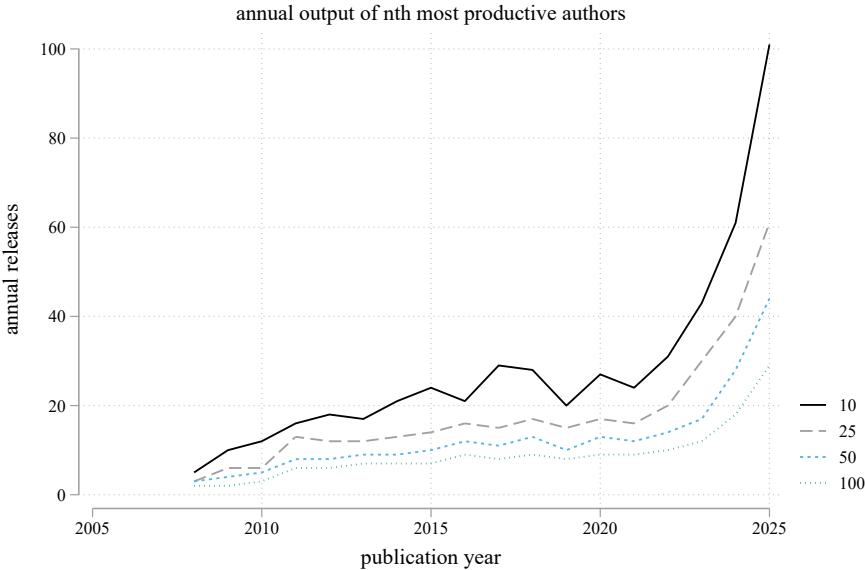
Notes: Panel (a) compares hypothetical pre-LLM (blue) and LLM-era (red) quality distributions, where the LLM era distribution has a lower mean and more products. The Panels show how an increase in the number of products, with a lower average quality, can affect the qualities at particular percentiles and rank positions. Using these distributions, Panel (b) shows that the 5th percentile product has lower quality in the LLM era, while Panel (c) shows that a product at a particular rank position has higher quality in the LLM era.

Figure 3: Overall entry effects



Notes: The left panel shows aggregate entry across all categories. May 2024 is excluded because queries deliver oddly large results in many categories for that month. The right panel shows the ratios to the maximum monthly new publications if 2025 to the average monthly publications during 2020-2022, for each category.

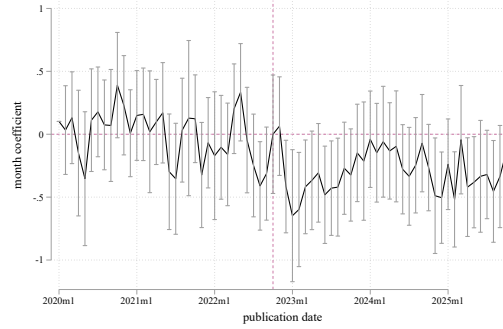
Figure 4: Annual releases of most prolific authors




Notes: Annual number of books published by most prolific authors in our census sample. For example, the solid black line shows the number of books produced by the 10th most productive author in each year.

Figure 5: Number of eventual book ratings across publication dates


(a) Overall average



(b) Conditional on rank percentiles




(c) Conditional on rank positions



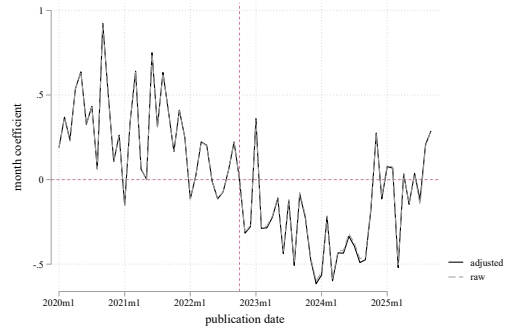
Notes: Month effects from regressions of \tilde{r}_j on month effects and category dummies, weighted to reflect total releases in the categories. Panel (a) is the overall average. Panel (b) includes fixed effects for rank percentile deciles and includes the top 50 deciles. Panel (c) includes fixed effects for rank position ranges within categories by 100, up to 2000.

