Fewer Buyers, Higher Pay? Monopsony and Product Market Power in the U.S. Publishing Industry

Maximiliano Machado*

November 3, 2026

The most recent version is available here.

Abstract

In 2022, a U.S. federal court blocked the merger of Penguin Random House and Simon & Schuster, the first U.S. merger halted solely on the basis of harm to an input market. Standard monopsony analyses hold firms' input valuations fixed, yet a merger that raises downstream pricing power raises the profits an input generates, and with it every firm's willingness to pay. This paper quantifies both channels in a unified framework and applies it to the blocked merger. Using a new dataset linking retail book sales to the corresponding manuscript acquisitions, I first document that, after the most recent merger between large publishers, advances declined for authors with a bestseller but increased for debut and mid-list authors. Reduced competition for manuscripts alone cannot explain these changes. I then develop and estimate a structural model that links ascending-bid auctions for manuscripts to differentiated-products demand and multiproduct price competition in the retail book market. Simulating the blocked merger, consumer surplus falls by \$19.9 million (-2.5%) while aggregate advances rise by \$20.3 million (10.2%): advances fall for the manuscripts the merging publishers jointly contested but rise across the much larger remainder of the market, where heightened downstream profitability lifts every publisher's bids. An upstream-only analysis of the same estimated model predicts a decline in author compensation, mis-signing the effect of a merger on authors.

Keywords: Merger Analysis, Competition Policy, Monopsony, Auctions, Publishing Industry
JEL Codes: L13, L41, J23, D44

*University of Toronto, Rotman School of Management. Email: maxi.machado@rotman.utoronto.ca
I thank my supervisors Victor Aguirregabiria, El Hadi Caoui, Heski Bar-Isaac and Yao Luo for the guidance and patience in the process of writing this article, as well as Samsun Knight and Yakov Bart for the valuable feedback. I also benefited from the comments by Rob Clark, Mitsuru Igami, Ig Horstmann, Matt Mitchell, Zi Yang Kang, Guembi Park, Elena Prager, Jeff Qiu, Shishir Shakya, Eduardo Souza-Rodrigues as well as all the participants at the IO Lunch Seminar at UofT, the Southern Economic Association Meeting 2024, the 1st PhD Conference at Boston College, and the 20th Economics Graduate Student Conference at Washington University in St. Louis. This research was funded by the TD MDAL at the Rotman School of Management. All errors remain my own.

1 Introduction

In October 2022, a U.S. federal court blocked the proposed merger between Penguin Random House (PRH) and Simon & Schuster (SS), the first merger in U.S. history blocked solely on the basis of harm to an input market: the market for anticipated best-seller manuscripts.[1] The ruling is consistent with Hemphill and Rose (2018), who argue that harm to sellers is sufficient to support antitrust liability against a merger of competing buyers. Together with the 2023 Merger Guidelines' explicit instruction to evaluate competitive effects in input markets (DOJ and FTC, 2023), it reflects a broader reorientation of U.S. antitrust policy toward labor and supplier monopsony, evidenced by a string of recent enforcement actions (U.S. Department of Justice, 2016, 2018, 2021b,a, 2022b, 2023).

Yet the frameworks used to evaluate such cases, including the one relied upon by the Department of Justice in the PRH–SS case, typically treat the input market in isolation and hold fixed what inputs are worth to buyers. If a merger raises the combined entity's pricing power in the product market, it raises the profit each input generates, and with it the prices firms are willing to pay for inputs. This feedback operates within the input market itself: it moves the advances paid to the same authors whose outside options the merger weakens. Monopsony-only analyses miss it and may therefore misstate the harm to upstream suppliers.[2] They are also silent on the consumer side, which is the traditional object of merger review (Setzler, 2025).

This paper asks how accounting for downstream profitability changes the estimated effects of a merger on suppliers, relative to an analysis that holds input valuations fixed. By linking product-market competition to publishers' bids, I quantify how the PRH–SS merger would have changed aggregate advances and their distribution across authors. Because the same model describes the retail market, it also measures the merger's effect on consumers, which an input-market analysis cannot (Setzler, 2025). The exercise does not presume that gains outside the relevant input market can offset competitive harm within it.[3]

The U.S. trade publishing industry is an ideal market for this purpose, where the same firm that acquires the rights to a manuscript also sets the retail price of the resulting book. How much a firm is willing to pay for a manuscript reflects the downstream profit it expects the title to generate. The industry is also highly concentrated on both sides of this structure, with five publishers—known as the Big Five—accounting for the majority of U.S. trade sales and competing for the most commercially promising manuscripts.[4] The object of upstream competition is the advance, i.e., a lump-sum payment made at contract signing and unconditional on subsequent sales, that constitutes the primary form of author compensation (U.S. Department of Justice, 2022a). When several publishers pursue the same manuscript, competition among

[1] The court defined the relevant market as manuscripts commanding advances above $250,000, the segment it labeled "anticipated best-sellers," in which the merging parties frequently competed head-to-head.

[2] Throughout, "monopsony" is used in the sense of the literature: buyer power arising from concentration among purchasers, not literally a single buyer. The blocked merger would have reduced the number of major bidders for manuscripts from five to four. As Section 2.1 discusses, the market for a given manuscript can nonetheless approach the single-buyer case when the merging parties are the only interested bidders.

[3] See DOJ and FTC (2023), Guideline 10 and Section 3.3. Guideline 10 states that harm to competition in an upstream market is not offset by benefits in a separate downstream market.

[4] The Big Five are Penguin Random House, HarperCollins, Hachette, Macmillan, and Simon & Schuster. Sections 2 and 3 detail their structure and market shares.

them determines how much the author is paid.

To study how merger effects propagate across the two sides of this market, I assemble a novel dataset that matches downstream retail sales for thousands of titles to the corresponding upstream rights transactions, including the advances paid to authors. With these data I build and estimate a structural model that nests an ascending-bid auction for manuscripts inside a differentiated-products demand model for books, and use it to simulate the blocked merger.

I begin with evidence from the 2013 merger of Penguin and Random House, the most recent consolidation among the largest U.S. publishers. After the merger, advances fell for authors with a previous best-seller, the segment in which the merging firms competed most, but rose for debut and mid-list authors, with both changes concentrated at the merged firm. Reduced competition for manuscripts can explain the first pattern but not the second, which points to higher downstream profits raising publishers' bids. The evidence cannot separate the two forces or measure their size, which motivates the structural model.

In the downstream model, consumers choose among differentiated titles and publishers are multi-product firms competing in retail prices. Upstream, the same publishers bid for manuscripts in ascending auctions, valuing each title at the expected profit from publishing it rather than leaving it to a rival, adjusted for risk. A merger therefore removes a bidder from some auctions and, by changing retail prices, changes every publisher's valuation in all of them. Demand is identified from cost shifters, namely paper and ink prices and publishing-sector wages interacted with page counts, and from variation in the set of competing titles. Given the expected profits implied by demand, publishers' valuations are identified from which publisher wins each manuscript and the advance it pays.

In the simulated PRH–SS merger, aggregate advances rise by 10.2%, whereas an analysis that holds publishers' valuations fixed predicts a decline of 4.3%: ignoring the downstream channel reverses the sign of the merger's effect on authors as a group. The aggregate conceals sharply uneven incidence. Authors whose manuscripts both merging publishers competed for, largely established best-sellers, receive lower advances, as the DOJ anticipated, while all other authors receive higher advances, because higher retail prices raise every publisher's willingness to pay. Meanwhile, consumer surplus falls by 2.5% and total welfare by 0.3%. An assessment confined to the input market would therefore correctly identify the authors harmed, but miss both the gains to other authors and the loss to consumers.

To validate the auction model, I compare its upstream-only predictions with those of the DOJ's expert witnesses, who used a different methodology and confidential bid logs (U.S. Department of Justice, 2022a). The model reproduces the direction of their estimates and the larger harm to Simon & Schuster's authors than to Penguin Random House's, which supports using it for the full counterfactual.

More broadly, the paper speaks to how firms share their rents with the people whose work they commercialize. When each piece of work is tied to a specific product and its price is set by buyers competing for it, compensation reflects the profit that product is expected to generate, so changes in product-market power reach suppliers directly. This describes much creative and innovative work, from screenplays acquired by film studios to drug candidates licensed by pharmaceutical firms, and in such markets a merger's effect on suppliers cannot be assessed

apart from its effect on downstream profits.

Related literature. This paper relates to three strands of literature. First, it contributes to work on the economics of book publishing, which has focused on the downstream side. Other work has used retail data to estimate structural demand models for books, in order to investigate the effects of reviews (Reimers and Waldfogel, 2021), the role of copyright (Reimers, 2019) and fixed-price legislation (Kawaguchi et al., 2025; Genakos et al., 2025). The closest study of the industry is Shi (2024), who uses a matching model to show that the 2013 Penguin–Random House merger transferred non-monetary surplus from rival publishers and authors to the merged firm, with established authors most affected. I focus instead on the monetary terms of the relationship: advances are determined by publishers' competition for each manuscript and by the profit it is expected to generate downstream, so a merger can change what each author is paid. I model the upstream market for manuscript rights jointly with the downstream market for retail books, which to my knowledge has not previously been done in publishing or in any analogous creative industry.

Second, the paper contributes to the literature on mergers in auction markets, where consolidation is modeled as a reduction in the number of competing bidders or as a coalition that bids cooperatively (Loertscher and Marx, 2019; Froeb et al., 2017; Li and Zhang, 2015), and the standard prediction is that mergers reduce winning bids and harm sellers. The mechanism is closest to Gowrisankaran et al. (2015), who study how downstream hospital mergers affect upstream bargaining outcomes, though there the upstream mechanism is bilateral bargaining rather than an auction. Empirically, structural work on mergers in auction markets is scarce. This paper provides a merger-simulation framework in which an upstream auction is nested with a differentiated-products demand-and-pricing model downstream.

Third, the paper relates to the literature on monopsony, which has documented employer power in labor markets and its quantitative importance for wages and welfare (Azar et al., 2022; Yeh et al., 2022; Card, 2022), and which has studied the wage effects of horizontal mergers directly: Prager and Schmitt (2021) on hospital consolidation, Arnold (2025) on labor-market concentration, and Rubens (2023) on the Chinese tobacco industry. On the theoretical side, Bar-Isaac et al. (2024) show that acquisitions can themselves serve as instruments of monopsony, eliminating the closest competitor for specialized talent, in essence, the theory of harm the DOJ advanced in the PRH–SS case.

Closest in spirit to this paper are Berger et al. (2023), who develop merger guidelines that account for monopsony power and apply them to PRH–SS within a macroeconomic framework, and Setzler (2025), who recovers monopsony parameters from reduced-form merger evidence in the U.S. hospital industry. Existing merger-simulation frameworks (Nevo, 2000; Miller and Weinberg, 2017) treat the input side as fixed, while existing monopsony–merger analyses treat the output side as fixed. Only recently has work begun to relax both restrictions: Setzler (2025) in hospitals and Kroft et al. (2025) in construction. This paper does so in a setting where the upstream mechanism is an auction, providing the empirical counterpart to the 2023 Merger Guidelines' instruction to evaluate competitive effects on both sides of the value chain.

2 Mechanism and Industry Background

2.1 How a Merger Between Buyers Affects Input Prices

To fix ideas about how a horizontal merger affects the price paid to the seller of an input, consider N firms bidding in an ascending auction for a single input (e.g., a manuscript), which the winner will use to produce and sell a differentiated product. Each firm's valuation is the downstream profit it expects to earn from the input. Standard auction theory implies that the winner is the bidder with the highest valuation, and that the price paid equals the valuation of the second-highest bidder. A merger between two of the bidders changes this outcome through two channels.

The competition channel. Suppose four firms with valuations $v_{(1)} > v_{(2)} > v_{(3)} > v_{(4)}$ bid for the input, so the price the seller receives is $v_{(2)}$. If the merger combines the firms with valuations $v_{(1)}$ and $v_{(2)}$, and the merged firm preserves the higher of the two as its post-merger valuation, then the second-highest remaining bidder is the firm with valuation $v_{(3)}$, and the price falls from $v_{(2)}$ to $v_{(3)}$ (Panel (a) of Figure 1). This is the standard prediction: removing the runner-up reduces the price the seller receives. Critically, the prediction holds only when the merger involves both the winning bidder and the runner-up. If the merger involves bidders below the runner-up (e.g., those with valuations $v_{(3)}$ and $v_{(4)}$), the price is unchanged, because neither party determined it to begin with.

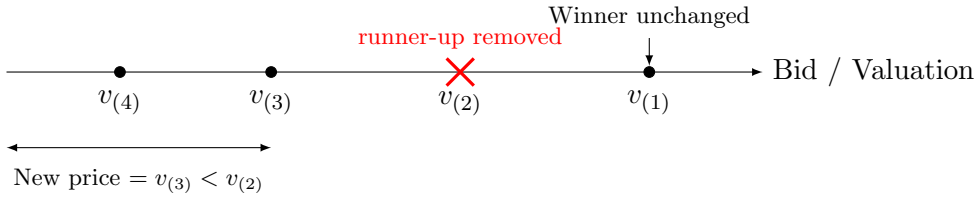
The downstream channel. The competition channel takes valuations as fixed. However, a merger between two firms changes the downstream environment as well. Consolidation can raise profits in the product market, through greater pricing power or through cost reductions. Either force raises the profit the merged firm expects from an input it acquires, and therefore its valuation in the upstream auction. In the four-bidder example, suppose the merger involves the firms with valuations $v_{(2)}$ and $v_{(3)}$ and raises the merged firm's valuation to $v'_{(2)}$, with $v_{(1)} > v'_{(2)} > v_{(2)}$. The price rises from $v_{(2)}$ to $v'_{(2)}$ (Panel (b)): the merger increases the price even though a bidder has been eliminated. The merged firm need not overtake the winner; a higher runner-up bid alone raises the price.

Moreover, the channel is not confined to the merging parties. Softer downstream price competition can increase the expected profit of every firm, so non-merging firms' valuations can shift upward as well. Prices would therefore rise even in auctions where neither merging party participates.

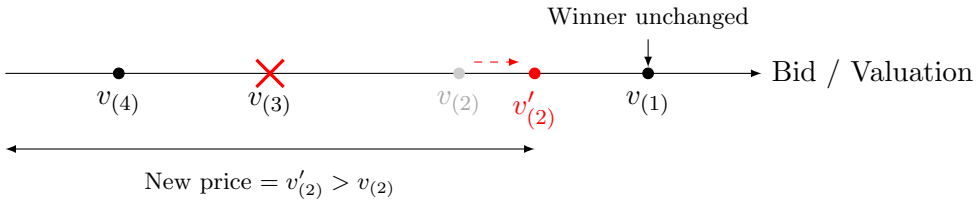
Implications. The two channels can operate in opposite directions, and which one dominates depends on (i) how often the merging firms are the top two bidders, and (ii) how much downstream consolidation raises profits across the industry. A purely upstream analysis measures only the first channel. It may therefore overstate the harm to sellers of inputs the merging firms jointly contest, and it misses the gains to sellers of inputs they do not. Quantifying the relative magnitude of the two channels is the central task of the structural model in Section 5.

The mechanism does not require inputs to be allocated through a literal auction. When an input is sold by direct bargaining, the runner-up's valuation still disciplines the outcome through

Figure 1: How a Horizontal Merger Affects Buyers: Two Channels



(a) **Competition channel**: the merger removes the runner-up.



(b) **Downstream pricing-power channel**: the merger does not remove the runner-up, but raises the merged firm's valuation. The advance rises from $v_{(2)}$ to $v'_{(2)}$.

Notes. Stylized illustration of the two channels through which a horizontal merger affects the price paid in an ascending auction with four bidders. In Panel (a), the merger combines the firms with valuations $v_{(1)}$ and $v_{(2)}$ and the merged firm preserves the higher valuation, with the auction price falling to $v_{(3)}$. In Panel (b), the merger combines the firms with valuations $v_{(2)}$ and $v_{(3)}$ and raises the merged firm's valuation through downstream pricing power, with the auction price rising to $v'_{(2)}$.

the seller's outside option: a seller negotiating with one firm, knowing that a rival would pay a high price if talks failed, bargains from a stronger position. The downstream channel therefore raises the seller's compensation under bargaining as well, so long as at least one firm outside the merging parties forms part of the seller's outside option. Appendix A formalizes this argument and shows that the sign and incidence of the merger's effects do not depend on the acquisition protocol, which governs magnitudes only.[5]

In the publishing industry, the input is a manuscript, the sellers are authors, the buyers are publishers, and the price is the advance: a lump-sum payment made when the contract is signed. The rest of this section describes the industry and the process through which manuscripts are acquired.

2.2 Industry Structure

In October 2022, a U.S. federal court blocked the proposed merger between Penguin Random House (PRH) and Simon & Schuster (SS), the first and only merger in U.S. history halted solely on the basis of harm to an input market, in this case, the market for anticipated best-seller manuscripts.[6] This is consistent with Hemphill and Rose (2018), who argue that harm to sellers is sufficient to support antitrust liability against a merger of competing buyers. The ruling, together with the 2023 Merger Guidelines' explicit instruction to evaluate competitive effects in input markets (DOJ and FTC, 2023), reflects a broader reorientation of U.S. antitrust policy

[5] Appendix A shows that, among the bargaining protocols considered, the auction yields the largest loss on jointly contested inputs, so the losses reported in Section 8 are upper bounds.

[6] The relevant market was defined as manuscripts commanding advances above $250,000 — the segment the court labeled "anticipated best-sellers" — in which the merging parties frequently competed head-to-head.

toward labor and supplier monopsony, evidenced by a string of enforcement actions in recent years (U.S. Department of Justice, 2016, 2018, 2021b,a, 2022b, 2023).

The U.S. trade publishing industry is highly concentrated. Five firms—Penguin Random House (PRH), HarperCollins, Hachette, Macmillan, and Simon & Schuster (SS), collectively known as the "Big Five"—account for more than 60% of trade book sales and roughly 50% of all manuscript acquisitions in a given year. Concentration is substantially higher at the top of the market: among deals classified by the DOJ as anticipated best-sellers, the Big Five account for around 75% of acquisitions. PRH is the largest of the five by both deal volume and total advance dollars paid, while SS oscillates between third and fifth place across years. Outside the Big Five, the market consists of a long tail of small and mid-size independent publishers, specialized presses, and self-published titles.

The Big Five are each organized hierarchically. At the top, the parent company sets strategy and provides shared resources. Below the parent are divisions, each with a distinct editorial focus. PRH, for example, includes the Penguin Publishing Group, the Random House Publishing Group, and Knopf Doubleday, among others. Divisions in turn contain multiple imprints, which function as semi-autonomous editorial brands. Acquisition decisions are made at the imprint level: editors at each imprint identify and bid for manuscripts that fit the imprint's editorial vision, and the imprint operates its own acquisition budget. Smaller publishers typically have a single editorial team and no divisional structure.

Despite this decentralized structure, a central institutional fact is that imprints within the same parent firm do not bid against one another for manuscripts: parent firms coordinate acquisition activity across their imprints to avoid internal competition for the same titles (U.S. Department of Justice, 2022a). This fact was central to the court's decision. The merging parties argued that their imprints would continue to compete with one another after the merger, while the evidence that parent firms suppress internal bidding led the court to reject that defense.

I therefore treat the parent firm as the relevant decision-making unit throughout the analysis: acquisition budgets and overall strategic direction are set at the parent level, where the parent internalizes the consequences of any one imprint's bid for another imprint's profits, consistent with the DOJ's analysis of the PRH–SS transaction.

2.3 The Acquisition Process

Authors who wish to publish with a traditional firm typically do so through a literary agent.[7] The agent solicits the manuscript, refines it editorially, and circulates it to editors at multiple publishing imprints, soliciting offers. When more than one editor expresses interest, the agent runs what is in effect an ascending auction: editors submit bids, the agent shares each bid (or its range) with the others, and the process continues until no remaining bidder is willing to

[7]The analysis below abstracts from agents as strategic actors. Agents are compensated with a share of the advance (roughly 15%), so their fee is a proportional commission that scales author take-home pay without altering the auction's comparative statics: the merger's effects on advances pass through one-for-one to authors, and the commission is a transfer that nets out of total welfare. Throughout, "author compensation" should be read as gross of the agent's commission. Agent heterogeneity in the quality of manuscripts represented is absorbed empirically by the agency controls in the reduced-form specifications and the imputation model (see Section 4 and Appendix D). What the model does not capture is any strategic response by agents to the merger itself, such as changes in submission strategies.

exceed the standing offer. The author signs with the imprint holding the highest valuation, at a price approximating the second-highest valuation. Alternatively, acquisition can occur through bilateral negotiation, either when a publisher pre-empts the auction with a take-it-or-leave-it offer or when it holds an option on the author's next work from a previous contract.[8]

The primary object of the negotiation is the advance: a lump-sum payment made by the publisher to the author at contract signing, unconditional on subsequent sales performance. Royalty rates are largely standardized within the industry and are not actively negotiated.[9] Royalties accrue to the author only after cumulative sales revenues recoup the advance, a threshold that only around 30% of titles reach (U.S. Department of Justice, 2022a). For the majority of titles, the advance therefore constitutes the entirety of the author's compensation. Approximately 93% of deals are won by the highest bidder (U.S. Department of Justice, 2022a), indicating that the advance is the dominant determinant of author choice and supporting the assumption, maintained throughout the analysis, that acquisitions can be modeled as price-determined auctions.

The model treats the acquisition of every title as an auction, following the DOJ's expert's own analysis of the market (U.S. Department of Justice, 2022a). That some titles are instead sold through direct bargaining is a legitimate concern. Section 2.1 and Appendix A show that the merger's effects are robust to the acquisition protocol. Representing non-auction allocation processes within an auction framework has precedent in settings such as local governments competing for firm relocations (Slattery, 2025), mortgage lending (Allen et al., 2019), and consumer credit (Cuesta and Sepúlveda, 2021).

One further institutional feature warrants mention: self-publishing. It is not a credible outside option for the authors and titles relevant to this paper. The DOJ presented evidence that established authors with substantial sales histories almost never self-publish (U.S. Department of Justice, 2022a). Consistent with this, self-published titles in my raw data account for 25% of observed titles but only 5% of total sales, concentrated far below the commercial segment in which the merging parties operate.[10] I therefore exclude self-publishing from the model and treat traditional publishers as the relevant market.

3 Data and Descriptive Statistics

The empirical analysis requires data on both sides of the market: the upstream transactions through which publishers acquire manuscripts, including the advances paid, and the downstream retail market in which the resulting books are sold, with the prices and quantities of each title. No single source covers both sides, so I combine two primary datasets: (i) deal-level records of rights acquisitions from Publishers Marketplace, which underpin both the reduced-form and structural analyses, and (ii) title-level Amazon sales and price data from Bookstat for January 2018 to August 2021, used to estimate the structural demand model. The structural estimation requires

[8] A standard contract clause grants the publisher the right of first refusal on the author's next title. The author submits first to that publisher and, if they agree on terms, does not offer the title to other firms.

[9] Standard royalty rates are roughly 7.5% on paperback editions, 15% on hardcover, and 25% on e-book editions.

[10] Self-published titles have median sales of 6 units, against 72 for traditionally published titles, between January 2018 and August 2021.

observing both the deal and the subsequent sales, so the working sample for the structural analysis is the intersection of the two datasets: titles for which I observe both the advance paid by the publisher and the retail sales. Sections 3.1 and 3.2 describe each dataset and the construction of the matched sample.

3.1 Upstream Market: Book Deals

The upstream data come from Publishers Marketplace (PM), a subscription-based platform on which literary agents announce book deals. PM is the standard source for systematic data on publishing transactions and has been used in related work by Peukert and Reimers (2022).[11] Coverage is most complete at the upper end of the advance distribution and for deals involving the Big Five publishers. Appendix C benchmarks the publisher composition of PM against the universe of comparable trade titles in Bookstat, the source I use to measure sales: the Big Five account for 41% of PM deals and 39% of Bookstat titles over the same window, with closely matched shares for each individual firm, indicating that PM is not selected on publisher identity.[12]

Book rights are typically sold as a bundle covering physical, electronic, and audio editions. The sample includes deals of this type in fiction, non-fiction, and children's categories, signed between 2008 and 2023. After dropping deals with missing publisher or category information, the analysis sample consists of 77,649 deal announcements, of which 41.8% involve a Big Five publisher. Multi-book contracts are recorded in PM as a single deal, and I treat them as such throughout: the unit of observation in the upstream data is the contract, not the title.[13]

A central limitation of PM is that advance amounts are not reported in dollars. Instead, each deal that discloses payment information is assigned to one of five categories—"nice," "very nice," "good," "significant," and "major"—each corresponding to a publicly known monetary range. Table 1 reports these ranges and the share of deals in each category. Approximately 23% of deals in the analysis sample carry a category assignment, the rest are reported without payment information. For deals with a disclosed category, I follow Peukert and Reimers (2022) and impute the advance as the midpoint of the category range, top-coding "major" deals (reported at \$500,000 and above) at \$750,000. For the remaining deals, I impute the advance from observable deal characteristics using the ensemble prediction model of Appendix D. The predicted advances are used only in the structural estimation, where they expand the sample available for identifying the auction-model parameters. The reduced-form analysis in Section 4 uses only deals with disclosed categories, reweighted to address selection into disclosure.

Table 2 reports the distribution of deals and mean advances by publisher group. PRH (or Penguin and Random House separately before the 2013 merger) alone accounts for 13.7% of disclosed deals, and SS for 5.2%. Panel (b) of the table restricts attention to deals with advances exceeding \$250,000, the threshold used by the DOJ to define an "anticipated best-seller" (U.S.

[11] Although reporting is voluntary, announcing a deal is widely viewed within the industry as a status signal: agents and authors have strong incentives to publicize successful sales, and major publishers report systematically through their publicists.

[12] Selection on advance levels cannot be checked directly, since advances are not observed outside PM. Section 3.2 documents selection in terms of sales, when comparing matched and unmatched deals.

[13] Multi-book contracts represent around 10% of the sample. In the matching between deals and sales described in Section 3.2, a multi-book contract is matched to each of the titles it covers downstream.

Department of Justice, 2022a). Within this segment, PRH and SS together account for 38% of deals, the concentration that is the proximate reason the proposed merger raised antitrust concerns, and the Big Five as a whole for more than 70%.

Table 1: Deal Categorization and Shares per Category

Deal Category	$ Range	Share
Nice	$\leq$ 50,000	51.4%
Very Nice	$>$ 50,000 & $\leq$ 100,000	11.5%
Good	$>$ 100,000 & $\leq$ 250,000	18.4%
Significant	$>$ 250,000 & $\leq$ 500,000	7.3%
Major	$>$ 500,000	11.5%

Notes: Deals data from Publishers Marketplace (18,321 disclosed deals out of 77,649 total deals between 2008 and 2023). Shares computed among disclosed deals. Deal midpoints used for advance imputation: Nice = $25,000; Very Nice = $75,000; Good = $175,000; Significant = $375,000; Major = $750,000.

Table 2: Descriptive Statistics at Publisher Level

Publisher	Share	Mean Advance
(a) Disclosed Deals (N = 18,321)		
Penguin Random House	0.137	$293,444
Simon & Schuster	0.052	$305,457
HarperCollins	0.108	$284,111
Hachette	0.032	$251,699
Macmillan	0.067	$239,256
Others	0.604	$93,444
(b) Disclosed Deals > $250,000 (N = 3,390)		
Penguin Random House	0.273	$606,081
Simon & Schuster	0.106	$606,494
HarperCollins	0.202	$594,846
Hachette	0.051	$599,564
Macmillan	0.100	$605,088
Others	0.269	$609,786

Notes: Deals data from Publishers Marketplace, 2008–2023. Observations at the deal level. Sample restricted to the deals with disclosed advance information (23.2% of total deals). Panel (b) further restricts to deals with advance above $250,000. "Others" includes all publishers outside the Big Five.

3.2 Downstream Market: Retail

The downstream data come from Bookstat and consist of daily ISBN-level records covering all editions (hardcover, paperback, e-book, and audio) of titles sold on Amazon US between January 2018 and August 2021. For each ISBN-day I observe units sold, transaction prices, the cumulative number of Amazon reviews and the average star rating, and time-invariant characteristics including publisher, imprint, genre, author, publication date, and binding format. The universe is restricted to titles that accumulated at least 10 reviews during the sample window and consists of approximately 200,000 unique ISBNs.

Combining these data with the deal-level records from PM is one of the empirical contributions of this paper. The matching is non-trivial because manuscripts do not have an ISBN at the time the rights are acquired.[14] I match deals to ISBNs through a multi-step procedure that combines title, author, publisher, and (approximate) publication-date matches against the Bookstat catalog, with manual review of the highest-value cases. Appendix B documents the procedure in detail.

Because the Bookstat universe covers only titles that clear a minimum review threshold on Amazon, the matched sample is selected toward more commercially visible books. Table C2 in Appendix C quantifies this selection by comparing PM deals matched to at least one Bookstat title against those matched to none. In fiction, 36.9% of reported deals match to at least one ISBN with positive sales. Matched deals have advances 17% higher on average than unmatched ones ($193,140 vs. $165,480), a 1.9 percentage-point higher best-seller share, and a 5.6 percentage-point higher Big Five share. The same patterns hold in non-fiction, where the overall match rate is lower (15.2%), likely reflecting the many non-fiction sub-categories (academic, professional, religious, regional) that fall outside the high-review trade segment Bookstat tracks.[15]

The matched sample, although selected, is appropriate for the welfare analysis, for a substantive reason rooted in the paper's welfare framework. In the model, the advance is a transfer from the publisher to the author rather than a creation of new surplus. The magnitude of any redistribution between publishers and authors that a merger induces is therefore bounded by the surplus the title generates. For a book that sells fewer than 100 copies and yields a few hundred dollars in publisher profit, even the most extreme redistribution, moving the entire profit from one party to the other, has negligible welfare consequences. For a book that generates millions in profit, the same proportional redistribution produces substantial welfare changes. The matched sample restricts attention to titles whose underlying surplus is large enough for redistribution between publishers and authors to be quantitatively meaningful. Titles outside the matched sample have low sales, low advances, and low underlying profits. Their omission reduces the count of authors whose surplus is analyzed but does not materially affect the total welfare calculation.

The structural estimation imposes two further restrictions on the matched sample. First, I drop titles with a median price below $2, or fewer than 20 units sold over the sample period, these are residual listings that do not represent active trade-publishing activity. Second, I exclude segments whose acquisition processes and sales patterns differ fundamentally from trade fiction and non-fiction: self-published titles (which never receive an advance), reference works, coloring books, and travel guides. Audio editions are also excluded.[16]

The unit of observation in the structural analysis is the title–month. A title aggregates all retained editions of the same work into a single product.[17] The aggregation reflects a substantive

[14] An ISBN is assigned only when the title is scheduled for publication, often a year or more after the deal.

[15] The match rate is also mechanically bounded by the temporal mismatch between the deals and sales windows: deals signed early in the window correspond to titles that may not have appeared in the Bookstat catalog until late in, or after, the sales window.

[16] Audio-book prices are recorded inconsistently in Bookstat, with a substantial fraction of observations missing or zero. Hardcover, paperback, and e-book editions are retained.

[17] An edition is a distinct published version of a work, identified by its own ISBN: each format (hardcover, paperback, e-book) constitutes at least one edition, and a format can contain several (e.g., a first and a revised printing, a movie tie-in paperback, a large-print version). Consumers choose among formats. Within a format,

feature of consumer choice: readers decide primarily whether to acquire a particular work, and only secondarily in which format. The demand model in Section 5.1 therefore treats titles, not editions, as the relevant unit of substitution.[18] Aggregation also materially reduces the dimensionality of the choice set, since not all titles are available in all formats, and many sell negligibly in some formats and substantially in others.[19]

Table 3 reports descriptive statistics for the fiction (Panel A) and non-fiction (Panel B) estimation samples at the title–month level. The Big Five publish 67% of fiction titles and 54% of non-fiction titles in the sample. PRH alone accounts for 22% of fiction titles and 23% of non-fiction titles, with the next-largest Big Five publisher being HarperCollins in fiction (18%) and Hachette in non-fiction (12%). Outside the Big Five, the market is fragmented: around 400 small publishers appear in each category, of which 88–90% publish fewer than five titles during the period. By sales, however, concentration is substantially higher: the Big Five account for 82% of units sold in fiction and 77% in non-fiction, reflecting the long-standing feature of this industry that a small number of firms publish a disproportionate share of commercially successful trade books.[20] The distribution of monthly sales is correspondingly skewed: the median fiction title sells 24 units in a typical month while the 99th percentile sells 2,756, whereas in non-fiction the corresponding figures are 29 and 3,526.

The final samples used to estimate the model consist of 4,882 fiction titles and 3,046 non-fiction titles tracked along the vertical chain between January 2018 and 2021.[21]

3.3 Additional Data Sources

Best-seller history. I supplement the deals data with an author-level indicator for prior appearance on the *New York Times* Hardcover Best-Seller list, constructed from the digitized lists compiled by Pruett (2022). Prior best-seller status is a strong predictor of both subsequent advances and subsequent sales. The indicator enters the reduced-form specifications and the construction of expected profits in the auction model. Throughout the paper, I refer to authors who appeared on the list before a given deal as "best-seller authors," a classification used in both the reduced-form analysis and the structural model. These data are used to complement the data from Publishers Marketplace, which already mentions whether an author has sold a best-seller title in the past. For those without this mention, I use Pruett (2022) to double-check.

I classify each manuscript into one of two segments based on the author's prior commercial record. *Best-sellers* are titles by authors who had previously published a *New York Times* best-seller; *non-best-sellers* comprise debut titles (the author's first book) and mid-list titles

editions are close to interchangeable.

[18] As a robustness check on this aggregation, Appendix L estimates a variant of the demand model in which print and e-book formats of the same work enter as separate products, so that substitution between formats is estimated rather than aggregated over. The estimated price sensitivity and the counterfactual effects of the merger are similar to the baseline. The title-level specification is retained because the work, not the format, is the primary object of consumer choice, and because it accommodates the titles sold in only one format without asymmetric treatment.

[19] For each title–month, the price is the average across editions weighted by units sold within the month.

[20] These shares are conditional on the estimation sample. Because Bookstat filters titles by review and the sample further requires a match to a deal, market shares computed over the full universe of titles would be lower.

[21] As explained later, some observations may be dropped in the supply-side estimation due to inconsistencies. This is explained in Section 7.2.

(by authors with prior books but no prior best-seller).[22] This author-based classification differs from the DOJ's advance-based definition of anticipated best-sellers (U.S. Department of Justice, 2022a), which would be unsuitable here because the dependent variable of interest is the advance itself.

In Table 4 I show the distribution of these three types of titles and the mean advances paid for each of the Big Five publishers and the pool of non-Big-Five publishers. The share of bestseller titles is higher among the Big Five than among other publishers, with the latter having a correspondingly larger share of midlist titles. Debut shares are similar across groups. In terms of advances, bestseller titles command advances substantially higher than those of debut and midlist titles. There is also a sizeable gap between Big Five and non-Big-Five publishers within each tier: while advances for bestsellers are broadly comparable across groups, Big Five publishers pay advances roughly three times larger than non-Big-Five publishers for both debut and midlist titles.

[22]This implies that the same author can have both best-seller and non-bestseller titles. Additionally, given that categories are based on author's status, those are exclusive categories, such that a debut cannot be a best-seller.

Table 3: Descriptive Statistics: Bookstat Sales Panel

Panel A: Fiction

	Mean	Median	p25	p75	p99
Units sold	188.28	24.00	6.00	100.00	2756.00
Price ($)	12.16	11.54	8.33	14.99	30.18
Months since publication	14.20	12.00	5.00	21.00	42.00

Genre (% of titles)		*Publisher (% of titles)*	
Literature	39.5	Penguin Random House	22.0
Teen & YA	27.2	HarperCollins	17.5
Mystery/Thriller	15.6	Hachette	8.4
Romance	8.8	Simon & Schuster	6.0
Sci-Fi & Fantasy	7.0	Macmillan	12.7
Comics & Graphic Novels	1.9	Other	33.3

Observations: 95,332 Titles: 4,882 Publishers: 401

Panel B: Non-fiction

	Mean	Median	p25	p75	p99
Units sold	249.27	29.00	7.00	118.00	3526.00
Price ($)	16.71	15.86	12.55	19.60	37.81
Months since publication	15.84	14.00	6.00	24.00	42.00

Genre (% of titles)		*Publisher (% of titles)*	
Politics/Soc. Sci.	30.0	Penguin Random House	23.1
Health & Self-Help	22.2	HarperCollins	8.4
Cookbooks	11.9	Hachette	11.8
Crafts & Arts	10.9	Simon & Schuster	5.1
Business & Tech	10.7	Macmillan	6.0
Other	14.4	Other	45.5

Observations: 67,537 Titles: 3,046 Publishers: 431

Notes: The sample corresponds to the demand estimation sample: titles matched to a deal in PM, with at least 20 total units sold and a median price above $2 over the sample period. Observations are at the title×month level for the period 01/2018-08/2021. Price is the monthly quantity-weighted average across a title's editions. Publisher shares consolidate imprints to Big Five parent groups, and "Other" includes all remaining publishers.

Cost shifters. For the demand estimation, I use monthly producer price indices for paper and for printing ink, and mean wages in the publishing sector, all from the U.S. Bureau of Labor Statistics. Interacted with title characteristics, these series serve as cost-side instruments for retail price in the demand model (Section 5.1).

4 Descriptive Evidence: Penguin–Random House Merger

In July 2013, Penguin and Random House merged to form Penguin Random House (PRH), reducing the Big Six to the current Big Five. This is the most recent prior consolidation of two top-tier U.S. trade publishers, and the natural retrospective benchmark for the blocked PRH–SS merger. To my knowledge, no existing work documents the merger's effects on author

Table 4: Deal Tier Composition and Mean Advance by Publisher Group

	N	**Bestseller**		**Debut**		**Midlist**	
		Share	Mean	Share	Mean	Share	Mean
PRH	2,501	13.9%	476.8	31.4%	268.5	54.7%	261.9
HarperCollins	1,789	13.4%	462.7	31.0%	277.3	55.6%	250.3
Simon & Schuster	939	15.8%	494.1	33.7%	284.7	50.6%	260.8
Hachette	586	10.4%	448.8	34.0%	238.1	55.6%	222.3
Macmillan	1,228	12.8%	388.1	28.5%	214.3	58.7%	221.2
Non-Big-5	11,278	4.7%	365.1	30.7%	89.1	64.6%	79.9
All	18,321	8.1%	425.9	30.9%	156.2	61.0%	138.3

Notes: Sample restricted to book deals with a disclosed advance, sourced from Publishers Marketplace. Tiers are mutually exclusive: *Bestseller* (NYT bestseller author), *Debut* (first-time author, non-bestseller), *Midlist* (all other non-bestsellers). Mean advance unweighted, in USD thousands (nominal).

compensation or publisher pricing.[23] Section 4 uses the upstream deals data to document how advances changed at PRH relative to other publishers, and across segments of the manuscript market that differ in how intensely the merging firms competed.

The evidence is descriptive rather than causal: a horizontal merger affects every firm in the industry simultaneously, directly for the merging parties, and through equilibrium responses for their rivals, so no group of publishers or titles is left unexposed, and the comparisons below are not difference-in-differences estimates of a treatment effect. The design instead exploits a contrast that follows from Section 2.1: the two channels make divergent predictions across segments of the manuscript market. The competition channel operates only where the merging firms plausibly competed head-to-head for the same manuscripts. However, the downstream channel operates everywhere, because all segments feed the same downstream product market. A pure-monopsony account therefore predicts advance declines concentrated at PRH and largest where the merging firms competed most intensely, with no movement elsewhere. The two-channel account predicts relative declines in that segment and relative gains in the rest, sharpest within PRH itself. Section 4 shows that the segments can be distinguished empirically, and that the observed patterns match the second prediction and are difficult to reconcile with the first.

Two further considerations support this reading. First, the timing of the merger was set by corporate decisions at the parent level, not by the trajectories of individual deals, and the specifications below condition on observable deal characteristics, so sharp changes in relative outcomes around the merger date plausibly reflect changes in firm behavior rather than composition. Second, the descriptive patterns motivate and constrain the structural model of Section 5, which must reproduce their signs and incidence, and which in turn provides the counterfactual decomposition the descriptive evidence cannot.

I document changes in advances paid by PRH relative to other publishers between 2008 and 2023, and across two segments of the manuscript market that differ in how intensely the merging parties competed before the merger.

[23] Shi (2024) studies the 2013 merger from a different angle, focusing on the division of non-monetary surplus between authors and publishers. That analysis does not examine advances or retail prices.

The design requires segments that differ in how much the merging firms competed for the same manuscripts. Bidding is not observed, so I use concentration of completed acquisitions as a proxy. Table 5 reports Herfindahl–Hirschman indices (HHI) for each segment before and after the 2013 merger. The index is computed each year from each publisher's share of the deals signed in that segment, and then averaged over the pre-merger and post-merger years. The unit is the deal as announced, and Penguin and Random House are treated as separate publishers before the merger.[24] The index measures the concentration of completed acquisitions, not bidder participation or head-to-head competition, which the data do not record.

The best-seller segment was already moderately concentrated before the merger (HHI of 1,659.5), and its HHI rose by 749.3 points after the merger, to 2,408.9: a change that the 2010 Horizontal Merger Guidelines, in force at the time, describe as potentially raising significant competitive concerns, and that the 2023 Guidelines would presume to be anticompetitive.[25] The non-best-seller segment is far less concentrated: its HHI rose by 88.9 points, to 447.5, well below the 1,500 level under which the 2010 Guidelines consider markets unconcentrated and mergers unlikely to have adverse competitive effects. The competition channel, if it operates anywhere, operates in the first segment, while the downstream channel operates in both, because both segments feed the same downstream product market.

Table 5: HHI and ΔHHI After the PRH Merger

Market	HHI pre-merger	HHI post-merger	ΔHHI
Best Seller	1659.5	2408.9	749.3
Non Best Seller	358.5	447.5	88.9

Notes: Herfindahl–Hirschman indices computed from publisher shares of deals signed in each segment and year, using Publishers Marketplace data. The unit is the deal; multi-book contracts count once. Pre- and post-merger values are averages of yearly indices over 2008–2013 and 2014–2023. Penguin and Random House are treated as separate publishers before the merger. The best-seller segment includes publishers with at least 50 deals with best-seller authors; the non-best-seller segment includes all firms with deals with non-best-seller authors. The index measures concentration of completed acquisitions, not bidder participation.