Figure 6: Quality effects: “high-low” approach


(a) Overall average




(b) Overall average, adj. vs. raw



(c) Conditional on percentiles



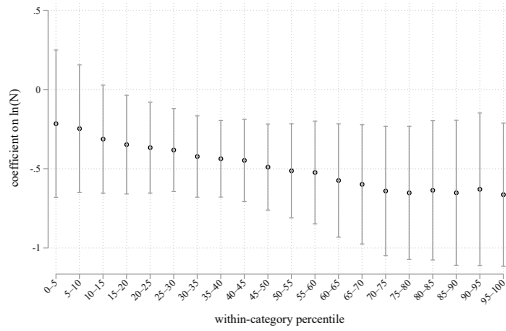
(d) Conditional on ranks



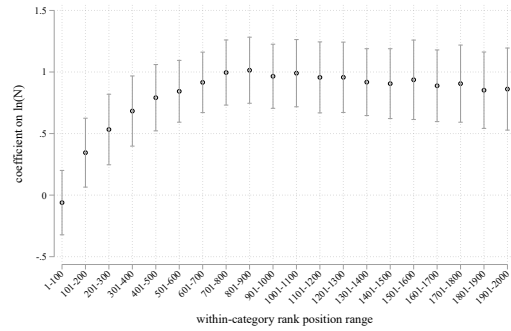
Notes: The figures report the coefficients on monthly adjusted usage effects for high-entry-growth categories relative to the others. Panel (a) reports the average effect. Panel (b) compares effects using the adjusted vs unadjusted usage measure. Panel (c) reports effects conditional on rank percentiles, and Panel (d) reports effects conditional on rank positions (1-100, etc). All regressions include category fixed effects and month effects, and standard errors are clustered on category \times release month.

Figure 7: Coefficients on $\ln(N)$

(a) By rank percentile




(b) By rank position



Notes: Coefficients from regressions of our adjusted usage measure on the log of monthly releases per category, with month and category fixed effects and standard errors clustered on category \times release month. In Panel (a), the coefficients are percentile-specific. The regression includes month and category fixed effects, as well as fixed effects for each 5-percentile group. In Panel (b), the coefficients are rank-position specific, and the regression includes month, category, and rank group fixed effects.

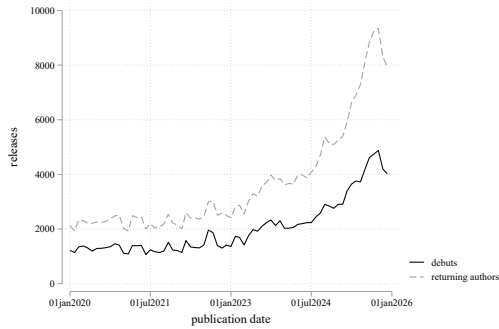
Figure 8: Log ratings distributions for 2020-22 vs 2023-25




Notes: Distributions of log adjusted ratings for 2020-22 vs 2023-25. Distributions are based on weighted data, and they are not normalized to one.

Figure 9: Who is producing the new books?


(a) Debut vs returning authors



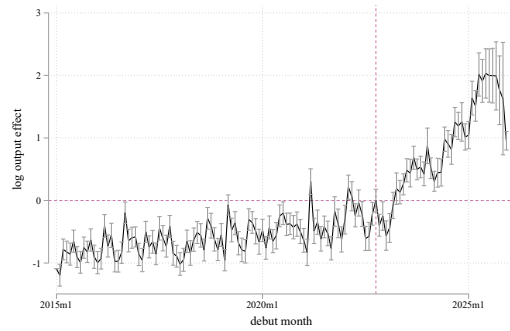
(b) Productivity of incumbent authors



(c) Vintage effects – top 500



(d) Vintage effects – outside top 500



Notes: Panel (a) shows debuts and releases by returning authors. Panel (b) shows calendar time fixed effects from regressions of log releases, by time and author age (months since first release) on time and age fixed effects. Panels (c) and (d) show author debut month fixed effects from regressions of log releases by debut month and author monthly age since debut on age and debut month fixed effects. Panel (c) includes only the top 500 books, by ratings received, from each release month. Panel (d) includes the remainder (books outside the top 500). All figures are based on the census sample.

Table 1: Summary statistics – categories random sample

	# books	# ratings	avg rating	price	Kindle Unl.
average	10,306,237	112.95	4.44	9.7	0.51
rank position 1	216,338	2,320.91	4.49	9.52	0.41
rank position 2	213,945	830.79	4.47	10.71	0.43
rank position 3	218,732	407.73	4.44	11.25	0.44
rank position 4	213,812	217.01	4.41	12.36	0.46
rank position 5	214,750	208.72	4.41	13.18	0.46
rank position 6	210,963	193.38	4.39	13.88	0.45
rank position 7	201,795	131.31	4.39	14.54	0.46
rank position 8	201,930	126.62	4.39	14.87	0.47
rank position 9	193,173	91.98	4.41	14.84	0.46
rank position 10	190,420	75.99	4.40	14.96	0.46
rank position 11	180,858	70.55	4.41	15.36	0.47
rank position 12	176,832	87.18	4.46	15.25	0.48
percentile decile 1	1,031,474	1,012.68	4.51	7.56	0.52
percentile decile 2	1,026,853	75.3	4.45	7.84	0.53
percentile decile 3	1,033,809	24.09	4.42	8.46	0.53
percentile decile 4	1,036,726	9.64	4.42	9.23	0.53
percentile decile 5	1,032,664	4.16	4.42	9.88	0.52

Notes: Descriptions of the random categories sample, which contains 333,224 observations. The means in the table we weighted to deliver the 10.3 million books released 2020-2025. The average describes the entire sample. The rank position entries refer to groups 100 within-category-and-month rank positions. For example, “rank position 1” refers to books ranked 1-100 (by number of eventual ratings received) among those in the category and release month. The percentile decile entries refer within-category-and-month percentiles. “Percentile decile 1” refers to books in the top (91-100) decile, 2 is 81-90, etc.