I trace the evolution of advances within each segment by estimating the following equation

$$y_{it} = X'_{it}\beta + \tau_t + \tau_{p(i)} + \tau_{a(i)} + \epsilon_{it} \tag{1}$$

separately for each segment, where i indexes a deal, t the year of the deal, X_{it} contains title and author characteristics (log number of prior books, a debut indicator, agency-size dummies, and genre), and τ_t, $\tau_{p(i)}$ and $\tau_{a(i)}$ are year, imprint and author fixed effects.[26] The dependent variable y_{it} is either the log advance, or an indicator for an advance above $500,000. Table 6

[24] The best-seller segment includes publishers that signed at least 50 deals with best-seller authors over the sample period; smaller publishers appear in this segment only occasionally, and their inclusion would add noise rather than competitors. The non-best-seller segment includes all publishers with at least one non-best-seller deals in the period. Because excluding small publishers mechanically raises the index, Appendix E reports it with alternative thresholds.

[25] The 2010 Horizontal Merger Guidelines state that mergers raising the HHI by more than 100 points in markets with a post-merger HHI between 1,500 and 2,500 potentially raise significant competitive concerns and often warrant scrutiny. The 2023 Merger Guidelines presume that a merger substantially lessens competition when the post-merger HHI exceeds 1,800 and the increase exceeds 100 points.

[26] Author fixed effects are only included in the regression for best-seller titles, where I observe authors repeatedly. For the rest, because most authors appear only once in the data, author fixed effects are not included.

summarizes the results as average post-merger changes: each cell reports the coefficient on a post-2013 dummy from a separate regression within the corresponding segment and publisher group. I run the design three ways: pooling all publishers, restricting to deals signed by PRH, and restricting to deals signed by the other Big Five publishers. The first specification asks how advances in each segment evolved across the industry after 2013, while the second and third ask how PRH's change differed from that of the rest of the Big Five. The estimated year effects $\hat{\tau}_t$ trace the corresponding dynamics, conditional on deal characteristics and publisher composition. Figure 2 plots them, normalizing 2013 to zero.

Table 6: Post-Merger Decomposition: Bestsellers, Non-Bestsellers, and Debut Authors

	All firms	**PRH only**	**Other Big-5**
	(1)	(2)	(3)
	(a) ln(advance)		
Bestseller	-0.147^{*}	-0.184^{*}	-0.094
	(0.075)	(0.103)	(0.129)
Non-bestseller	0.035	0.260^{***}	0.126
	(0.032)	(0.067)	(0.086)
Debut	0.051	0.234^{***}	0.105
	(0.045)	(0.088)	(0.116)
	(b) $P(\text{advance} > \$500{,}000)$		
Bestseller	-0.036	-0.118^{*}	0.045
	(0.042)	(0.070)	(0.064)
Non-bestseller	0.038^{***}	0.090^{***}	0.050^{***}
	(0.007)	(0.023)	(0.015)
Debut	0.063^{***}	0.144^{***}	0.068^{**}
	(0.012)	(0.033)	(0.027)
N (Bestseller)	1,318	322	589
N (Non-bestseller)	15,694	2,072	3,923
N (Debut)	5,054	746	1,394

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. SE clustered at imprint level. IPW-weighted. Each cell is a separate OLS regression of the outcome on a post-merger dummy (= 1 if year $\geq$ 2014), controlling for log prior books, and agency size dummies, with genre (17 categories) and imprint fixed effects. Regressions for best-seller cases use author fixed effects. The *Non-bestseller* row captures the trends of both *debut* and *mid-list*. Advances are expressed in real terms, deflated to 2012 USD.

Three patterns stand out. First, advances rose in the non-best-seller segment after 2013, with the increase concentrated at PRH: 26 log points, against 13 log points at other Big Five publishers, where the change is statistically indistinguishable from zero. Debut titles at PRH show a comparable increase of 23 log points. The probability of receiving an advance above \$500,000 in this segment rose industry-wide, by 4 percentage points, larger at PRH (9 points) than at other Big Five firms (5 points), and largest for PRH debuts (14 points).

Second, advances in the best-seller segment declined after 2013: by 15 log points pooled across firms, and by 18 log points within PRH, with no comparable change at other Big Five publishers. The pooled decline is therefore not an industry-wide movement, it is concentrated in the merged firm whose deals pull the pooled estimate down, while best-seller advances elsewhere in the Big Five are essentially flat. Panel (a) of Figure 2 shows this: the PRH series departs from

zero after 2013 while the corresponding series for the other Big Five does not. The probability of an above-$500,000 advance shows the same concentration: a 12-point decline within PRH (Panel (b)), again with no comparable change at other Big Five firms.

Third, the segment-specific divergence is sharpest within PRH. Across segments, PRH's advances rose by 26 log points for non-best-sellers and fell by 18 log points for best-sellers, a differential of more than 40 log points. The same comparison at other Big Five publishers shows a much weaker segment differential, with neither change distinguishable from zero.

The event studies also speak to the validity of the design. In both panels of Figure 2, the best-seller differential is flat in the years before the announcement and emerges only after 2013, which is the timing a merger-driven change implies and a pre-existing trend does not. In Appendix F, Figure F1 shows the same dynamics when the comparison group is restricted to debut authors alone.

Taken together, these patterns match the two-channel prediction of Section 2.1 and are difficult to reconcile with a pure-monopsony account. A pure-monopsony account predicts declines in advances at PRH, largest in the segments where the merging firms competed most intensely, and no movement elsewhere. What the data show is that advances at PRH rose most in the segment where the merging firms competed least, and declined in the segment where they competed most: a reduction in upstream competition that bites in the best-seller segment, and an increase in downstream pricing power that lifts advances across the other segments and is most visible where the competition channel is absent.

5 Structural Model

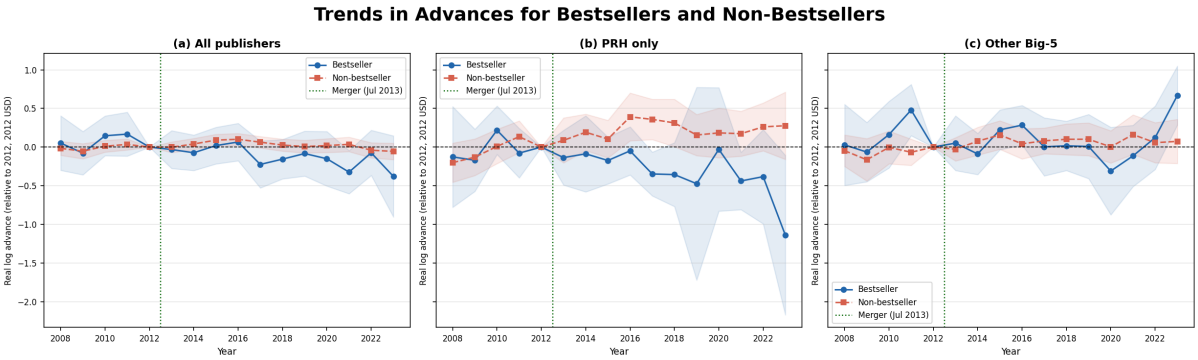
The evidence in Section 4 documents that the 2013 merger of Penguin and Random House was associated with changes upstream, in the patterns the two channels of Section 2.1 imply. To quantify the contribution of each channel and to evaluate the welfare consequences of the blocked PRH–SS merger in 2022, I develop a structural model of the publishing industry.

Publishers compete for the rights to manuscripts in ascending auctions, then choose retail prices for the titles they acquire, and consumers make purchase decisions based on those prices and titles' characteristics. The model is recursive: the auction outcomes determine each firm's portfolio, the portfolio determines the firm's pricing problem, and the expected profits from pricing discipline the firm's bids in the auction. A counterfactual merger propagates through this recursion, changing the set of bidders upstream and the intensity of competition downstream simultaneously, and the welfare consequences for authors, publishers, and consumers can be computed under the new equilibrium.

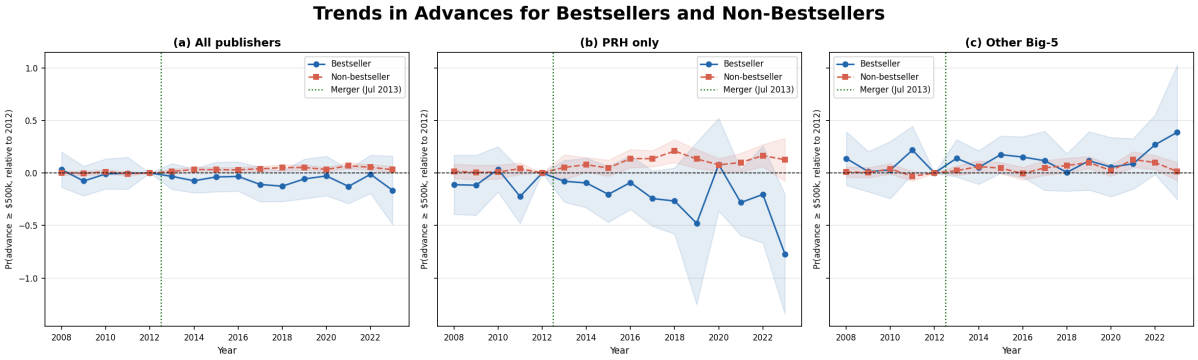
Figure 3 summarizes the timing. Before the game begins, publishers observe the set of manuscripts available in the period and form valuations. They then bid in ascending auctions, and the auction outcomes determine each publisher's portfolio. Publishers set retail prices for the titles in their portfolios, and consumers make purchase decisions. I describe each component in reverse order of the timing, since each stage's decisions are disciplined by the anticipated outcomes of the stages that follow: Section 5.1 develops the consumer demand system, Sections 5.2.1 and 5.2.2 the publisher pricing problem and the upstream auction, and Section 5.3 defines the

Figure 2: Event-study: Advances Paid to Best-seller Authors

(a) Log Advances



(b) Probability of an Advance Above $500,000



Notes: Event-study estimates for titles by best-seller authors relative to debut and mid-list authors, around the 2013 merger. Panel (a) reports the change in log advances (inflation-adjusted). Panel (b) reports the change in the probability that a deal exceeds $500,000. Within each panel, the three columns show, from left to right: all publishers, the merging parties, and the remaining Big Five. Standard errors are clustered at the imprint level. Estimates are inverse-probability-weighted to address selection into the disclosed-advance subsample.

welfare objects.

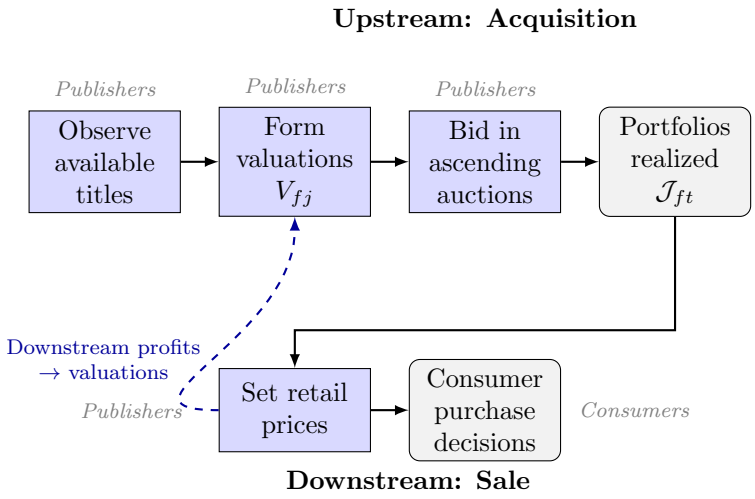
5.1 Demand

Consumers choose among differentiated titles in a nested-logit demand system with random coefficients. All demand-side objects are estimated separately for fiction and nonfiction, with each category constituting its own demand system, with its own parameters and markets.[27] The unit of observation is the title–month, where a title aggregates all retained editions of the same work, with prices and characteristics averaged across editions, weighted by units sold.[28] A market is a category–month pair over the sample window (January 2018 to August 2021).

[27]The separation reflects the segmentation of the industry, publishers organize imprints, editors, and acquisition decisions by category, and the upstream analysis of Section 6.2 is likewise estimated by category. It allows price sensitivity and substitution patterns to differ freely across the two. The cost is that direct fiction–nonfiction substitution is ruled out, the genre nests within each category capture the substitution margins the data can identify.

[28]The edition-level aggregation reduces the dimensionality of the choice set without imposing a substantive constraint, since consumers in the model choose at the title level: a reader decides whether to buy a particular novel, not which format of that novel to buy. Format choice within a title is treated as a separate decision outside the demand model.

Figure 3: Model Structure



Because Amazon prices and inventory are national, markets have no geographic dimension. Market size is the U.S. adult reading population, following Reimers and Waldfogel (2021).

Each period t, consumers choose whether to buy one of the available products in $\mathcal{J}$, or the outside option (not buying a book). The utility consumer i obtains from purchasing title j in period t is

$$U_{ijt} = \delta_{jt} + \mu_{ijt} + \zeta_{ig(j)} + (1-\rho)\epsilon_{ijt}, \tag{2}$$

where δ_{jt} is the mean utility component common across consumers, μ_{ijt} captures individual-level taste heterogeneity, $g(j)$ indexes the genre nest containing title j, and $\zeta_{ig(j)}$ is a nest-level taste shock with the unique distribution, implied by the nesting parameter ρ, such that $\zeta_{ig(j)}+(1-\rho)\epsilon_{ijt}$ is an i.i.d. Type I Extreme Value error. The outside option is no purchase in the period, normalized to mean utility zero and placed in its own nest.

The mean utility is

$$\delta_{jt} = \bar{\alpha} p_{jt} + \mathbf{Z}'_{jt}\boldsymbol{\gamma} + \xi_j + \xi_{g(j)t} + \xi_{jt}, \tag{3}$$

where p_{jt} is price, $\mathbf{Z}_{jt}$ is a vector of observed time-varying title–month characteristics, ξ_j and $\xi_{g(j)t}$ are title and month-genre fixed effects, and ξ_{jt} is the structural demand shock. The vector $\mathbf{Z}_{jt}$ contains the title's average star rating interacted with a Big Five indicator and with title's age. These interactions allow perceived-quality dynamics to differ between Big Five and other publishers, and across time. Because the specification includes title fixed effects, all time-invariant characteristics are absorbed by ξ_j and enter the model only through interactions with time-varying variables or through the random coefficients below. The structural demand shock ξ_{jt} is the source of endogeneity in the model: ξ_j absorbs persistent title-level quality, so ξ_{jt} captures month-to-month deviations from the title's average appeal that are observed by publishers and consumers but not by the econometrician.

Individual-level heterogeneity is

$$\mu_{ijt} = \sigma_p \nu_i^p \, p_{jt} + \sigma_{\text{best}} \nu_i^{\text{best}} x_{jt}^{\text{best}} + \sigma_{\text{debut}} \nu_i^{\text{debut}} x_{jt}^{\text{debut}} + \sigma_{\text{mid}} \nu_i^{\text{mid}} x_{jt}^{\text{mid}}, \tag{4}$$

where x_{jt}^{best}, x_{jt}^{debut}, and x_{jt}^{mid} are indicators for the author's status (best-seller, debut, or mid-list)

and ν_i^k are i.i.d. standard-normal draws. Consumer i's effective price sensitivity is therefore $\alpha_i = \bar{\alpha} + \sigma_p \nu_i^p$, distributed $N(\bar{\alpha}, \sigma_p^2)$. The author-status random coefficients capture heterogeneity in how consumers value books by authors of different commercial profiles (e.g., some consumers concentrate their purchases on established authors, others are more open to debuts) and are essential for generating realistic substitution patterns within and across the author-type segments studied in the reduced-form analysis and targeted by the merger counterfactual.

Genre nests group titles by content category, and the distribution of those was shown in Table 3. The nesting parameter $\rho \in [0, 1)$ governs within-nest substitution: $\rho = 0$ recovers standard logit substitution, while $\rho \to 1$ implies that a price increase on a title induces substitution almost entirely toward other titles in the same genre rather than toward other genres or the outside option.

The market share of title j in market t is obtained by integrating the choice probability over the distribution of individual taste shocks. Conditional on the random-coefficient draws $\boldsymbol{\nu}_i = (\nu_i^p, \nu_i^{\text{best}}, \nu_i^{\text{debut}}, \nu_i^{\text{mid}})$, the probability that consumer i chooses title j in market t is

$$s_{ijt}(\boldsymbol{\nu}_i) = \frac{\exp\left(\frac{\delta_{jt}+\mu_{ijt}}{1-\rho}\right)}{\left[\sum_{k\in g(j)} \exp\left(\frac{\delta_{kt}+\mu_{ikt}}{1-\rho}\right)\right]^{\rho} \left[\sum_{g\in\mathcal{G}} \left(\sum_{k\in g} \exp\left(\frac{\delta_{kt}+\mu_{ikt}}{1-\rho}\right)\right)^{1-\rho} + 1\right]}, \tag{5}$$

where the inner sum runs over titles within j's nest, the outer sum runs over all nests $\mathcal{G}$, and the trailing $+1$ is the outside option. The market share of title j in market t is

$$s_{jt} = \int s_{ijt}(\boldsymbol{\nu}_i)\, dF(\boldsymbol{\nu}), \tag{6}$$

I recover δ_{jt} numerically using the contraction mapping of Berry et al. (1995), extended to the random-coefficients nested logit as in Grigolon and Verboven (2014). The structural demand shock ξ_{jt} is then the residual of δ_{jt} after partialling out price, observed characteristics, and fixed effects, and is the source of endogeneity the instrumental-variables strategy in Section 6.1 addresses.

As a robustness exercise, I estimate a two-level nested logit model in the spirit of Reimers (2019) and Genakos et al. (2025), treating electronic and physical formats as distinct nests rather than aggregating across formats as in the baseline specification. Appendix L presents the demand estimates and Appendix R the corresponding counterfactual welfare effects, both confirming that the main findings under the RCNL specification are robust to this alternative modeling approach.

5.2 Supply

5.2.1 Downstream Problem

The supply side describes publishers' pricing and bidding decisions. Publishers acquire manuscripts in stage 1 to form their portfolios and set retail prices in stage 2. The advance paid for an acquired title is a sunk cost that does not affect retail pricing. The only source of uncertainty in the model is the demand shock ξ_{jt}, which publishers do not observe at the time of bidding but

observe before pricing. Following the backward-induction logic of the exposition, I begin with the pricing stage.

Stage 2: Pricing. Given its portfolio $\mathcal{J}_{ft}$ in market t, publisher f chooses retail prices $\{p_{jt}\}_{j\in\mathcal{J}_{ft}}$ to maximize variable profit:[29]

$$\max_{\{p_{jt}\}_{j\in\mathcal{J}_{ft}}} \sum_{j\in\mathcal{J}_{ft}} (p_{jt} - mc_{jt})\, s_{jt}(\mathbf{p}_t), \tag{7}$$

where $s_{jt}(\mathbf{p}_t)$ is the market share of title j in market t, which depends on its own price and the prices of all other titles in the market, collected in $\mathbf{p}_t$. The first-order conditions yield a system of equations that, in matrix form, delivers a closed-form expression for marginal costs:

$$\mathbf{mc}_t = \mathbf{p}_t + \left[\Theta \circ \nabla_{p_t}^{\top}\right]^{-1} \mathbf{s}_t, \tag{8}$$

where $\mathbf{s}_t$ is the vector of market shares, $\nabla_{\mathbf{p}_t}$ is the Jacobian of market shares with respect to prices ($\nabla_{\mathbf{p}_t,k\ell} = \partial s_{kt}/\partial p_{\ell t}$), Θ is the ownership matrix with $\Theta_{k\ell} = 1$ if titles k and ℓ are owned by the same publisher and 0 otherwise. The ownership matrix encodes the multiproduct nature of the pricing problem: a publisher internalizes the substitution between its own titles when setting prices, but not the substitution between its titles and those of rivals. Equation (8) recovers mc_{jt} for each title from observed prices and the demand-side estimates of $\nabla_{\mathbf{p}_t}$ and $\mathbf{s}_t$.

5.2.2 Upstream Problem

Stage 1: Bidding. Each month t, a set of new manuscripts Ω_t is offered to publishers, and each publisher $f \in \mathcal{F}$ participating in the auction for $j \in \Omega_t$ must decide how much it is willing to pay for the rights. For tractability, I assume publishers acquire on a monthly basis and that there is no delay between acquisition and publication, so that every title whose first unit is sold in month t is acquired at the start of that month.[30]

At the bidding stage, publisher f values title j at the marginal contribution the title would make to its portfolio, relative to the counterfactual in which a rival publishes it. If f's portfolio expands from $\mathcal{J}_{ft}^0$ to $\mathcal{J}_{ft}^0 \cup \{j\}$, its variable profit in market t becomes

$$\pi_{ft}\big(\mathcal{J}_{ft}^0 \cup \{j\}, \mathcal{J}_{-ft}\big) = \sum_{k\in\mathcal{J}_{ft}^0\cup\{j\}} (p_{kt} - mc_{kt})\, s_{kt}\big(\mathcal{J}_{ft}^0 \cup \{j\}, \mathcal{J}_{-ft}\big)\,, \tag{9}$$

where $\mathcal{J}_{-ft}$ denotes the union of rival portfolios in market t, held fixed at the bidding stage

[29] Under wholesale contracts, retailers rather than publishers set retail prices, a familiar issue in studies of brand competition using retail data (Nevo, 2000). The model is consistent with retailers setting prices if the retailer's markup is proportional, constant across titles, and unchanged by the merger, so that the publisher receives a fixed share λ of the retail price. The publisher's problem is then $\max \sum_j (\lambda p_{jt} - c_{jt})\, s_{jt}(\mathbf{p}_t) = \lambda \max \sum_j (p_{jt} - c_{jt}/\lambda)\, s_{jt}(\mathbf{p}_t)$, which has the same first-order conditions as (7): equilibrium prices are unchanged, the marginal cost recovered in (8) is c_{jt}/λ, and publisher profits are λ times the profits computed here. For e-books, the assumption holds by contract: under agency agreements the publisher sets the retail price and the retailer keeps a fixed commission. Appendix G considers bargaining between publishers and the retailer.