Table 2: Summary statistics – census of eight subcategories

	# books	ratings/book	avg rating	price	Kindle Unl.
Average	479,648	166.6	4.33	14.62	0.40
Business & Money	92,046	37.15	4.32	38.42	0.23
History	90,389	70.33	4.36	21.00	0.22
Literature & Fiction	18,731	82.96	4.28	5.18	0.56
Mystery, Thriller & Suspense	56,636	530.12	4.36	5.07	0.50
Romance	72,742	302.01	4.36	4.41	0.61
Science Fiction & Fantasy	91,692	152.92	4.36	4.50	0.52
Sports & Outdoors	23,744	66.72	4.30	8.02	0.25
Travel	33,668	29.11	4.07	8.01	0.42

Notes: The table describes the “census” data, the complete list of books published in selected subcategories of the listed categories, January 2008 - December 2025. The particular subcategories are economics, history (world), alternative history, mystery (women sleuths), romantasy, science fiction adventure, sports & outdoors (sports biographies), and travel (US). Subcategories are chosen to be small enough to make data collection feasible. See the text.

Table 3: Entry and adjusted and raw usage measures

	average		percentile		rank position	
	(1) adjusted	(2) raw	(3) adjusted	(4) raw	(5) adjusted	(6) raw
log monthly pubs	-0.482*** (0.133)	-0.479*** (0.130)	-0.488*** (0.133)	-0.485*** (0.130)	1.345*** (0.177)	1.335*** (0.174)
Observations	9735079	9735079	9735079	9735079	9735079	9735079

Notes: Regressions of adjusted and raw usage measures – which are the $\text{asinh}(\text{ratings})$ or $\text{asinh}(\text{adjusted ratings})$ on the log of the number of monthly releases in the category. All regressions include category and month fixed effects; and standard errors are clustered at the category level. The rank position regressions also include fixed effects for adjusted usage rank positions 1-100, 101-200, etc. within category-month. The percentile regressions also include fixed effects for percentiles of adjusted usage within category-month (0-5,6-10, etc).

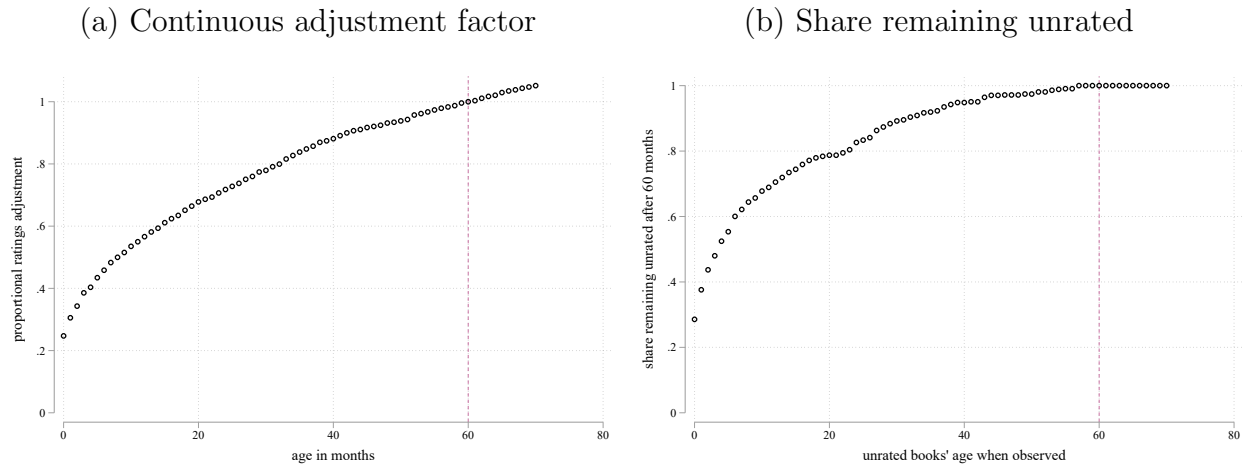
Table 4: Consumer surplus effects of LLMs

	# books	% Δ CS
Quality effect	3,558,869	0.87
<u>Quantity + quality effect:</u>		
Observed books	6,747,368	1.21
Double # of pre-LLM books	7,117,738	1.24
Triple # of pre-LLM books	10,676,610	1.53

Notes: The table shows the proportionate change in CS for the LLM era (2023-2025) relative to the pre-LLM era (2020-2023) under different assumptions about the rate of creation in the LLM era. Pre-LLM CS is based on actual books released 2020-2022, and the “observed books” entry reflects actual books released 2023-2025. The other rows reflect random draws of releases from the LLM-era quality distribution.