[30] In practice, titles are typically published 12 to 24 months after rights acquisition. Collapsing this timing into a single monthly decision abstracts from the dynamic problem of when to release an acquired title, which is outside the scope of this paper.

apart from the allocation of j. I assume that, if f does not acquire j, the title is published by the rival most specialized in its genre within f's segment (Big Five, non-Big Five), so that j remains in the market under either outcome and competes with other products from f's portfolio. Formally, let $g(j)$ denote j's genre and $\iota(f)$ denote f's size stratum (Big Five or non–Big Five); the counterfactual publisher is

$$f'(f,j) \;=\; \underset{f'\in\iota(f),\, f'\neq f}{\arg\max} \; n_{f't}(g(j))\,, \tag{10}$$

where $n_{f't}(g)$ is the number of titles firm f' has previously published in genre g as before t.[31] The marginal value of acquiring j, summed over the markets in which the title is sold, is

$$\Delta_{fj} = \sum_{t\geq t_0} \Big[\pi_{ft}\big(\mathcal{J}^0_{ft}\cup\{j\},\mathcal{J}_{-ft}\big) - \pi_{ft}\Big(\mathcal{J}^0_{ft},\,\mathcal{J}^{+j}_{-ft}\Big)\Big]\,, \tag{11}$$

where $\mathcal{J}^{+j}_{-ft}$ denotes the rival configuration with j appended to the portfolio of $f'(f,j)$, t_0 is the first month in which j is sold and sales are scaled to full-market, full-lifecycle units as described in Appendix M. Because j is present in the market under both branches of equation (11), the cannibalization of f's own catalog by j largely nets out, and Δ_{fj} captures the direct margin f earns on j together with the defensive value of denying the title to the rival that would otherwise publish it.

Two further restrictions are imposed for tractability. First, publishers do not re-solve the pricing problem in every future market when evaluating a candidate title. Firm f solves for the optimal price of j in its first month on sale and carries that price through the subsequent markets. The counterfactual publisher $f'(f,j)$ likewise sets a single optimized launch price for the title it acquires. This restriction is innocuous in this setting for two reasons. Observed prices are stable over a title's first months on sale, so a single optimized launch price is a good approximation to the price path that matters for the valuation, and at the bidding stage the publisher has not yet observed the demand shocks that would move future prices, so a bid rule based on a fixed launch price is the natural description of how valuations are formed. The restriction also reduces the computational burden substantially, since the Monte Carlo integration over demand shocks is solved in one market per title rather than in every market the title will be sold in.

Second, f takes rival portfolios $\mathcal{J}_{-ft}$ as fixed. This is reasonable in a market where individual portfolios contain hundreds of titles and a single acquisition has a small effect on aggregate competition. Appendix H documents that publisher acquisition patterns are stable across years, supporting the assumption that rivals do not shift their catalogs in response to f's actions.

The remaining source of uncertainty in Δ_{fj} is the demand shock ξ_{jt}, which publishers do not observe at the bidding stage. Expected profits are obtained by integrating Δ_{fj} over the empirical distribution $\hat{F}_\xi$ recovered from the demand estimation:

$$\mathbb{E}_\xi[\Delta_{fj}] = \int \Delta_{fj}(\xi)\, d\hat{F}_\xi(\xi). \tag{12}$$

[31] The rule assigns j to a publisher comparable in scale to f whose catalog it fits best: a Big Five bidder is replaced by the other Big Five publisher most active in j's genre, and a non–Big Five bidder by the most genre-specialized firm in the non–Big Five tail. Holding the market presence of j fixed across the two outcomes isolates the incremental value f captures by publishing j itself rather than ceding it to a competitor.

I assume the distribution of ξ_{jt} is stationary, so that $\hat{F}_\xi$ recovered from the 2018–2021 sample is the relevant distribution for the publisher's expectation regardless of the year of acquisition, and estimate a separate $\hat{F}_\xi$ for each title tier (best-seller, mid-list, and debut). This way, at the moment of bidding firms know title and author characteristics ξ_j, genre-time effects $\xi_{g(j)t}$ (i.e., seasonality per genre), but have uncertainty about demand shocks (ξ_{jt}).

Publishers have CARA preferences, $u(x) = -\exp(-r_h\, x)$, where $r_{h(j)}$ is a tier-specific coefficient of absolute risk aversion. If publisher f acquires title j at advance a, its net payoff is

$$x_{fj}(a) = \Delta_{fj} - FC_{fj} - a, \tag{13}$$

where the incremental profit Δ_{fj} is random because the demand shocks ξ are not yet realized, while the fixed cost FC_{fj} and the advance are known when bidding. The valuation V_{fj} is the highest advance f is willing to pay: the advance at which f is indifferent between acquiring j and not acquiring it,

$$\mathbb{E}_\xi[u(\Delta_{fj} - FC_{fj} - V_{fj})] = u(0). \tag{14}$$

Because FC_{fj} and V_{fj} are not random and CARA utility is exponential, they factor out of the expectation, and solving (14) gives

$$V_{fj} = CE_{r_{h(j)}}(\Delta_{fj}) - FC_{fj}, \tag{15}$$

where $CE_r(\Delta) \equiv -\frac{1}{r}\log \mathbb{E}_\xi\left[e^{-r\Delta}\right]$ is the certainty equivalent of Δ_{fj}, i.e., the sure amount that gives f the same expected utility as the risky profit. In words, f is willing to pay the certainty-equivalent value of the title's profits net of the cost of bringing it to market. For risk-averse publishers, $CE_{r_{h(j)}}(\Delta_{fj}) < \mathbb{E}_\xi[\Delta_{fj}]$, and the gap is the risk premium f requires to bear the uncertainty in j's sales.

When Δ_{fj} is normally distributed, the certainty equivalent takes the familiar mean–variance form, and (15) becomes

$$V_{fj} = \mathbb{E}_\xi[\Delta_{fj}] - \tfrac{r_{h(j)}}{2}\,\mathrm{Var}_\xi(\Delta_{fj}) - FC_{fj} \tag{16}$$

where $\frac{r_{h(j)}}{2}$ is the discount per unit of variance that f applies to a title of tier $h(j)$. When Δ_{fj} is not normal, (16) is the second-order (Arrow-Pratt) approximation to $CE_{r_{h(j)}}(\Delta_{fj})$.[32] The fixed cost FC_{fj} collects the costs of bringing j to market (e.g., marketing, design, editorial time) net of revenue not captured by Δ_{fj}. I then specify

$$FC_{fj} = \phi_{k(f)} + \psi_{h(j)} + \chi_{g(j)} + \varepsilon_{fj}, \tag{17}$$

[32]The derivation applies utility to the incremental profit from title j. Applying it instead to f's total portfolio profit leads to the same valuation under a mild condition. Let Π_f^0 denote f's portfolio profit when a rival publishes j, and $\Pi_f^1 = \Pi_f^0 + \Delta_{fj}$ its profit when f acquires j. The valuation is then the increase in the portfolio's certainty equivalent net of the fixed cost, $V_{fj} = CE_{r_{h(j)}}(\Pi_f^1) - CE_{r_{h(j)}}(\Pi_f^0) - FC_{fj}$. When portfolio profits are normally distributed, this equals $\mathbb{E}_\xi[\Delta_{fj}] - \frac{r_{h(j)}}{2}\left[\mathrm{Var}_\xi(\Delta_{fj}) + 2\,\mathrm{Cov}_\xi(\Pi_f^0, \Delta_{fj})\right] - FC_{fj}$, which coincides with (16) when the title's incremental profit is uncorrelated with the profit of the rest of the portfolio. Because a single title has a small effect on the profit of a catalog of hundreds of titles, the covariance term is small relative to $\mathrm{Var}_\xi(\Delta_{fj})$.

where $k(f)$ indexes the publisher group (each of the Big Five, and all other publishers pooled), and $\phi_{k(f)}$, $\psi_{h(j)}$, and $\chi_{g(j)}$ are publisher-group, tier, and genre fixed effects.[33] The term $\varepsilon_{fj} \sim N(0, (\sigma^{\varepsilon}_{k(f),h(j)})^2)$ is a publisher-specific cost shock, known to f but not to its rivals or the econometrician. It captures costs specific to the house, such as the editorial time the title requires given the imprint's current list and its fit with the publisher's marketing and distribution capacity.

Each title $j \in \Omega_t$ is sold in an ascending (English) auction. The data record the identity of the winning publisher and the transaction price, but not the identities of the losing bidders, the number of bidders, or the bids they submitted.

Substituting the fixed-cost specification (17) into the valuation (16), publisher f's valuation of title j is

$$V_{fj} = \underbrace{\mathbb{E}_{\xi}[\Delta_{fj}] + \beta_{h(j)} \operatorname{Var}_{\xi}(\Delta_{fj}) - \phi_{k(f)} - \psi_{h(j)} - \chi_{g(j)}}_{\equiv \bar{V}_{fj} \text{ (common knowledge)}} - \underbrace{\varepsilon_{fj}}_{\text{private value}}, \tag{18}$$

where $\beta_{h(j)} = -\frac{r_{h(j)}}{2}$. The component $\bar{V}_{fj}$ is common knowledge among bidders, while the cost shock ε_{fj} is known only to f. Across titles, valuations are correlated because publishers share components of $\bar{V}_{fj}$: a title with higher expected demand raises every publisher's valuation. Within an auction, however, the only source of variation in valuations conditional on the observable characteristics X_j of the title and the bidders is the cost shock, which is drawn independently across publishers. Letting $\mathcal{N}_j$ denote the set of N_j bidders for title j, valuations are therefore independent but not identically distributed:

$$V_{fj} \mid X_j \sim \mathcal{N}\left(\bar{V}_{fj},\, \sigma^2_{k(f),h(j)}\right), \qquad f \in \mathcal{N}_j. \tag{19}$$

This is an independent private values (IPV) environment: each publisher's private information affects only its own valuation. The substantive restriction is that no publisher holds private information bearing on its rivals' valuations, such as a private signal about the title's demand, which would generate a winner's curse. Such information is limited here because the manuscript or proposal, the author's track record, and sales of comparable titles are observed by every publisher invited to bid.[34]

Because the advance is paid with certainty and values are private, remaining active until the price reaches one's valuation is a weakly dominant strategy for any degree of risk aversion. Let $V_{(1)j} \geq V_{(2)j} \geq \cdots \geq V_{(N_j)j}$ denote the ordered valuations of the bidders in $\mathcal{N}_j$. The winner w_j is the publisher with the highest valuation, and the advance equals the valuation of the last rival to drop out:

$$V_{w_j j} = V_{(1)j}, \qquad a_j = V_{(2)j}. \tag{20}$$

Given (19), the likelihood of each observed pair (w_j, a_j) has a closed form: the winner's valuation exceeds a_j, one rival's valuation equals a_j, and all remaining rivals' valuations fall below it.

[33] Because fixed costs are defined net of unmodeled revenue, these components can take either sign.

[34] With only winners and transaction prices from ascending auctions, common-value and affiliated-value models are not identified (Athey and Haile, 2002). Appendix J relaxes independence by letting ε_{fj} share a title-level component known to all bidders; because the degree of correlation is not identified from one price per auction, I impose it at several levels as a sensitivity analysis.

Section 6.2 derives the likelihood.

Bidder set. The data do not record which publishers bid in each auction, so the bidder set must be specified. This requires two choices: which publishers are potential rivals for a given title, and how many bidders each auction has. Neither is identified from winners and prices alone, so I fix both by assumption and report the sensitivity of the results to each.

I assume the upstream market is segmented by publisher type: titles are contested either among Big Five publishers or among other publishers, but not across the two groups. The segment is determined before the auction, by the agent's submission strategy and the profile of the title, and the type of the winning publisher reveals it. Two institutional features support this assumption. First, agents submit widely and the Big Five are almost always among the recipients of a submission, so a title acquired by a smaller publisher is one for which no Big Five publisher chose to compete. Second, when a Big Five publisher does compete, smaller publishers rarely outbid it, consistent with the higher Big Five group effects estimated in Section 7. For each auction, the winner's rivals are drawn from the potential bidders in the winner's segment. Rivals are drawn without replacement, with probability proportional to each publisher's share of titles acquired in the segment; the largest publisher is therefore the most frequent rival.[35]

For the number of bidders, the baseline sets $N_j = 3$ in every auction, so that a title contested by both merging parties retains one rival after the merger. As sensitivity analyses, I re-estimate the model with $N_j = 2$, under which jointly contested titles lose all upstream competition after the merger (the most pessimistic case for authors), and with N_j increasing in the advance, reflecting stronger competition for more expensive titles. Section 8 and Appendix O report the counterfactuals under each specification.

Throughout, the model treats the supply of manuscripts as exogenous, with authors' output not responding to advances, such that the merger does not change the number of titles offered. I do not model manuscript production in order to keep the model tractable. Appendix I shows that the number of titles published by PRH did not change relative to other publishers after the 2013 merger, which is consistent with this assumption.

5.3 Welfare

The two sides of the model generate surplus for three sets of agents: consumers, who derive utility from the titles they purchase; publishers, who earn variable profit on sales net of the advances they pay to acquire manuscripts; and authors, whose compensation is the advance. This subsection defines each object and states the welfare accounting identity that the counterfactual analysis in Section 8 applies. The identity is a property of the model at any market structure, it is not specific to the merger.

Consumer surplus. Given the nested-logit demand system of Section 5.1, consumer i's expected utility in market t is the inclusive value of the choice set, and aggregate consumer

[35] As a robustness check, I also draw rivals from clusters defined by the title's tier and genre, with probabilities proportional to each publisher's past output in that tier–genre cell, so that publishers specialized in a genre are more likely to compete for titles in it. This mainly affects the composition of rival sets among smaller publishers, which tend to specialize. Results are reported in Appendix P.

surplus is obtained by integrating over the random coefficients and scaling by market size M_t:

$$CS_t \;=\; M_t \int \frac{1}{|\alpha_i|} \log\left(1 + \sum_{g \in G} \left[\sum_{k \in g} \exp\left(\tfrac{\delta_{kt}+\mu_{ikt}}{1-\rho}\right)\right]^{1-\rho}\right) \mathrm{d}F(\nu), \tag{21}$$

where $\alpha_i = \bar{\alpha} + \sigma_p \nu_i^p$ is consumer i's price coefficient and the leading 1 is the outside option.[36] Total consumer surplus is $CS = \sum_t CS_t$.

Publisher surplus. Publisher f's variable profit in market t is $\Pi_{ft}^{\text{var}} = \sum_{j \in \mathcal{J}_{ft}} (p_{jt} - mc_{jt})\, s_{jt}\, M_t$, as in the pricing problem (7). Summing over markets and netting the advances paid for the titles in its portfolio, publisher f's surplus is

$$\Pi_f \;=\; \sum_t \Pi_{ft}^{\text{var}} \;-\; \sum_{j \in J_f} a_{fj}, \tag{22}$$

where a_{fj} is the advance paid for title j bu firm f and $\mathcal{J}_f$ is the set of titles f acquires. I treat fixed costs of production as unaffected by the counterfactual and omit them from (22). The reported profit changes are therefore gross of fixed costs and should be read as changes in variable profit net of advances.

Author surplus. Because royalties are standardized and, for the large majority of titles, never recouped against the advance, the advance is the entirety of aggregate advances. Also, I abstract from any non-pecuniary author surplus from the publisher match. Hence, the total welfare of authors across the $\mathcal{J}$ products is:

$$A \;=\; \sum_{j \in \mathcal{J}} a_j, \tag{23}$$

the sum of winning advances across acquired titles.

Total welfare and the role of the advance. Aggregate welfare is the sum of the three components. Because each acquired title's advance enters once as a cost to the winning publisher in (22) and once as author income in (23), the advances cancel, and total welfare reduces to consumer surplus plus variable profit:

$$W \;=\; CS \;+\; \sum_f \Pi_f \;+\; A \;=\; \underbrace{\sum_t CS_t}_{\text{consumers}} \;+\; \underbrace{\sum_{f,t} \Pi_{ft}^{\text{var}}}_{\text{publishers}}. \tag{24}$$

The advance is a transfer from publishers to authors: it determines how the surplus generated by a title is divided between the firm and the author, but it does not enter total welfare. Equation (24) holds for any allocation of titles to publishers and any vector of advances. Reallocating a title or changing its advance changes the split between Π_f and A, and changes CS and Π^{var} through the ownership matrix Θ, but the form of the identity is unchanged.

[36] With a normally distributed price coefficient, surplus is not well-defined for consumers whose marginal utility of income is zero or positive. This does not arise here: the estimated dispersion of the price coefficient is small (0.011 in fiction, and zero in non-fiction), so every simulated consumer has a negative price coefficient bounded away from zero.

This decomposition is the lens through which I read the counterfactual. A merger affects total welfare only through the downstream terms: the change in consumer surplus and in variable profit that follows from the merged firm's enlarged, jointly-priced catalog. Its effect on advances, by contrast, is a redistribution between publishers and authors that nets out of W. I therefore report four objects: the change in consumer surplus, ΔCS; the change in variable profit, $\Delta\Pi^{\text{var}}$, whose sum with ΔCS is the total welfare effect; the change in aggregate advances, $\Delta A = \sum_j \Delta a_j$, which is the transfer; and publishers' retained surplus, $\Delta\Pi^{\text{var}} - \Delta A$, the welfare-relevant profit figure once the transfer to authors is accounted for.

The role of the valuation in welfare. It is worth being explicit about what the publisher valuation V_{fj} does and does not contribute to welfare. The valuation governs who acquires a title and at what advance, but it is not itself a component of surplus. Its common-knowledge part is anchored by f's expected incremental profit from publishing j rather than ceding it to a rival, $\mathbb{E}_\xi[\Delta_{fj}]$. The profit the title actually generates is counted in Π^{var} once it is sold, whichever publisher owns it, and its idiosyncratic part (ε_{fj}) never enters welfare at all. The reason is the auction mechanism. Authors sell to the highest bid, so ε_{fj} determines only the identity of the winner and the level of the advance, both of which are settled before books are sold. The title's contribution to surplus is exhausted by the consumer surplus and variable profit it subsequently generates in the product market, neither of which depends on the signal. A high shock lets a publisher win a manuscript and may raise the advance it pays, but it does not make the resulting book more valuable to consumers or more profitable to produce. The private value is a device that allocates titles and prices them, not a source of value.

This is the key structural difference from settings such as Slattery (2025), in which firms competing to be acquired by a buyer (e.g., a state government) confer value on that buyer's externalities so that the buyer's valuation is itself a welfare primitive. There, the auction allocates over the sum of firm profit and buyer value, both of which are real surplus, and the bidder's valuation enters total welfare. In the publishing market, the acquirer is the same agent that monetizes the title downstream, and the manuscript generates value for the publisher only through sales. The author confers no separate, non-pecuniary surplus on the firm beyond the profit the book earns. Consequently the valuation, and in particular its idiosyncratic component, shapes the distribution of surplus between publishers and authors through the advance, but leaves total welfare to be determined entirely by the downstream product market.

6 Estimation

6.1 Demand

I estimate the demand system by GMM, following Berry et al. (1995) and using the `pyblp` implementation of Conlon and Gortmaker (2020). The structural parameter vector is $\theta = (\theta_1, \theta_2, \rho)$, where $\theta_1 = (\bar{\alpha}, \boldsymbol{\gamma})$ collects the linear parameters on price and the characteristics in $\mathbf{Z}_{jt}$, $\theta_2 = (\sigma_p, \sigma_{\text{best}}, \sigma_{\text{debut}}, \sigma_{\text{mid}})$ collects the standard deviations of the random coefficients. Finally, ρ is the nesting parameter governing within-genre correlation. Fiction and non-fiction are estimated separately.

For each candidate value of (θ_2, ρ), estimation proceeds in two steps. First, I invert the observed market shares to recover the mean utility $\delta_{jt}(\theta_2, \rho)$ for each title–month, using the contraction mapping of Berry et al. (1995) extended to the random-coefficients nested logit as in Grigolon and Verboven (2014). Second, given δ_{jt}, the structural demand shock is the residual after partialling out price, observed characteristics, and fixed effects:

$$\xi_{jt}(\theta_2) = \delta_{jt}(\theta_2) - \bar{\alpha} p_{jt} - \mathbf{Z}'_{jt}\boldsymbol{\gamma} - \xi_j - \xi_t. \tag{25}$$

The GMM objective interacts ξ_{jt} with the instrument vector described next, and the estimator searches over (θ_2, ρ) to minimize it. Two objects on the right-hand side of equation (25) are endogenous: the price p_{jt}, which publishers set after observing ξ_{jt}, and, implicitly, through the share inversion, the within-group share that identifies ρ, which is a function of the demand shocks of all titles in the nest. I therefore require instruments for both, and construct them as follows.