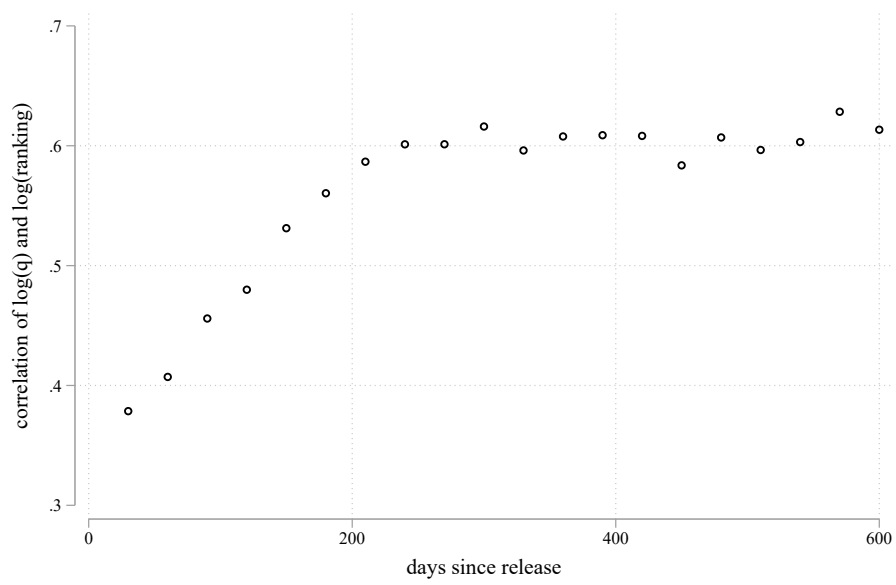
A Appendix figures and tables

Figure A.1: Ratings adjustments by book age at collection



Notes: The left panel shows the proportional continuous adjustments factors for ratings. We divide raw ratings by the adjustment to create our 60-months-old “adjusted rating” (\tilde{r}_j). See text. The right panel shows the adjustments for books with zero ratings as of age when observed. Relatively few books observed with zero ratings at month zero – about a quarter – would still have zero ratings at 60 months, while 80 percent of books observed with zero ratings 20 months after publication would still have zero ratings at month 60.


Figure A.2: Quantities and the number of ratings



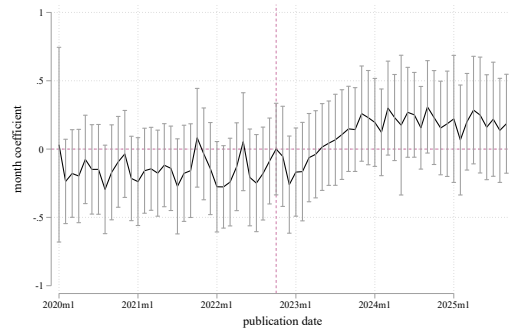
Notes: The Figure shows the correlation between log cumulate sales and the log of rankings as of various number of days after publication. Included ebooks were published in April 2017 and Januaries of 2018-2020.

Figure A.3: Quality effects conditional on rank position, rank position sample

(a) temporal approach



(b) High-low approach



Notes: Panel A shows month effects from regressions of inverse hyperbolic sine of adjusted ratings on category, rank position, and month fixed effects using the rank position sample. Panel B shows coefficients on interactions of month effects with an indicator for the top quartile of categories, by growth in entry. Other variables include month, category, and rank position effects.

Table A.1: Entry, star ratings, and prices

	average		percentile		rank position	
	(1)	(2)	(3)	(4)	(5)	(6)
	stars	ln(price)	stars	ln(price)	stars	ln(price)
log monthly pubs	-0.0692 (0.0467)	-0.0893 (0.0589)	-0.0792 (0.0499)	-0.0892 (0.0589)	-0.0474 (0.0418)	-0.0640 (0.0567)
Observations	4843538	9672374	4843538	9672374	4843538	9672374

Notes: Regressions of average star ratings and log-prices on the log of the number of monthly releases in the category. All regressions include category and month fixed effects; and standard errors are clustered at the category level. The percentile regressions in columns (3) and (4) also include fixed effects for percentiles of adjusted usage within category-month (0-5,6-10, etc). The rank position regressions, in columns (5) and (6) also include fixed effects for adjusted usage rank positions 1-100, 101-200, etc. within category-month.