I use two sets of excluded instruments. The first consists of three cost shifters: the monthly producer price indices for paper and for printing ink, and the mean monthly wage in the publishing sector, each interacted with the title's page count. Production costs vary over time with input prices and across titles with physical size, so these interactions shift marginal cost differentially across titles within a month. Because the specification includes month fixed effects, the aggregate level of input prices is absorbed, and identification comes from cross-title differences in exposure to cost movements. These instruments shift prices but do not directly shift within-nest shares, so they mainly identify the price coefficient $\bar{\alpha}$.

The second set follows Berry (1994) and Gandhi and Houde (2019) and measures how close each title is to its rivals within the same nest. For each characteristic x^k (page count, months since publication, and star rating), let $\mathcal{R}_{jt} = \mathcal{J}_{g(j)t} \setminus \{j\}$ denote the other titles in j's genre nest in month t. I construct

$$z^k_{jt} = \sum_{j' \in \mathcal{R}_{jt}} \left(x^k_{jt} - x^k_{j't}\right)^2, \tag{26}$$

$$z^{kl}_{jt} = \sum_{j' \in \mathcal{R}_{jt}} \left(x^k_{jt} - x^k_{j't}\right)\left(x^l_{jt} - x^l_{j't}\right), \tag{27}$$

Here z^k_{jt} measures how far j is from its rivals in characteristic k, and z^{kl}_{jt} captures how differences in two characteristics move together. A title whose rivals are close in characteristics faces tighter competition, which affects both its markup, and hence its price, and its share within the nest. These instruments are therefore informative about both $\bar{\alpha}$ and the nesting parameter ρ, and their interactions separate the random-coefficient dispersions θ_2 from ρ. Finally, the number of titles in j's nest in month t shifts the within-nest share directly and mainly identifies ρ.

6.2 Supply

I estimate the parameters of the valuation equation (18) by maximum likelihood, using the identity of the winning publisher and the transaction price (the advance) for each auctioned manuscript. Collecting the parameters as $\theta^A = \left(\{\beta_h\}, \{\phi_{k(f)}\}, \{\psi_{h(j)}\}, \{\chi_{g(j)}\}, \{\sigma^{\varepsilon}_{k(f),h(j)}\}\right)$, the

estimation sample consists of the fiction and non-fiction matched deals described in Section 3.2, with the advance measured as the disclosed-category imputed value of Appendix D. As discussed in Section 5.2.2, the baseline sets $N_j = 3$ in every auction, with the rivals of the observed winner drawn from the corresponding market segment. Appendix O reports results under $N_j = 2$ and under a number of bidders that increases with the advance.

Under independent private values in an ascending auction, remaining active up to one's valuation is weakly dominant, so the transaction price equals the second-highest valuation, $a_j = V_{(2)j}$. Let w denote the winning publisher in auction j and $\mathcal{N}_j$ the set of participants. The likelihood contribution of auction j is the joint density of the event that the winner's valuation exceeds a_j, some rival's valuation equals a_j, and all remaining rivals' valuations fall below a_j:

$$\mathcal{L}_j(\theta) = \underbrace{[1 - F_w(a_j\,;\theta)]}_{\substack{\text{winner's valuation} \\ \text{exceeds } a_j}} \underbrace{\sum_{r \in \mathcal{N}_j \setminus \{w\}} f_r(a_j\,;\theta) \overbrace{\prod_{k \in \mathcal{N}_j \setminus \{w,r\}} F_k(a_j\,;\theta)}^{\substack{\text{remaining rivals} \\ \text{below } a_j}}}_{\substack{\text{one rival's valuation} \\ \text{equals } a_j}}. \tag{28}$$

where $F_f(\cdot)$ and $f_f(\cdot)$ are the CDF and PDF of publisher f's valuation, implied by the normality of ε_{fj} around the deterministic component $\bar{V}_{fj}$ in equation (18). The sum over r accounts for the fact that the identity of the price-setting runner-up is not observed: any of the winner's rivals may be the one whose valuation equals the price. The overall log-likelihood is

$$\ell(\theta) = \sum_j \log \mathcal{L}_j(\theta), \tag{29}$$

and I estimate θ^A by maximizing $\ell(\theta^A)$, with standard errors computed from the numerical Hessian.

The estimation is a two-step procedure. The inputs $\mathbb{E}_\xi[\Delta_{fj}]$ and $\mathrm{Var}_\xi(\Delta_{fj})$ entering $\bar{V}_{fj}$ are computed in a first step from the demand estimates of Section 6.1 and the Monte Carlo procedure of Section 5.2.2, and are treated as data in the auction likelihood rather than estimated jointly.

6.3 Identification

Consider the demand side first. The parameters requiring instruments are the price coefficient $\bar{\alpha}$ (and its dispersion σ_p) and the nesting parameter ρ. The threat to each is the structural demand shock ξ_{jt}. The specification's fixed effects define what that shock is: ξ_j absorbs persistent title-level quality and $\xi_{g(j)t}$ absorbs aggregate demand conditions for each genre, so ξ_{jt} is the transitory, title-specific component of appeal (e.g., a review in a prominent outlet, a media appearance, a viral recommendation) observed by publishers and consumers but not by the econometrician. Publishers set prices after observing ξ_{jt}, so price is correlated with it, and OLS on equation (3) would attenuate $\bar{\alpha}$ toward zero.

The cost-side instruments address this by isolating price variation driven by production costs rather than demand. Their relevance rests on cost pass-through: when paper, ink, or publishing-sector wages rise, the marginal cost of longer books rises by more, and Bertrand–Nash pricing transmits the differential into retail prices. Their exclusion rests on two features

of the specification working together. First, month fixed effects absorb the aggregate level of input prices, so any correlation between input-price cycles and aggregate book demand is irrelevant. Second, page count itself is time-invariant and absorbed by the title fixed effect, so the instrument cannot proxy for a preference for longer books. The identifying requirement is only that the transitory demand shock ξ_{jt} is uncorrelated with the interaction of aggregate input-price movements and page count. This is, publicity shocks and viral moments do not systematically hit long books precisely when paper is expensive.

The within-group share that identifies ρ is endogenous by construction: it is a function of the demand shocks of every title in the nest. The nest-composition instruments are valid here for a reason specific to this industry's timing. A title's month-t nest composition is determined by acquisition and publication decisions taken 12 to 24 months earlier, and by the time ξ_{jt} realizes, the set of competing titles and their characteristics are predetermined. Nest composition can therefore shift a title's within-group share without responding to the title's transitory demand shock. The same predetermined variation identifies the random-coefficient dispersions: as Gandhi and Houde (2019) emphasize, how a title's share responds to the proximity of rivals in characteristics space is informative about taste heterogeneity, and the interactions of the nest instruments provide the corresponding moments.

The dispersions σ_{best}, σ_{debut}, and σ_{mid} are identified from how substitution patterns differ across author-status segments, whether a price increase on a best-seller sends consumers to other best-sellers or across segments, which is exactly the substitution the merger counterfactual must get right. The remaining linear parameters in θ_1 are identified from within-title variation over the life-cycle, conditional on the fixed effects.

On the supply side, the auction structure converts two observables per manuscript (i.e., the winner's identity and the transaction price) into information about the valuation distributions. Under the model, the price is the second-highest of N_j draws and the winner is the publisher with the highest draw, so each auction contributes an order-statistic observation from a known family. Which publisher wins, and how often, conditional on the observable components of valuations, identifies the publisher fixed effects $\phi_{k(f)}$: a firm that wins more auctions than its expected profits warrant must have a higher intercept. The level of transaction prices relative to expected profits identifies the tier intercepts $\psi_{h(j)}$, with the scale anchored by the unit coefficient on $\mathbb{E}_\xi[\Delta_{fj}]$ in equation (16): valuations are denominated in dollars of expected variable profit, so the intercepts and risk parameters inherit dollar units rather than being identified only up to scale.

The risk parameters β_h are identified from how advances covary with profit variance conditional on profit mean: two titles with equal $\mathbb{E}_\xi[\Delta_{fj}]$ but different $\text{Var}_\xi(\Delta_{fj})$ command different advances if and only if bidders price risk. This requires independent variation in the two moments. In the data, the conditional variation is substantial, driven by differences in genre volatility and life-cycle profiles among titles with similar expected sales. Finally, the signal variances $\sigma^{\varepsilon}_{k(f),h(j)}$ are identified from the dispersion of transaction prices around their predicted means within publisher–tier cells, together with the frequency of "upset" wins in which a publisher with a lower observable valuation prevails: cost-shock dispersion generates both wider price residuals and more frequent upsets.

These arguments rest on one maintained assumption, which parallels the exogeneity conditions on the demand side: the private signal is mean-independent of the observable components of the valuation,

$$\varepsilon_{fj} \perp \mathbb{E}_\xi[\Delta_{fj}], \mathrm{Var}_\xi[\Delta_{fj}], \chi_{g(j)} \tag{30}$$

so that $\bar{V}_{fj} = \mathbb{E}[V_{fj} \mid X_{fj}]$ and all systematic variation in valuations loads on the observables. Economically, the condition requires that the cost shock ε_{fj} be unrelated to a title's measured profitability once tier, genre, and publisher are conditioned on. The main way it could fail is if publishers observe a component of commercial potential that the demand model does not capture and that correlates with $\mathbb{E}_\xi[\Delta_{fj}]$. The validation exercise of Section 7.3, in which the model reproduces the competitive patterns implied by the DOJ experts' confidential bid-log data, provides an external check that the maintained structure captures the market's actual competitive behavior.

Two limits of identification are worth stating plainly. The number of bidders N_j is not identified from data on winners and prices alone (Athey and Haile, 2002), it is fixed by assumption, and Appendix O describes the sensitivity analysis. And the first-step inputs $\mathbb{E}_\xi[\Delta_{fj}]$ and $\mathrm{Var}_\xi(\Delta_{fj})$ are generated regressors, so the supply-side identification is conditional on the demand model. Misspecification upstream propagates downstream, which is an additional reason the external validation in Section 7.3 carries weight.

7 Results

7.1 Demand

Table 7 reports the demand estimates for fiction and non-fiction. The mean price coefficient has the expected signs, with consumers disliking higher prices. Price-sensitiveness coefficient in Fiction markets is roughly twice as large as in Non-fiction markets, consistent with the denser field of close substitutes in fiction. First-stage IV results are reported in Appendix K. The nesting parameter is high in both categories ($\rho = 0.69$ in fiction, 0.77 in non-fiction): when a title's price rises, consumers substitute overwhelmingly toward other titles in the same genre rather than toward other genres or the outside option. This matters directly for the merger simulation, with the closer the substitution among titles, the more of a price increase's lost demand is recaptured by the merged firm's own catalog.

Among the random coefficients, the dispersion in tastes for best-selling authors is large and, in fiction, precisely estimated ($\sigma_{\text{best}} = 9.17$): some consumers concentrate their purchases on established authors while others are close to indifferent to best-seller status.[37] This heterogeneity is what allows the model to generate realistic substitution within and across the author segments at the center of the counterfactual. The dispersions on price and on debut status are estimated at or near zero. The data do not detect taste heterogeneity along these dimensions, and the model effectively collapses to homogeneous preferences there.

The estimates imply median own-price elasticities of -6.1 in fiction and -6.3 in non-fiction, more elastic than the near-zero estimates of Reimers and Waldfogel (2017), Reimers and Waldfogel

[37] In dollar terms, the implied dispersion in willingness to pay for best-seller status in fiction is $\sigma_{\text{best}}/|\bar{\alpha}| \approx \$49$.

(2021), and Kawaguchi et al. (2025), but well below the very elastic demands in Genakos et al. (2025) and Ghose et al. (2006). The elasticity level matters beyond demand fit, because it disciplines the supply side: markups are inversely related to elasticities, and per-book markups must be large enough that expected downstream profits can rationalize the advances observed upstream. A demand system implying near-perfectly elastic demand would leave margins too thin for any realistic sales volume to justify six-figure advances.

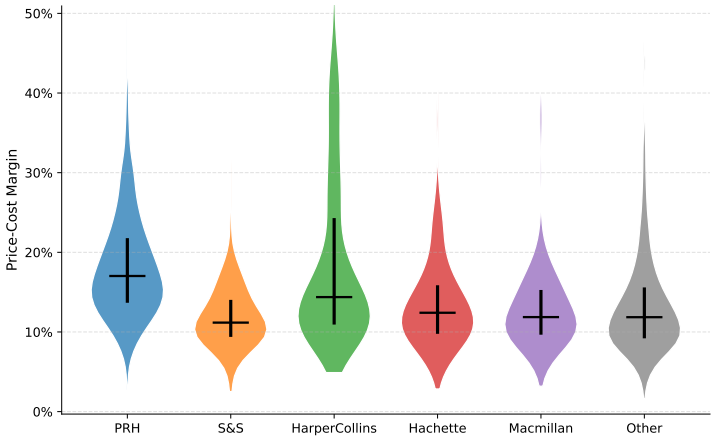
Figure 4 shows the distribution of price–cost margins by publisher, recovered from the first-order conditions in equation (8). The median margin is 17%, above the single-product Lerner benchmark of $1/|\eta| \approx 15\%$ implied by the median elasticities, reflecting the multi-product internalization encoded in the ownership matrix and the tight within-genre substitution. Across firms, PRH has the highest median margin and Simon & Schuster the lowest among the Big Five, a ranking that foreshadows the asymmetry in the firms' valuations estimated below. Appendix N presents the results from the cost-function estimation.

Table 7: Demand Estimates

	Fiction		Nonfiction	
	Estimate	(SE)	Estimate	(SE)
A. Mean utility				
Price ($\bar{\alpha}$)	−0.189	(0.060)	−0.091	(0.005)
B. Standard deviations of random coefficients (σ)				
Price	0.045	(0.029)	0.000	(0.000)
Bestseller	9.167	(1.524)	3.044	(1.124)
Debut	0.490	(3.900)	0.000	(0.000)
Midlist	1.237	(1.196)	0.000	(5.383)
C. Nesting parameter				
Within-genre correlation (ρ)	0.691	(0.028)	0.771	(0.011)
sProduct FE	Yes		Yes	
Genre × Market FE	Yes		Yes	
Median own-price elasticity	−6.12		−6.26	
Observations	95,332		67,537	
Products	4,882		3,046	
Firms	6		6	
Genres (nests)	6		7	

Notes: Robust standard errors in parentheses. Product and genre-market fixed effects are absorbed. Instruments are Gandhi–Houde differentiation IVs, cost shifters and number of products per group. Random coefficients are specified on price and the three book-type dummies (bestseller, debut, midlist). ρ is the within-nest correlation parameter.

Figure 4: Price-cost Margins by Firm



Notes: Distribution of price–cost margins $(p_{jt} - mc_{jt})/p_{jt}$ at the title–month level, with marginal costs recovered from the Bertrand–Nash first-order conditions in equation (8) using the demand estimates of Table 7. "Other" pools all publishers outside the Big Five.

7.2 Supply

Tables 8 and 9 report the estimated parameters of the valuation equation (18) for fiction and non-fiction. As on the demand side, the two categories are estimated separately, since what publishers value in a manuscript may differ across them. All parameters are denominated in millions of dollars, with the scale anchored by the unit coefficient on expected profits.

The firm fixed effects measure how much each publisher is willing to pay beyond what a title's measured profitability implies. In non-fiction, PRH stands clearly apart: its fixed effect (2.06) is roughly four times HarperCollins' (0.52), the second highest, with the remaining Big Five clustered below. In fiction, HarperCollins has the highest fixed effect (1.64), with PRH second (1.08) and Simon & Schuster the lowest among the Big Five (0.25). In fiction, SS sits at the bottom of the Big Five while it's the third-highest in non-fiction, which foreshadows the asymmetry of the merger's effects in Section 8.

The tier intercepts are estimated conditional on expected profits and profit variance, so they must be read as ceteris paribus objects: they do not say that mid-list titles command higher advances than best-sellers, but rather that, holding measured profitability fixed, Mid-list titles have the highest intercept in both categories, the lowest is best-sellers in fiction and debuts in non-fiction. The ordering is natural once the intercepts are understood as absorbing tier-specific costs and values that variable profit does not capture. Acquiring a best-seller entails commitments that a mid-list acquisition does not (e.g., larger marketing campaigns, tours, and production outlays) so a best-seller expected to sell no more than a mid-list title is a strictly worse purchase: the publisher bears a best-seller's fixed costs for a mid-list's revenues. Debuts carry the largest promotional burden of all, since demand must be built from no existing readership. Read this way, the intercepts partly reflect tier-specific fixed costs on willingness to pay, and their ordering, is what a model with unmeasured, tier-varying fixed costs would predict.

Publishers are risk-averse in every tier of both markets: all six variance coefficients are negative, so greater profit variance lowers a publisher's willingness to pay, holding expected profit

fixed. The incidence of risk aversion differs across categories. In fiction, variance is penalized most heavily for best-sellers (-1.41) and least for mid-lists (-0.46): conditional on paying a best-seller premium, publishers want predictability, while tolerating dispersion in the middle of the list. In non-fiction debuts are penalized most (-1.34) and mid-list least (-0.73).

The cost-shock standard deviations $\sigma^{\varepsilon}_{k(f),h(j)}$ are large relative to the systematic components of valuations: within publisher–tier cells, they range from roughly 0.26 to 1.93 in fiction and 0.29 to 1.95 in non-fiction. Most of the dispersion in valuations across bidders for the same title therefore comes from the cost shocks rather than from systematic differences across firms, which is precisely what generates upsets, price dispersion within observably similar auctions, and the runner-up variation that the merger counterfactual reallocates.

Figures 5 and 6 compare the distribution of simulated second-highest valuations with the observed advances. The model matches mean prices closely, overpredicting by 8.0% in fiction and 8.5% in non-fiction. The simulated distribution is more skewed than the observed one: it places more mass near the lowest tier while tracking closely the upper tail, where the auctions that drive the welfare analysis are concentrated. Since the counterfactual's welfare magnitudes are dominated by mid- and high-advance titles (as discussed in Section 3.2), the fit where it matters is the upper tail.

Figure 5: Distribution of Observed and Predicted Prices in Fiction

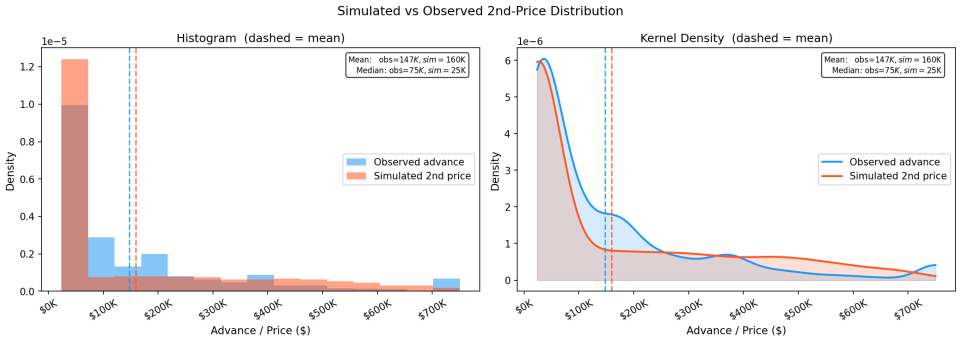


Figure 6: Distribution of Observed and Predicted Prices in Non-fiction

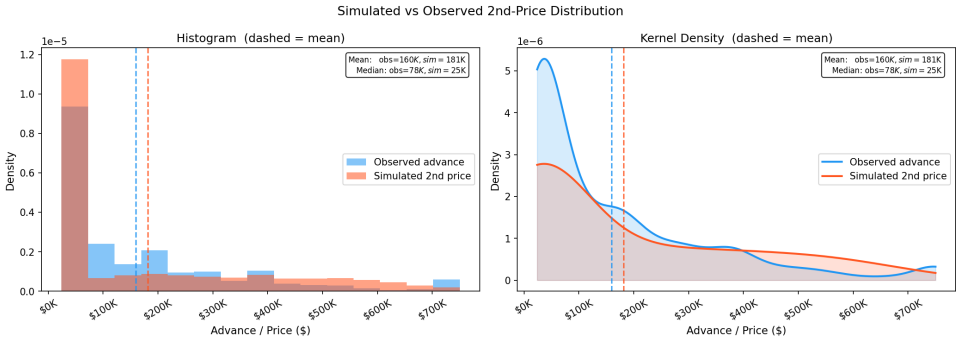


Table 8: Structural Auction Model: Estimated Parameters for Fiction Titles

		Estimate	Std. Error
A. Tier intercepts			
Midlist		0.044	(0.033)
Debut		-0.193	(0.032)
Bestseller		-0.328	(0.058)
B. Firm fixed effects (ref: non-Big5)			
Hachette		0.776	(0.102)
HarperCollins		1.644	(0.070)
Macmillan		0.427	(0.035)
Penguin Random House		1.077	(0.033)
Simon & Schuster		0.246	(0.039)
C. Observable covariates			
β_2: X2=Var (midlist)		-0.457	(0.000)
β_2: X2=Var (debut)		-0.537	(0.013)
β_2: X2=Var (bestseller)		-1.412	(0.124)
D. Signal SDs ($\sigma^{\varepsilon}_{f\tau}$)			
non-Big5	Midlist ($\sigma^{\varepsilon}_{\text{nb5,mid}}$)	1.654	(0.036)
	Debut ($\sigma^{\varepsilon}_{\text{nb5,deb}}$)	0.673	(0.035)
	Bestseller ($\sigma^{\varepsilon}_{\text{nb5,best}}$)	0.769	(0.093)
Hachette	Midlist ($\sigma^{\varepsilon}_{\text{Hac,mid}}$)	1.162	(0.058)
	Debut ($\sigma^{\varepsilon}_{\text{Hac,deb}}$)	1.181	(0.115)
	Bestseller ($\sigma^{\varepsilon}_{\text{Hac,best}}$)	0.893	(0.128)
HarperCollins	Midlist ($\sigma^{\varepsilon}_{\text{HC,mid}}$)	1.928	(0.059)
	Debut ($\sigma^{\varepsilon}_{\text{HC,deb}}$)	1.557	(0.016)
	Bestseller ($\sigma^{\varepsilon}_{\text{HC,best}}$)	1.171	(0.077)
Macmillan	Midlist ($\sigma^{\varepsilon}_{\text{Mac,mid}}$)	0.851	(0.008)
	Debut ($\sigma^{\varepsilon}_{\text{Mac,deb}}$)	0.556	(0.048)
	Bestseller ($\sigma^{\varepsilon}_{\text{Mac,best}}$)	0.488	(0.074)
Penguin Random House	Midlist ($\sigma^{\varepsilon}_{\text{PRH,mid}}$)	1.358	(0.000)
	Debut ($\sigma^{\varepsilon}_{\text{PRH,deb}}$)	0.992	(0.049)
	Bestseller ($\sigma^{\varepsilon}_{\text{PRH,best}}$)	0.740	(0.053)
Simon & Schuster	Midlist ($\sigma^{\varepsilon}_{\text{S\&S,mid}}$)	0.661	(0.040)
	Debut ($\sigma^{\varepsilon}_{\text{S\&S,deb}}$)	0.262	(0.041)
	Bestseller ($\sigma^{\varepsilon}_{\text{S\&S,best}}$)	0.394	(0.050)
Auctions		4,341	
Log-likelihood		-8,054.2	
Pred./actual price ratio		1.080	

Notes: Maximum-likelihood estimates of the valuation equation (18) for fiction titles. All parameters are in millions of dollars, with the scale anchored by the unit coefficient on expected variable profit $E_\xi[\Delta_{fj}]$ (not reported). Standard errors are from the numerical Hessian (observed Fisher information matrix); they condition on the first-step demand estimates, treating $E_\xi[\Delta_{fj}]$ and $\text{Var}_\xi(\Delta_{fj})$ as data, and so do not reflect first-stage sampling error. The predicted/actual price ratio is the mean simulated second-highest valuation divided by the mean observed advance. Entries shown as 0.000 are below 0.0005.

Table 9: Structural Auction Model: Estimated Parameters for Non-fiction Titles

		Estimate	Std. Error
A. Tier intercepts			
Midlist		0.501	(0.072)
Debut		0.152	(0.058)
Bestseller		0.276	(0.082)
B. Firm fixed effects (ref: non-Big5)			
Hachette		0.379	(0.062)
HarperCollins		0.524	(0.099)
Macmillan		0.087	(0.044)
Penguin Random House		2.062	(0.070)
Simon & Schuster		0.462	(0.123)
C. Observable covariates			
β_2: X2=Var (midlist)		-0.728	(0.000)
β_2: X2=Var (debut)		-1.338	(0.019)
β_2: X2=Var (bestseller)		-0.848	(0.005)
D. Signal SDs ($\sigma_{f\tau}$)			
non-Big5	Midlist ($\sigma_{\text{nb5,mid}}$)	1.954	(0.046)
	Debut ($\sigma_{\text{nb5,deb}}$)	0.939	(0.009)
	Bestseller ($\sigma_{\text{nb5,best}}$)	0.523	(0.065)
Hachette	Midlist ($\sigma_{\text{Hac,mid}}$)	0.772	(0.008)
	Debut ($\sigma_{\text{Hac,deb}}$)	0.654	(0.051)
	Bestseller ($\sigma_{\text{Hac,best}}$)	1.558	(0.234)
HarperCollins	Midlist ($\sigma_{\text{HC,mid}}$)	1.024	(0.062)
	Debut ($\sigma_{\text{HC,deb}}$)	0.586	(0.056)
	Bestseller ($\sigma_{\text{HC,best}}$)	1.708	(0.249)
Macmillan	Midlist ($\sigma_{\text{Mac,mid}}$)	0.555	(0.042)
	Debut ($\sigma_{\text{Mac,deb}}$)	0.286	(0.039)
	Bestseller ($\sigma_{\text{Mac,best}}$)	0.777	(0.009)
Penguin Random House	Midlist ($\sigma_{\text{PRH,mid}}$)	1.458	(0.000)
	Debut ($\sigma_{\text{PRH,deb}}$)	1.209	(0.051)
	Bestseller ($\sigma_{\text{PRH,best}}$)	1.374	(0.112)
Simon & Schuster	Midlist ($\sigma_{\text{S\&S,mid}}$)	1.005	(0.078)
	Debut ($\sigma_{\text{S\&S,deb}}$)	0.610	(0.071)
	Bestseller ($\sigma_{\text{S\&S,best}}$)	1.079	(0.204)
Auctions		2,816	
Log-likelihood		-5,183.3	
Pred./actual price ratio		1.085	

Notes: Maximum-likelihood estimates of the valuation equation (18) for non-fiction titles. All parameters are in millions of dollars, with the scale anchored by the unit coefficient on expected variable profit $E_\xi[\Delta_{fj}]$ (not reported). Standard errors are from the numerical Hessian (observed Fisher information matrix); they condition on the first-step demand estimates, treating $E_\xi[\Delta_{fj}]$ and $\text{Var}_\xi(\Delta_{fj})$ as data, and so do not reflect first-stage sampling error. The predicted/actual price ratio is the mean simulated second-highest valuation divided by the mean observed advance. Entries shown as 0.000 are below 0.0005.

7.3 Validation

Before turning to the full counterfactual, I use the model to reproduce the exercise the DOJ's expert witnesses conducted in the PRH–SS case, which serves as an external check on the auction model. The DOJ's experts estimated the merger's effect on advances using a confidential dataset of publishers' internal bid logs and their own empirical framework (U.S. Department of Justice, 2022a). Their analysis treats the upstream market in isolation, so the only operative channel is the removal of one merging party from auctions in which both participated. I replicate that exercise within my model by shutting down the downstream feedback: valuations are held at their baseline values, and the merger's only effect is to replace PRH and SS with a single bidder wherever both were present. Because the two analyses share neither data nor structural assumptions, agreement between them is informative about whether the model captures the competitive structure of the manuscript market.

Table 10 reports the results under the baseline $N_j = 3$, while Appendix Q.1 reports the case of $N_j = 2$ and the specification with N_j increasing in the advance. Panel B confirms the isolation is doing what it should: advances change only in auctions where both merging parties participated, and are exactly unchanged elsewhere, the downstream channel that would move them is switched off. Within the affected set, the harm is substantial, and diluting over each firm's full portfolio of acquisitions yields firm-level declines of 6.8% on PRH titles and 22.7% on SS titles, with total aggregate advances falling by \$8.7 million (−4.3%).

The DOJ's experts estimated firm-level declines of 4–7% at PRH and 12–19% at SS. My estimate for PRH falls inside their range, and both analyses deliver the same asymmetry, with SS's authors being harmed far more than PRH's, for the same reason: PRH, the larger firm with systematically higher valuations, is the runner-up in SS's auctions far more often than SS is in PRH's, so removing PRH from SS's auctions bites harder than the reverse. My estimate for SS exceeds the top of their range, which likely reflects the runner-up frequencies implied by the bidder-set assumption: with rivals drawn from the Big Five and PRH holding the highest mean valuations, the model makes PRH SS's runner-up almost mechanically, whereas the DOJ's bid logs record the actual frequency. Note the direction of this discrepancy: overstating PRH's presence as SS's runner-up overstates the harm to SS's authors, so the model errs on the side of the DOJ's own concern.

The agreement in both magnitude and asymmetry, obtained from independent data and independent modelling choices, is evidence that the auction model recovers the competitive structure that the court's own experts identified. The next section shows how these estimates change when the downstream channel is switched back on and the merger's effect on pricing power is allowed to transmit into bids.

The decomposition in Panel B also shows what the firm-level averages conceal: the entire effect is borne by the jointly contested titles, where advances fall by 39.0%, while every other auction is untouched. The DOJ's firm-level figures average this concentrated harm over each party's full catalog, the model recovers the incidence underneath the average.

Table 10: Merger Counterfactual With Upstream Isolation (N=3)

Panel A: By baseline auction winner			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
PRH	-\$14,655	-4.4	-6.8%
Simon & Schuster	-\$39,184	-4.2	-22.7%
Hachette	\$0	0.0	0.0%
HarperCollins	\$0	0.0	0.0%
Macmillan	\$0	0.0	0.0%
Non-Big5	\$0	0.0	0.0%
Overall	-\$7,097	-8.7	-4.3%
Panel B: By merging firms' participation in baseline			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
Both PRH & SS	-\$85,023	-8.7	-39.0%
One of PRH/SS	\$0	0.0	0.0%
Neither	\$0	0.0	0.0%
Overall	-\$7,097	-8.7	-4.3%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

8 Counterfactual Merger: PRH-SS

8.1 Setting

I simulate the merger between Penguin Random House and Simon & Schuster, blocked by the U.S. antitrust authorities on monopsony grounds. In the model, the merger operates through the two channels of Section 2.1. Downstream, the merged entity prices its enlarged portfolio jointly, which raises equilibrium prices and changes every publisher's expected profit from every title. These changes translate into willingness to pay in the upstream market. Upstream, PRH and SS no longer compete independently: in auctions where both firms participated in the baseline, they are replaced by a single merged bidder whose valuation is the higher of the two constituent firms'.

The counterfactual is run on titles first sold on or after January 1, 2021. The merger agreement was reached in late 2020, so 2021 is the first year in which acquisitions would have been made by the merged entity. Under the model's assumption of no gap between acquisition and publication (Section 5.2.2), the titles acquired by a merged PRH–SS are exactly those first appearing in 2021. Restricting the treated set in this way keeps the exercise free of the release-timing dynamics the model abstracts from: titles already published before 2021 were acquired under the pre-merger market structure.[38]

The simulation proceeds in three steps, mirroring the model's timing in reverse. First, I

[38] In practice, titles published in 2021 would have been acquired one to two years earlier, before the agreement. The exercise should be read as comparing the market for the 2021 cohort under the two ownership structures, not as a prediction of the transition path.

re-solve the downstream pricing equilibrium under the merged ownership matrix for all titles on sale in 2021, holding demand primitives and marginal costs fixed: each title retains its own recovered marginal cost mc_{jt}, so the merged entity operates PRH's titles at PRH's costs and SS's titles at SS's costs, with no merger-specific efficiencies.[39] This yields counterfactual prices, shares, consumer surplus, and variable profits. Second, I recompute every publisher's expected profit $\mathbb{E}_\xi[\Delta_{fj}]$ and profit variance for every 2021 title under the new pricing equilibrium.[40] Through equation (16), these shift all publishers' valuations. Third, I re-run each auction with the updated valuations and the post-merger bidder set, and recover the counterfactual advances.

The third step requires each bidder's baseline valuation, which the data do not reveal directly. The ascending format, however, implies tight restrictions on where each valuation lies, and I recover them by their conditional expectations under the estimated truncated normals. The runner-up's valuation is known exactly: in an ascending auction the runner-up's exit determines the transaction price, so $\hat{V}_{rj} = a_j$, where the runner-up is taken to be the non-winning bidder with the highest mean valuation $\bar{V}_{fj}$. The winner's valuation must exceed a_j, so I assign its conditional expectation

$$\hat{V}_w \;=\; \mathbb{E}[V_w \mid V_w > a_j] \;=\; \mu_w + \sigma_w \frac{\phi\left(\frac{a_j - \mu_w}{\sigma_w}\right)}{1 - \Phi\left(\frac{a_j - \mu_w}{\sigma_w}\right)}, \tag{31}$$

and each losing bidder's valuation must fall below a_j, so I assign $\hat{V}_{kj} = \mathbb{E}[V_{kj} \mid V_{kj} < a_j]$ analogously. In the counterfactual auction, each surviving bidder's valuation is its recovered baseline value plus the change in its common-knowledge component from step two. The counterfactual advance is the second-highest valuation among the surviving bidders, floored at \$25,000, which is the midpoint of the lowest advance category, below which no reported deal is observed.

The merged entity's primitives are constructed as follows. For titles both firms contested in the baseline, I compute each constituent firm's valuation and assign the merged bidder the higher of the two: the merged entity inherits, title by title, the identity of whichever arm values the manuscript more.[41] For titles only one merging firm contested, the merged bidder is simply that firm, with its valuation updated for the counterfactual pricing equilibrium like every other bidder's, nothing else about those auctions changes. Titles neither firm contested keep their baseline bidder sets, and their advances move only through the industry-wide revaluation of step two. The construction assumes the merger creates no valuation primitive beyond the constituents': the merged entity is, for each manuscript, its higher-valuing arm.

Because the identities of the losing bidders are drawn rather than observed, the exercise does not deliver title-level effects. It does deliver effects by auction type (i.e., titles for which both merging parties bid, titles for which one did, and titles for which neither did) which is the

[39] It is conservative for the paper's central result: any true cost synergy would raise the merged firm's per-title profits further, amplifying the downstream channel and the advance increases it generates.

[40] The counterfactual publisher in equation (10) is determined under the merged ownership structure. When the most specialized publisher for a title bid on by PRH is SS, or the reverse, the title is assigned to the next most specialized Big Five publisher. For non-merging bidders, the merged entity is treated as a single publisher, with the genre counts of PRH and SS combined.

[41] In most jointly contested auctions this is PRH, reflecting its larger firm fixed effect.

decomposition the two channels make meaningful. Consistent with the reduced-form evidence of Section 4, the jointly contested auctions are concentrated among best-seller titles, while the merging parties rarely both participate in auctions for debut or mid-list manuscripts.

Welfare is computed over each title's full commercial life and the full U.S. market. Because the sales data cover only Amazon and observe each title for only part of its life cycle,[42] observed sales are scaled by the market-coverage and life-cycle adjustment factors of Appendix M before the welfare objects of Section 5.3 are evaluated.

In the counterfactual, the winner of an auction can change. Two forces can reallocate a title. First, the merged firm may value some titles less than its constituent firms did: a title one of the merging firms acquired in the data can now take sales from the other firm's catalog, lowering the merged firm's valuation enough for a rival to win. Second, higher downstream prices raise every publisher's valuation, and a rival whose valuation rises by more than the baseline winner's can overtake it. In both cases the title goes to the bidder with the highest counterfactual valuation, and the advance equals the second-highest.

8.2 Results

Table 11 reports the merger's effects on advances, and Tables 12 and 13 its effects in the retail market, all under the baseline of $N_j = 3$ bidders per auction, while Appendix O reports the alternative bidder specifications. The two sets of outcomes cover different titles. Retail prices, consumer surplus, and variable profits are computed for all titles on sale between January and August 2021, including titles published earlier, since the merged firm reprices its entire catalog; these amounts are scaled to the full U.S. market and to each title's remaining commercial life as described in Appendix M. Advances are computed only for the 2021 cohort, the titles first sold in 2021, since only these are acquired after the merger. Percentage changes are relative to the baseline value of the corresponding set of titles.

Consolidation enlarges the merged firm's jointly priced catalog and softens retail price competition. Consumer surplus falls by \$19.9 million (−2.5%), with a larger proportional loss in fiction (−3.4%) than in non-fiction (−2.0%). The merged firm's variable profit rises by \$8.2 million (6.6%), but the non-merging publishers gain proportionally more, between 9% and 11%: rivals benefit both from the industry-wide softening of price competition and from the demand that substitutes away from the merged firm's higher-priced titles, while bearing none of the cannibalization within the merged catalog. Industry variable profit rises by \$17.0 million (8.1%) in total, and welfare falls by \$2.9 million (−0.3%), because the consumer-surplus loss exceeds the industry profit gain.

The downstream profit gains feed back into manuscript valuations and therefore into advances, and the sign of the feedback depends on whether the merging firms competed in a given auction. Panel B of Table 11 groups auctions by the merging firms' baseline participation. For manuscripts in which both PRH and SS were active bidders, advances fall by 19.9%: the merger removes one of the bidders disciplining the price, and the competition channel dominates. The magnitude is telling, however, when set against the isolated-upstream exercise of Section 7.3, in which the

[42] Titles published in July 2021, for example, are observed for only two months, while titles published earlier are observed only over the later part of their life cycles.

same auctions lose 39.0%: allowing the downstream gains to feed into the surviving bidders' valuations offsets about half of the competition harm even where that harm is concentrated. For manuscripts in which only one of the merging firms participated, advances rise by 15.6%, and where neither participated, by 11.3%: in these auctions the merging parties were not setting the price, so the only operative force is the downstream channel, which lifts every publisher's willingness to pay as expected per-book profits rise across the market.

Because the jointly contested segment is small, its losses (−$4.4 million) are far smaller than the gains elsewhere ($24.8 million), and total aggregate advances rise by $20.3 million (10.2%), where the isolated-upstream analysis predicts a decline of $8.7 million (−4.3%) as shown in Section 7.3. A monopsony-only assessment does not merely understate the aggregate effect on authors, it mis-signs it.

Panel A, which groups auctions by the baseline winner, shows the same logic from the firms' side. Advances on titles won by SS fall most (−13.3%), because PRH—the larger firm, with systematically higher valuations—was disproportionately the runner-up in SS's auctions, so its removal as an independent bidder bites hardest there (Section 7.3). Advances on PRH's own titles rise by a modest 5.2%: the average blends gains on the titles PRH won without SS's participation with losses on the jointly contested set. Advances on titles won by the non-merging publishers rise by 14.7% to 22.2%, precisely because none of their titles carry the contested-segment losses. Their gains are the downstream channel operating alone.

Netting advances from variable profit gives the change in publishers' surplus over the year: the variable profit they earn on every title on sale in 2021, net of the advances they pay for the titles they acquire that year. On this measure, publishers' retained surplus falls by $3.3 million, as advances rise by $20.3 million against an increase in variable profit of $17.0 million.

The incidence of the retained-surplus decline is itself asymmetric, and it reverses the standard prediction for non-merging firms. Downstream, rivals free-ride on the merger in the usual way (Deneckere and Davidson, 1985): their variable profits rise proportionally more than the merged firm's. Upstream, the same revaluation that raises their profits raises their bids, and rivals pay $19.4 million more in advances with every non-merging publisher losing on net. On the other hand, the merged entity retains approximately $7.3 million. The merger is privately profitable, and its producer-side burden falls on the firms that did not merge: an input-market counterpart to raising rivals' costs (Salop and Scheffman, 1983), arising here as an equilibrium externality of the downstream consolidation rather than as an exclusionary strategy. A product-market analysis alone would count these rivals among the merger's beneficiaries. The joint analysis shows them bearing most of its producer-side cost, and supplies a reason, distinct from authors' interests, why rival publishers might oppose such a transaction.

Table 13 reports the average retail price changes from the counterfactual merger. Prices of titles owned by the merged entity rise by about 2.3% on average, in both fiction and non-fiction, as the merged firm internalizes substitution between the PRH and SS catalogs. Other publishers' prices rise by less than 0.5%: they respond to the merged firm's higher prices, but have no new substitution to internalize. These price changes are the source of the consumer surplus loss and the variable profit gains reported in Table 12.

Endogenous supply of manuscripts. The counterfactual holds fixed the set of manuscripts offered and the publishers' capacity to acquire them, so it abstracts from several margins through which a merger could affect authors and readers. Higher advances for uncontested titles could draw additional manuscripts, which would add variety for consumers but also intensify downstream competition and lower the profit, and hence the advance, per title. Lower advances on jointly contested titles could reduce the output of established authors. The merged firm could also change how many titles it acquires or how much it invests in editing and marketing them, and agents could change which publishers they submit to. These responses can move aggregate advances and welfare in either direction, and modeling them is left for future work. As a check on their likely size, the number of titles published by PRH did not change relative to other publishers after the 2013 merger (see Appendix I).

Table 11: Merger Counterfactual: Upstream Market (N=3)

Panel A: By baseline auction winner			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
PRH	\$11,358	3.4	5.2%
Simon & Schuster	-\$22,880	-2.5	-13.3%
Hachette	\$26,565	2.5	22.2%
HarperCollins	\$31,377	6.5	17.7%
Macmillan	\$30,609	4.0	16.5%
Non-Big5	\$16,776	6.4	14.7%
Overall	\$16,622	20.3	10.2%
Panel B: By merging firms' participation in baseline			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
Both PRH & SS	-\$43,382	-4.4	-19.9%
One of PRH/SS	\$29,178	16.9	15.6%
Neither	\$14,524	7.9	11.3%
Overall	\$16,622	20.3	10.2%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

9 Conclusion

This paper studies what merger analysis looks like when a firm's power in its input market and its power in its product market are evaluated jointly rather than in isolation. The setting is the U.S. trade publishing industry and the blocked merger of Penguin Random House and Simon & Schuster, the first U.S. merger halted solely on the basis of harm to an input market. I assemble a dataset linking manuscript acquisitions to the retail sales of the resulting books, document that the 2013 merger of Penguin and Random House moved advances in patterns a pure-monopsony account cannot generate, and build a structural model that nests an ascending-bid auction for manuscripts inside a differentiated-products demand-and-pricing model for books,

Table 12: PRH–Simon & Schuster Merger: Welfare Effects

A. Firm profits			
	Variable profits		Total profits
	$\Delta\pi$ (\$M)	$\Delta\pi$ (%)	$\Delta\Pi$ (\$M)
PRH + SS	8.15	6.60%	7.25
HarperCollins	2.84	11.11%	−3.66
Hachette	1.41	9.86%	−1.09
Macmillan	1.21	9.00%	−2.79
Other publishers	3.43	10.41%	−2.97
All firms	17.03	8.13%	−3.27
B. Consumer Surplus Change			
	ΔCS (\$M)	ΔCS (%)	
Fiction	−10.16	−3.42%	
Nonfiction	−9.78	−2.01%	
Combined	−19.94	−2.54%	
C. Total Welfare Change			
	ΔW (\$M)	ΔW (%)	
Fiction	−0.92	−0.23%	
Nonfiction	−1.99	−0.33%	
Combined	−2.91	−0.29%	

Notes: Results from Nash-Bertrand merger counterfactual using RCNL preferred estimates. Consumer surplus is computed via the RCNL logsum formula and scaled to full-market, full-lifetime units using the same sales-weighted factor applied to profits (lifecycle coverage weight divided by Amazon market share). Total welfare $\Delta W = \Delta\text{CS} + \Delta\pi$ (unweighted sum). Includes titles sold between January and August 2021.

so that a counterfactual change in market structure propagates through both sides of the market simultaneously.

The central finding is that the two channels through which a merger reaches aggregate advances operate in opposite directions, and their incidence is sharply uneven. For the manuscripts the merging parties jointly contested, eliminating head-to-head competition dominates and advances fall. For the much larger set of manuscripts they did not jointly contest, the merger's downstream pricing power raises every publisher's per-book profit and lifts advances. In the baseline simulation, total aggregate advances rise by \$20.3 million (10.2%), where an upstream-only analysis of the same estimated model predicts a decline of \$8.7 million (−4.3%). A monopsony-only assessment does not merely mismeasure the merger's effect on suppliers, it can mis-sign it, and it conceals the redistribution from the most commercially successful authors toward the broader author population, that the aggregate contains. Meanwhile, consumer surplus falls by more than industry profit rises, so total welfare declines: the merger's harm, in the model, runs through the traditional consumer channel rather than the input-market channel on which the case was brought.

These results characterize the information set the 2023 Merger Guidelines now instruct regulators to assemble, rather than overturning the court's decision. The DOJ's central concern is confirmed: the authors over whom the merging parties competed would have been harmed, and

Table 13: Retail Price Changes from the PRH–Simon & Schuster Merger

	BL price	CF price	Δp ($)	Δp (%)
A. All markets (fiction + non-fiction)				
PRH + Simon & Schuster	16.65	16.96	0.31	2.33%
HarperCollins	16.70	16.73	0.03	0.31%
Hachette	17.80	17.81	0.01	0.14%
Macmillan	16.22	16.23	0.01	0.17%
Other publishers	18.24	18.25	0.01	0.11%
B. Fiction				
PRH + Simon & Schuster	15.24	15.50	0.26	2.26%
HarperCollins	14.19	14.22	0.03	0.45%
Hachette	15.41	15.42	0.01	0.13%
Macmillan	15.20	15.21	0.01	0.20%
Other publishers	16.22	16.23	0.00	0.05%
C. Non-fiction				
PRH + Simon & Schuster	18.26	18.62	0.36	2.41%
HarperCollins	18.65	18.69	0.03	0.20%
Hachette	19.38	19.40	0.02	0.14%
Macmillan	18.53	18.54	0.01	0.12%
Other publishers	19.04	19.05	0.01	0.13%

Notes: Revenue-weighted average retail prices across all product-market observations in 2021 (January–August, 16 category-months). BL and CF denote baseline and counterfactual prices. Counterfactual prices solve the post-merger Bertrand–Nash equilibrium with PRH and Simon & Schuster treated as a single firm; all other publishers reprice optimally in response to the merger. Estimates use the baseline RCNL specification.

a regulator weighing consumers alongside authors could well have reached the same conclusion the court did. What the joint analysis adds is the incidence of the merger's effects: concentrated losses for the authors the merging parties contested, diffuse gains across the rest of the manuscript market, and a consumer-surplus loss that an input-market analysis never measures. The remaining question is distributional: the merger harms one group of authors and benefits another, and weighing one group against the other is a judgment the analysis informs but does not make.

The framework has limitations. First, the counterfactual holds the acquisition protocol, the set of manuscripts brought to market, and authors' participation fixed. A merger large enough to change what gets written, rather than what gets paid for it, would require a model of author entry beyond this paper's scope. Second, the data record the winner and the price of each deal but not the identities or number of losing bidders, which the ascending format cannot identify from these observables alone. I therefore fix the number of bidders per auction and verify robustness: the sign of the aggregate effect and the incidence across author groups are unchanged under the alternative bidder configurations considered, though magnitudes vary.

The mechanism studied here is not specific to publishing. It operates wherever inputs are differentiated, acquired through bidding or negotiation, and tied to the downstream products they become: film scripts, patents and licensing rights, broadcast deals and athlete contracts. In any such market, buyer power and product-market power must be analyzed jointly, because a merger changes both the competition for an input and the value of acquiring it. As the blocked

PRH–SS merger illustrates, an analysis of either side alone can misstate who gains, who loses, and by how much.

References

Allen, J., Clark, R., and Houde, J.-F. (2019). Search frictions and market power in negotiated-price markets. *Journal of Political Economy*, 127(4):1550–1598.

Aradillas-López, A., Gandhi, A., and Quint, D. (2013). Identification and inference in ascending auctions with correlated private values. *Econometrica*, 81(2):489–534.

Arnold, D. (2025). Mergers and acquisitions, local labor market concentration, and worker outcomes. *Unpublished*.

Athey, S. and Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, 70(6):2107–2140.

Azar, J. A., Berry, S. T., and Marinescu, I. (2022). Estimating labor market power. Technical report, National Bureau of Economic Research.

Bar-Isaac, H., Johnson, J. P., and Nocke, V. (2024). Acquihiring for monopsony power. *Management Science*.

Berger, D. W., Hasenzagl, T., Herkenhoff, K. F., Mongey, S., and Posner, E. A. (2023). Merger guidelines for the labor market. Technical report, National Bureau of Economic Research.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262.

Binmore, K., Rubinstein, A., and Wolinsky, A. (1986). The nash bargaining solution in economic modelling. *The RAND Journal of Economics*, pages 176–188.

Bondi, T. and Cabral, L. (2021). Amazon and the future of retail. Technical report, Working Paper.

Card, D. (2022). Design-based research in empirical microeconomics. *American Economic Review*, 112(6):1773–1781.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.

Conlon, C. and Gortmaker, J. (2020). Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161.

Cuesta, J. I. and Sepúlveda, A. (2021). Price regulation in credit markets: A trade-off between consumer protection and credit access. *Available at SSRN 3282910*.

Deneckere, R. and Davidson, C. (1985). Incentives to form coalitions with bertrand competition. *The RAND Journal of economics*, pages 473–486.

DOJ and FTC (2023). "2023 Merger Guidelines", Technical Report, U.S. Department of Justice and the Federal Trade Commission.

Froeb, L. M., Mares, V., and Tschantz, S. (2017). Horizontal mergers in optimal auctions. *Vanderbilt Owen Graduate School of Management Research Paper*, (2742187).

Gandhi, A. and Houde, J.-F. (2019). Measuring substitution patterns in differentiated-products industries. *NBER Working paper*, (w26375).

Genakos, C., Pagliero, M., Sabatino, L., and Valletti, T. (2025). Cultural exception? the impact of price regulation on prices and variety in the market for books.

Ghose, A., Smith, M. D., and Telang, R. (2006). Internet exchanges for used books: An empirical analysis of product cannibalization and welfare impact. *Information systems research*, 17(1):3–19.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203.

Grigolon, L. and Verboven, F. (2014). Nested logit or random coefficients logit? a comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics*, 96(5):916–935.

Hemphill, C. S. and Rose, N. L. (2018). Mergers that harm sellers. *The Yale Law Journal*, pages 2078–2109.

Horn, H. and Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, pages 408–419.

Kawaguchi, K., Qiu, Y. J., and Zhang, Y. (2025). Competitive effects of resale price maintenance through inventory: Evidence from publishing industry. Technical report, EAG Discussion Paper.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *Review of Economic Studies*, 78(1):293–327.

Kroft, K., Luo, Y., Mogstad, M., and Setzler, B. (2025). Imperfect competition and rents in labor and product markets: The case of the construction industry. *American Economic Review*, 115(9):2926–2969.

Li, T., Perrigne, I., and Vuong, Q. (2000). Conditionally independent private information in OCS wildcat auctions. *Journal of Econometrics*, 98(1):129–161.

Li, T. and Zhang, B. (2015). Affiliation and entry in first-price auctions with heterogeneous bidders: An analysis of merger effects. *American Economic Journal: Microeconomics*, 7(2):188–214.

Loertscher, S. and Marx, L. M. (2019). Merger review for markets with buyer power. *Journal of Political Economy*, 127(6):2967–3017.

McFadden, D. (1972). Conditional logit analysis of qualitative choice behavior.

Micci-Barreca, D. (2001). A preprocessing scheme for high-cardinality categorical attributes in classification and prediction problems. *ACM SIGKDD explorations newsletter*, 3(1):27–32.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122.

Miller, N. H. and Weinberg, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791.

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The Rand journal of economics*, pages 395–421.

Peukert, C. and Reimers, I. (2022). Digitization, prediction, and market efficiency: Evidence from book publishing deals. *Management Science*, 68(9):6907–6924.

Prager, E. and Schmitt, M. (2021). Employer consolidation and wages: Evidence from hospitals. *American Economic Review*, 111(2):397–427.

Pruett, J. (2022). New York Times Hardcover Fiction Bestsellers (1931–2020). *Post45 Data Collective*.

Reimers, I. (2019). Copyright and generic entry in book publishing. *American Economic Journal: Microeconomics*, 11(3):257–284.

Reimers, I. and Waldfogel, J. (2017). Throwing the books at them: Amazon's puzzling long run pricing strategy. *Southern economic journal*, 83(4):869–885.

Reimers, I. and Waldfogel, J. (2021). Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings. *American Economic Review*, 111(6):1944–1971.

Rubens, M. (2023). Market structure, oligopsony power, and productivity. *American Economic Review*, 113(9):2382–2410.

Salop, S. C. and Scheffman, D. T. (1983). Raising rivals' costs. *The American economic review*, 73(2):267–271.

Setzler, B. (2025). Labor and product market power, endogenous quality, and the consolidation of the us hospital industry. Technical report, National Bureau of Economic Research.

Shi, K. (2024). Mergers and mismatches in the labor market for creativity.

Slattery, C. (2025). Bidding for firms: Subsidy competition in the united states. *Journal of Political Economy*, 133(8):2563–2614.

U.S. Department of Justice (2016). United states v. anthem, inc. and cigna corp. Case No. 1:16-cv-01493.

U.S. Department of Justice (2018). United states v. cdk global, inc. and auto/mate, inc.

U.S. Department of Justice (2021a). United states v. aon plc and willis towers watson plc. Case No. 1:21-cv-01633.

U.S. Department of Justice (2021b). United states v. unitedhealth group inc. and change healthcare inc. Case No. 1:21-cv-02553.

U.S. Department of Justice (2022a). Memorandum opinion: United states v. bertelsmann se co. kgaa, penguin random house, llc, and viacomcbs, inc. Case No. 1:21-cv-02886-FYP, filed October 31, 2022.

U.S. Department of Justice (2022b). United states v. cargill meat solutions corp., sanderson farms inc., and wayne farms llc. Case No. 1:22-cv-01540.

U.S. Department of Justice (2023). United states v. kroger co. and albertsons companies, inc.

Wan, M. and McAuley, J. (2018). Item recommendation on monotonic behavior chains. In *Proceedings of the 12th ACM conference on recommender systems*, pages 86–94.

Wan, M., Misra, R., Nakashole, N., and McAuley, J. (2019). Fine-grained spoiler detection from large-scale review corpora. *arXiv preprint arXiv:1905.13416*.

Wolpert, D. H. (1992). Stacked generalization. *Neural networks*, 5(2):241–259.

Yeh, C., Macaluso, C., and Hershbein, B. (2022). Monopsony in the us labor market. *American Economic Review*, 112(7):2099–2138.

Appendix

A Robustness of Merger Effects to the Acquisition Protocol

The baseline model represents the acquisition of a title as an ascending-bid auction with independent private values, so that the advance equals the second-highest valuation among interested publishers. In practice, a substantial share of acquisitions is negotiated bilaterally (e.g., exclusive submissions, option books, and pre-emptive agreements) rather than run as formal auctions. This appendix shows that the mechanism behind the paper's comparative statics does not depend on the protocol, and characterizes when the sign of the merger's effect does.

I embed the auction in a family of bilateral bargaining protocols indexed by the author's bargaining weight and establish three results. First, the auction is the member of the family in which the author has no bargaining power, so the advance under bargaining is weakly higher than under the auction. Second, on titles where the merging parties were not the two highest-valuation bidders, the author weakly gains under every bargaining weight. Third, on titles where the auction predicts a loss, the change in the advance under bargaining is weakly larger than under the auction for every bargaining weight, and it becomes a gain once the author's weight exceeds a title-specific threshold. Holding valuations fixed, the auction is therefore the member of the family that is least favorable to authors in the merger.

A.1 Environment

Fix a title and suppress its index. A finite set of publishers $\mathcal{F} = \{1, \dots, N\}$, $N \geq 2$, is interested in the title. Publisher f's valuation $v_f \geq 0$ is its willingness to pay: in the notation of Section 5, the certainty equivalent of its incremental profit from the title net of its fixed cost, including the private cost shock.[43] Let $v^{(1)} \geq v^{(2)} \geq \cdots \geq v^{(N)}$ denote the order statistics of $(v_f)_{f \in \mathcal{F}}$, and let w denote the highest-valuation publisher. The author's payoff from a deal is the advance A; her payoff from no deal with any publisher is normalized to zero.

Two firms $\mathcal{M} = \{m, g\} \subset \mathcal{F}$ merge. Post-merger valuations are $(\tilde{v}_f)$, and the merged entity participates as a single buyer with valuation $\tilde{v}_{\mathcal{M}} \equiv \max\{\tilde{v}_m, \tilde{v}_g\}$, as in the baseline model. Post-merger order statistics, $\tilde{v}^{(k)}$, are taken over the post-merger buyer set (the merged entity and the non-merging publishers).

Assumption 1 (Downstream channel)**.** *$\tilde{v}_f \geq v_f$ for every $f \in \mathcal{F}$.*

Assumption 1 restates, at the level of a single title, the downstream mechanism of the paper: the merger raises equilibrium retail prices, which weakly raises every publisher's valuation of the title.

Lemma 1 (Order statistics)**.** *Under Assumption 1, $\tilde{v}^{(1)} \geq v^{(1)}$. If the two highest-valuation publishers are not both merging parties, then $\tilde{v}^{(2)} \geq v^{(2)}$. If they are, then $\tilde{v}^{(2)} \geq v^{(3)}$; specifically,*

1. *if the merged entity remains the highest-valuation buyer, $\tilde{v}_{\mathcal{M}} \geq \max_{f \notin \mathcal{M}} \tilde{v}_f$, then $\tilde{v}^{(2)} = \max_{f \notin \mathcal{M}} \tilde{v}_f \geq v^{(3)}$;*

[43] Nothing below uses the internal structure of v_f. The results apply to any valuation vector.

2. otherwise a non-merging publisher wins and $\tilde{v}^{(2)} \geq \tilde{v}_{\mathcal{M}} \geq v^{(1)}$.

Proof. Let a and b be the two highest-valuation publishers before the merger. If $\{a, b\} \neq \mathcal{M}$, they correspond to two distinct buyers after the merger (one of them possibly the merged entity), with valuations at least v_a and v_b, so the top two order statistics weakly rise. If $\{a, b\} = \mathcal{M}$, then $\tilde{v}_{\mathcal{M}} = \max\{\tilde{v}_m, \tilde{v}_g\} \geq \max\{v_m, v_g\} = v^{(1)}$ by Assumption 1, and the best non-merging publisher has valuation $\max_{f \notin \mathcal{M}} \tilde{v}_f \geq \max_{f \notin \mathcal{M}} v_f = v^{(3)}$. In case (i), the merged entity wins and the runner-up is the best non-merging publisher. In case (ii), a non-merging publisher wins and the merged entity's valuation is a lower bound on the second-highest, so $\tilde{v}^{(2)} \geq \tilde{v}_{\mathcal{M}} \geq v^{(1)} \geq v^{(3)}$. □

Partition. A title is *jointly contested*, $j \in \mathcal{J}_B$, if the two highest-valuation publishers are the merging parties; otherwise $j \in \mathcal{J}_A$.

This is the partition that organizes the counterfactual results in Section 8: $\mathcal{J}_B$ collects the titles for which the merger eliminates the head-to-head contest between PRH and S&S.

A.2 Protocols

I compare two protocols for the sale of the title.

(P1) Ascending auction (baseline). All interested publishers bid. Under independent private values the title is sold to w at the second-highest valuation,

$$A^{\text{AUC}} = v^{(2)}. \tag{32}$$

(P2) Bilateral bargaining with a competitive outside option. The author negotiates with a single publisher under generalized Nash bargaining with author weight $\alpha \in [0, 1]$. If the negotiation breaks down, the title goes to the best alternative publisher, who pays up to its valuation. The excluded publisher earns nothing from the title, so its disagreement payoff is zero.[44] The author negotiates with the highest-valuation publisher w: with any other publisher f, her best alternative is w itself, worth $v^{(1)} \geq v_f$, so there are no gains from a deal. With w, her disagreement payoff is $d = v^{(2)}$.

Lemma 2 (The auction is a floor)**.** *Under (P2), the title is sold to* w *and the advance is*

$$A^{\text{NB}}(\alpha) = v^{(2)} + \alpha \left(v^{(1)} - v^{(2)} \right) \;\geq\; v^{(2)} = A^{\text{AUC}}, \tag{33}$$

with equality if and only if $\alpha = 0$ *or* $v^{(1)} = v^{(2)}$*. The ascending auction is therefore the zero-bargaining-power member of the family (P2): an author with no bargaining power receives exactly her disagreement payoff, the second-highest valuation. Because the title goes to the highest-valuation publisher under both protocols, the allocation, and therefore total welfare, in which advances net out as transfers, is the same under (P1) and (P2).*

[44] This is the relevant threat in this market: a pre-empt gives one house the chance to buy a title before it is offered more widely, and the agent's alternative to accepting it is to take the submission to other publishers (U.S. Department of Justice, 2022a). That the author obtains the alternative publisher's full valuation is the most favorable case for her; Corollary 1 considers a weaker outside option. The zero disagreement payoff is consistent with the valuations of Section 5, which already measure a publisher's profit relative to a rival publishing the title.

Proof. The surplus from a deal with w is $v^{(1)} - d = v^{(1)} - v^{(2)} \geq 0$, of which the author receives the share α on top of her disagreement payoff $d = v^{(2)}$. The inequality follows from $\alpha \geq 0$ and $v^{(1)} \geq v^{(2)}$. □

The advance under bargaining thus decomposes into a competitive floor, $v^{(2)}$, delivered by the presence of a rival, and a share $\alpha(v^{(1)} - v^{(2)})$ of the winner's rent. The auction pays the floor and none of the rent.

A.3 Main result

For any object x, write $\Delta x \equiv \tilde{x} - x$ for its change across the merger. Before and after the merger, the author's disagreement payoff is the second-highest valuation, $d = v^{(2)}$ and $\tilde{d} = \tilde{v}^{(2)}$, so $\Delta d = \Delta v^{(2)}$, which is also the change in the auction advance.

Proposition 1 (Merger effects under bargaining)**.** *Let Assumption 1 hold and fix $\alpha \in [0, 1]$.*

(i) Decomposition. *Under (P2),*

$$\Delta A^{\mathrm{NB}} = \underbrace{\alpha\, \Delta v^{(1)}}_{\textit{downstream channel}} + \underbrace{(1-\alpha)\, \Delta v^{(2)}}_{\textit{competition channel}}, \qquad \Delta A^{\mathrm{AUC}} = \Delta v^{(2)}, \tag{34}$$

so the auction is the $\alpha = 0$ case.

(ii) Titles on which the auction advance does not fall *($\Delta v^{(2)} \geq 0$). These include every title in $\mathcal{J}_A$. On them, $\Delta A^{\mathrm{NB}} \geq 0$ for every α.*

(iii) Titles on which the auction advance falls *($\Delta v^{(2)} < 0$). These are jointly contested titles on which the merged entity remains the highest-valuation buyer, and*

$$\Delta v^{(2)} = \underbrace{\left(\max_{f \notin \mathcal{M}} \tilde{v}_f - v^{(3)}\right)}_{\textit{spillover} \geq 0} - \underbrace{\left(v^{(2)} - v^{(3)}\right)}_{\textit{runner-up gap} \geq 0}. \tag{35}$$

On them, the advance under bargaining falls if and only if

$$\alpha < \alpha^* \equiv \frac{-\Delta v^{(2)}}{\Delta v^{(1)} - \Delta v^{(2)}} \in (0, 1]. \tag{36}$$

(iv) Comparison with the auction. *For every title,*

$$\Delta A^{\mathrm{NB}} - \Delta A^{\mathrm{AUC}} = \alpha\big(\Delta v^{(1)} - \Delta v^{(2)}\big). \tag{37}$$

On the titles in (iii), this is nonnegative, and strictly positive if $\alpha > 0$. On the titles in (ii), it is nonnegative if and only if $\Delta v^{(1)} \geq \Delta v^{(2)}$.

Proof. (i) Apply (33) before and after the merger. In both cases the author negotiates with the highest-valuation buyer, and her disagreement payoff is the second-highest valuation, so $A^{\mathrm{NB}} = v^{(2)} + \alpha(v^{(1)} - v^{(2)})$ and $\tilde{A}^{\mathrm{NB}} = \tilde{v}^{(2)} + \alpha(\tilde{v}^{(1)} - \tilde{v}^{(2)})$. Subtracting gives (34); setting $\alpha = 0$ gives the auction.

(ii) By Lemma 1, $\Delta v^{(1)} \geq 0$ for every title and $\Delta v^{(2)} \geq 0$ on $\mathcal{J}_A$. If $\Delta v^{(2)} \geq 0$, both terms of (34) are nonnegative.

(iii) By (ii), $\Delta v^{(2)} < 0$ requires $j \in \mathcal{J}_B$. On $\mathcal{J}_B$, if a non-merging publisher overtakes the merged entity, Lemma 1 gives $\tilde{v}^{(2)} \geq v^{(1)} \geq v^{(2)}$, so $\Delta v^{(2)} \geq 0$. Hence the merged entity remains the highest-valuation buyer, and Lemma 1 gives $\tilde{v}^{(2)} = \max_{f \notin \mathcal{M}} \tilde{v}_f$. Subtracting $v^{(2)}$, then adding and subtracting $v^{(3)}$, gives (35). The spillover is nonnegative by Assumption 1, since $v^{(3)} = \max_{f \notin \mathcal{M}} v_f$ on $\mathcal{J}_B$, and the gap is nonnegative by the ordering of valuations. For the threshold, write $\Delta A^{\mathrm{NB}} = \Delta v^{(2)} + \alpha(\Delta v^{(1)} - \Delta v^{(2)})$. Because $\Delta v^{(1)} \geq 0 > \Delta v^{(2)}$, this is strictly increasing in α. It equals $\Delta v^{(2)} < 0$ at $\alpha = 0$ and $\Delta v^{(1)} \geq 0$ at $\alpha = 1$, so it is negative exactly when $\alpha < \alpha^*$.

(iv) Subtract $\Delta A^{\mathrm{AUC}} = \Delta v^{(2)}$ from (34). On the titles in (iii), $\Delta v^{(1)} \geq 0 > \Delta v^{(2)}$, so the bracket is positive. On the titles in (ii), the sign is that of $\Delta v^{(1)} - \Delta v^{(2)}$. □

For example, let pre-merger valuations be $(100, 90, 80)$, with the merging parties first and second, and suppose the merger raises their valuations to $(120, 110)$ and the outsider's to 85. The auction advance falls from 90 to 85. Under bargaining with $\alpha = 1/2$, it rises from 95 to 102.5. Here $\Delta v^{(1)} = 20$, $\Delta v^{(2)} = -5$, and $\alpha^* = 5/25 = 0.2$.

The merger enters every protocol in the family through the same two objects: the change in the winner's valuation, which carries the downstream channel, and the change in the runner-up's valuation, which carries the competition channel. The bargaining weight sets how much each one counts. Where the auction advance does not fall, both changes are nonnegative, and the author gains under every protocol. Where it falls, the two changes have opposite signs, and the author loses only if her weight is below α^*. Whether the merger harms authors on these titles therefore depends on bargaining power as well as on valuations. What does not depend on it is the ranking in part (iv): at given valuations, the auction is the protocol under which the merger is least favorable to authors on the titles where it harms them.

A.4 An alternative treatment of the outside option

Corollary 1 (Outside option principle)**.** *Suppose the outside option enters as a constraint rather than a threat point (Binmore et al., 1986), so that* $A^{\mathrm{OOP}} = \max\{\alpha\, v^{(1)},\ v^{(2)}\}$. *Then (a)* $A^{\mathrm{OOP}} \geq A^{\mathrm{AUC}}$*; (b) on* $j \in \mathcal{J}_A$, $\Delta A^{\mathrm{OOP}} \geq 0$*; and (c) on titles with* $\Delta v^{(2)} < 0$, $\Delta A^{\mathrm{OOP}} \geq \Delta A^{\mathrm{AUC}}$.

Proof. (a) is immediate. (b) By Lemma 1 both arguments of the maximum weakly rise on $\mathcal{J}_A$. If $A^{\mathrm{OOP}} = v^{(2)}$ before the merger, then $\Delta A^{\mathrm{OOP}} \geq \tilde{v}^{(2)} - v^{(2)} = \Delta v^{(2)}$. If $A^{\mathrm{OOP}} = \alpha v^{(1)}$, then $\Delta A^{\mathrm{OOP}} \geq \alpha \Delta v^{(1)} \geq 0 > \Delta v^{(2)}$. □

A.5 Discussion

Remark 1 (Sequential bargaining without an auction)**.** *If the author could not hold an auction and instead bargained with publishers one at a time, her advance would be a weighted sum of all valuations and could fall below* $v^{(2)}$ *when her bargaining weight is low. The two channels keep their signs in that case, but the auction is no longer a floor. I do not pursue this protocol because agents in this market can and do run auctions, so an author would not accept less than the auction price.*

Remark 2 (Levels versus differences, at given valuations)**.** *The two bounds above hold for a given valuation vector. In levels, the auction gives the author the lowest advance in (P2) (Lemma 2). In merger differences, it gives the smallest change on titles where the auction advance falls (Proposition 1(iv)). These are not bounds on the estimates in Section 8. Those estimates are obtained by treating observed advances as second-highest valuations, so if some advances were negotiated, the estimated valuations themselves differ (Remark 3), and so does the counterfactual built on them. This appendix shows that the mechanism does not depend on the protocol and that the auction is the least favorable protocol for authors at given valuations. It does not sign the bias of the estimated counterfactual.*

Remark 3 (Implication for estimation)**.** *Since* $A^{\mathrm{NB}}(\alpha) \geq v^{(2)}$*, with equality at* $\alpha = 0$*, an advance negotiated under (P2) is an upper bound on the runner-up's valuation. Interpreting observed advances as second-highest valuations, as the estimation does, therefore overstates valuations for deals that were negotiated rather than auctioned. This appendix does not sign the effect of that bias on the counterfactual.*

Remark 4 (Levels versus differences)**.** *Two statements above concern different objects. In levels, the auction gives the author the lowest advance within (P2) (Lemma 2). In merger differences, the auction gives the largest loss on jointly contested titles (Proposition 1(iv)). Together they imply that the losses on jointly contested titles reported in Section 8 are upper bounds within (P2).*

B Combining Deals With Sales Data

This appendix describes the procedure used to match data on deals from Publishers Marketplace with sales data from Bookstat. The main challenge is due to the fact that when a deal is registered there is no ISBN code yet. The main information we can use for this relates to the book's title, though at the time of the deal it may differ from the title upon publication.

To overcome these challenges, I use the website isbndb.com to retrieve ISBNs for books reported in the deals. By leveraging the API provided by the website, I search for the deals reported using the book titles, authors' names, and year of publication to find a match. The search pattern is structured as follows:

1. First, search by title and author name. If a match is found, retain the ISBN.
2. If no match is found, search by title only (i.e., the author's name may be incorrect or misspelled). If a match is found and the author's name is similar enough to the one in the deal, retain the ISBN.
3. If no match is found, repeat step 2 but search for books whose title contains the title reported in the deal (e.g., instead of finding "Misery," find "Misery: A Novel"). If a match is found and the author's name is similar enough, retain the ISBN.
4. If no match is found, search by author name and check whether the author published a book within two years of the deal. If a match is found, retain the ISBN.

After completing this process, I manually review any unmatched titles to account for cases where the book was published more than two years after the deal or where deal information (e.g., names, years) was incorrect.

C Representativeness of the Deals Data

Reporting to Publishers Marketplace is voluntary, which raises the concern that the deals data are a selected sample of the market's transactions. Selection could operate along two dimensions: who reports (if some publishers' deals are systematically over- or under-represented, the estimated competitive structure of the upstream market would be distorted), and what is reported (if disclosure depends on the size of the advance). This appendix addresses the first dimension by benchmarking the publisher composition of PM against an external universe of published titles. The second cannot be checked directly, since advances are not observed outside PM, but Section 3.2 provides indirect evidence using sales as a proxy when comparing matched and unmatched deals.

The benchmark universe comes from Bookstat, which records titles sold on Amazon independently of any publisher's decision to announce deals. To make the two sources comparable, I restrict the Bookstat catalog to titles whose characteristics match the scope of the deals data: traditionally published trade titles in fiction and non-fiction, first published during 2018-2021. Because titles are published 12 to 24 months after the rights sale, PM deals signed in 2016-2020 correspond to Bookstat titles appearing in 2018-2021. The comparison aligns the two accordingly. The unit differs mechanically across sources so the comparison should be read as composition shares within each source rather than as a count reconciliation.

Table C1 reports the share of each source accounted for by each of the Big Five publishers and by all remaining publishers. The full PM sample contains 78,949 deal announcements after dropping deals with missing publisher or category information, the analysis sample of Section 3.1 contains 77,649. The Big Five account for 41% of PM deals and 39% of the comparable Bookstat titles, and the order of firms is consistent, with PRH leading and followed by Harper, and Simon & Schuster being among the ones with the lowest shares.

If reporting to PM were driven by publisher identity (e.g., if major houses publicized systematically more, or small presses systematically less, than their true acquisition activity) the two compositions would diverge. That they match closely indicates that PM's coverage is not selected on publisher identity, and that the relative acquisition activity of the firms in the model reflects the market rather than reporting behavior. Two caveats bound the claim. First, the Bookstat universe itself conditions on a minimum review count (Section 3.2), so both sources under-represent the smallest publishers relative to the full population of published titles. The comparison validates PM against the commercially active segment that the analysis concerns, not against every title published. Second, composition parity does not rule out selection on unobserved deal characteristics within publisher, in particular the advance itself, which is why the disclosure-selection evidence of Section 3.2 and the reweighting of the reduced-form estimates remain necessary.

Table C1: Big Five Publisher Shares: Deals vs. Sales Data

	(1) **PM Deals (all years)** (N = 78,949)	(2) **PM Deals (2016–2020)** (N = 26,618)	(3) **Bookstat (2018–2021)** (N = 128,723)
Penguin Random House	16.3%	15.6%	13.1%
HarperCollins	7.4%	8.0%	8.7%
Simon & Schuster	5.5%	5.2%	4.0%
Hachette	4.9%	5.1%	7.9%
Macmillan	7.0%	7.2%	5.3%
Big Five (total)	41.0%	41.0%	39.0%
Other / Independent	59.0%	59.0%	61.0%

Notes. *PM Deals*: all book deals recorded in Publishers Marketplace (N = 78,949). The second column restricts to deals signed in 2016–2020, the years most likely to correspond to titles observed in the 2018–2021 sales panel. *Bookstat*: unique ISBNs from the 2018–2021 Amazon sales panel, restricted to traditional trade publishers. Filters applied: publisher total units ≥ 100; exclude non-book/non-trade Amazon categories (audiobooks, textbooks, foreign languages, crafts, engineering, medical, law, etc.); exclude audio formats; exclude hybrid/subsidy publishers (e.g. Morgan James, She Writes Press) and self-published titles. Additionally, I exclude ISBNs with missing US sales over the full panel period.

C.1 Matched and Unmatched Deals

The second dimension of selection, disclosure related to deal size, cannot be tested against advances directly, since advances are unobserved outside PM. Table C2 provides the indirect evidence referenced in Section 3.2, using commercial visibility as the proxy: it compares the characteristics of PM deals that match to at least one title in the Bookstat sales data against those that match to none. Matched deals are larger and more concentrated among major publishers: in fiction, they carry advances 17% higher on average, a 1.9 percentage-point higher best-seller share, and a 5.6 percentage-point higher Big Five share, with the same patterns in non-fiction. The differences confirm that matching selects toward the commercially visible segment of the market, which is the segment the structural analysis concerns. They also confirm that disclosure and commercial success are related, which is why the reduced-form estimates of Section 4 are reweighted for selection into disclosure rather than treating disclosed deals as representative.

Table C2: Mean Values for Matched and Unmatched Deals.

	Fiction		Non-fiction	
	Unmatched	Matched	Unmatched	Matched
Advance ($)	165,480	193,140	170,051	188,784
Debut	0.233	0.192	0.346	0.260
Best-seller	0.093	0.112	0.044	0.053
Big Five	0.413	0.469	0.423	0.476
PRH	0.141	0.196	0.158	0.203
Simon & Schuster	0.050	0.042	0.052	0.046
N	8,355	4,882	17,012	3,046
Match rate	36.9%		15.2%	

Notes: Deals data from Publishers Marketplace. Sales data from Bookstat. Children's titles excluded. Matched deals are those linked to at least one ISBN-13 in the Bookstat sales data. Observations are at the deal level.

D Imputing Undisclosed Advance Amounts

Publishers Marketplace reports payment information only as one of five coarse categories, and only for roughly 23% of deals (see Section 3.1). The structural estimation requires a dollar-valued advance for every deal in the matched sample, which poses two tasks: predicting the category for deals that do not disclose one, and converting categories into dollar values. This appendix describes the prediction models, their validation, and the conversion. The imputed advances are used only in the structural estimation. The reduced-form analysis of Section 4 uses disclosed deals exclusively, reweighted for selection into disclosure.

Training Sample

The models are trained on the 11,154 post-2014 deals that disclose a category and used to impute the 39,399 post-2014 deals that do not. The restriction to post-2014 data is deliberate. Section 4 documents that the 2013 merger of Penguin and Random House shifted both the level and the segment composition of advances. A model trained on data straddling the merger would pool two different advance-generating regimes, and its predictions would blend them. After 2014, no consolidation among the Big Five publishers occurred, so the mapping from deal characteristics to advances is plausibly stable over the training window. The window also contains the 2016–2021 deal window of the structural sample, so no imputation extrapolates outside the training support.

Estimation Objects

Let $c_j \in \{1, \ldots, 5\}$ denote the advance category of deal j, ordered from "nice" to "major," m_c the midpoint of category c's dollar range (see Table 1), and $\mathbf{x}_j$ the feature vector described in Table D1. I estimate two formulations. The multiclass formulation recovers the conditional category probabilities

$$\hat{p}_{jc} \;=\; \widehat{\Pr}\left(c_j = c \mid \mathbf{x}_j\right), \qquad c = 1, \ldots, 5, \tag{38}$$

by minimizing the multinomial log loss $-\sum_j \sum_c \mathbf{1}[c_j = c] \log \hat{p}_{jc}$ over gradient-boosted tree ensembles. The regression formulation recovers a conditional-mean function

$$\widehat{\log a}_j \;=\; \hat{f}\left(\mathbf{x}_j\right), \tag{39}$$

trained on $\log m_{c_j}$ with squared-error loss. Gradient-boosted trees fit a sequence of shallow decision trees, each to the residuals of the ensemble so far, and are the standard approach for prediction on tabular data with many categorical features. I use two implementations, LightGBM (Ke et al., 2017) and XGBoost (Chen and Guestrin, 2016).

Features

Each formulation is estimated on two feature sets. The base set contains the deal's observable characteristics as recorded in PM and the auxiliary data of Section 3.3, together with a core set of hand-built interactions. The enriched set adds five further raw characteristics, target encodings of the high-cardinality categorical variables (out-of-fold means of the outcome within each category level (Micci-Barreca, 2001)) and additional interaction and polynomial terms. Table D1 lists every variable.

Table D1: Features Used in the Advance-Imputation Models

Variable	Definition	Base	Enriched
Panel A: Deal and title characteristics			
Publisher	Acquiring imprint / cleaned publisher name	✓	✓
Literary agency	Agency representing the author	✓	✓
Genre	Fiction / non-fiction / children's	✓	✓
Number of books	Books included in the contract	✓	✓
Competitive situation	Auction / pre-empt / exclusive submission	✓	✓
Signing year	Calendar year of the announcement	✓	✓
Co-agent	Deal involved a co-agent	✓	✓
Repeat imprint	Author previously published with this imprint	✓	✓
Panel B: Author characteristics			
Backlist size	Log number of the author's prior titles	✓	✓
Best-seller status	Prior appearance on any best-seller list	✓	✓
Debut indicator	First book by the author	✓	✓
Panel C: Core interactions (both sets)			
Interactions	Publisher × best-seller, agency × best-seller, competition × best-seller, competition × debut, year × publisher, genre × books, year × books	✓	✓
Panel D: Enriched-only additions			
Editor	Acquiring editor		✓
Rights territory	World / World English / US–Canada / other		✓
Publication season	Season of scheduled publication		✓
Best-seller list	NYT / other list / none		✓
Pub-year gap	Scheduled publication year − deal year		✓
Target encodings	Out-of-fold mean outcome by publisher, editor, agency, rights territory, genre, competition, and best-seller list		✓
Polynomial terms	Year^2, log-backlist^2, year × log-backlist, and further cross-products		✓
Additional interactions	Territory × publisher, territory × genre, editor × best-seller, season × publisher, best-seller list × publisher, pub-gap × genre		✓

Notes: Target encodings (TE) are computed out of fold: the encoding for observations in fold k uses only observations outside fold k, preventing leakage of the outcome into the features.

Models, Tuning, and Stacking

The two formulations and two feature sets yield six models, labeled M1–M6 in Table D2: the multiclass formulation is estimated with LightGBM, and the regression formulation with both LightGBM and XGBoost, each on the base and enriched feature sets. A seventh model, M7, is a stacking ensemble (Wolpert, 1992): a Ridge regression meta-learner trained on the out-of-fold dollar predictions of the three best-performing base models (M4–M6). The boosting models share a common set of regularization hyperparameters (shrinkage 0.02, 31 leaves, row and column subsampling of 0.8). The leading single model is additionally refined over a small grid of leaf-count and learning-rate values chosen by five-fold cross-validation. Selected values are reported in the replication package.

Evaluation and Model Selection

Performance is evaluated on out-of-fold (OOF) predictions from five-fold cross-validation: each deal's prediction comes from a model estimated without that deal, so all reported fit is out of sample by construction. To place the two formulations on a common footing, the classifiers'

predictions are first converted to dollars through equation (40) below, and the reported metric is the R^2 of OOF predicted log advances against disclosed log midpoints. Table D2 reports the comparison.

Table D2: Model Comparison: Out-of-Fold Performance

	Implementation	Formulation	Features	CV R^2	SD across folds
M1	LightGBM	Multiclass	Base	0.688	0.008
M2	LightGBM	Regression	Base	0.677	0.007
M3	XGBoost	Regression	Base	0.589	0.010
M4	LightGBM	Multiclass	Enriched	0.709	0.009
M5	LightGBM	Regression	Enriched	0.716	0.003
M6	XGBoost	Regression	Enriched	0.709	0.006
M7	Stack of M4–M6	—	—	0.721	0.005

Notes: Five-fold cross-validated R^2 of out-of-fold predicted log advances against disclosed log category midpoints, $N = 11{,}154$ post-2014 disclosed deals. Classifier predictions are converted to dollars via equation (40) before evaluation. SD across folds is the standard deviation of the fold-level R^2. M7 stacks the three best base models (M4–M6) via a Ridge meta-learner.

The best single model is the LightGBM regression classifier on the enriched features (M5). The stacking ensemble M7 improves on it modestly and produces the imputations used in the paper.

Conversion to Dollars

The models operate in logs. The classifiers return category probabilities, which I summarize by the probability-weighted average of log midpoints,

$$\widehat{\log a_j} \;=\; \sum_{c=1}^{5} \hat{p}_{jc} \log m_c, \tag{40}$$

while the regression models and the stacking ensemble M7 predict $\widehat{\log a_j}$ directly. The imputed log advance is M7's fitted value for undisclosed deals and the log midpoint $\log m_{c_j}$ for disclosed deals, with "major" deals top-coded at \$750,000 (Section 3.1). The structural model uses the dollar advance $a_j = \exp(\widehat{\log a_j})$, which for disclosed deals returns the category midpoint exactly. Because the models are estimated in logs and the advance enters as $\exp(\widehat{\log a_j})$, the imputation recovers a conditional geometric mean of the category midpoints rather than their arithmetic mean, with no retransformation adjustment. The two coincide for disclosed deals and differ only for undisclosed deals with an uncertain predicted category, where the geometric mean is modestly lower.

Validity

The imputation's validity rests on a conditional-ignorability assumption: conditional on the features in Table D1, the category distribution of undisclosed deals matches that of disclosed ones. The assumption is not directly testable, but the available diagnostics bear on it. The

density of imputed advances for undisclosed deals reproduces the overall shape of the disclosed distribution, with a mass concentrated at the lower categories, thinning toward the top (see Figure D1, panel E) and mean imputed advances align with mean disclosed advances within each rights territory (panel J). The model is also well calibrated across the deciles of predicted advances (panel F). The imputed advances are, moreover, economically sensible along dimensions not used to validate them: they are higher for the Big Five than for independent publishers (panel G), stable over the sample period (panel H), and higher for fiction and non-fiction than for children's books (panel I).

Fit Across the Advance Distribution

Table D3 reports fit by category, and Figure D1 the corresponding calibration and distributional diagnostics.

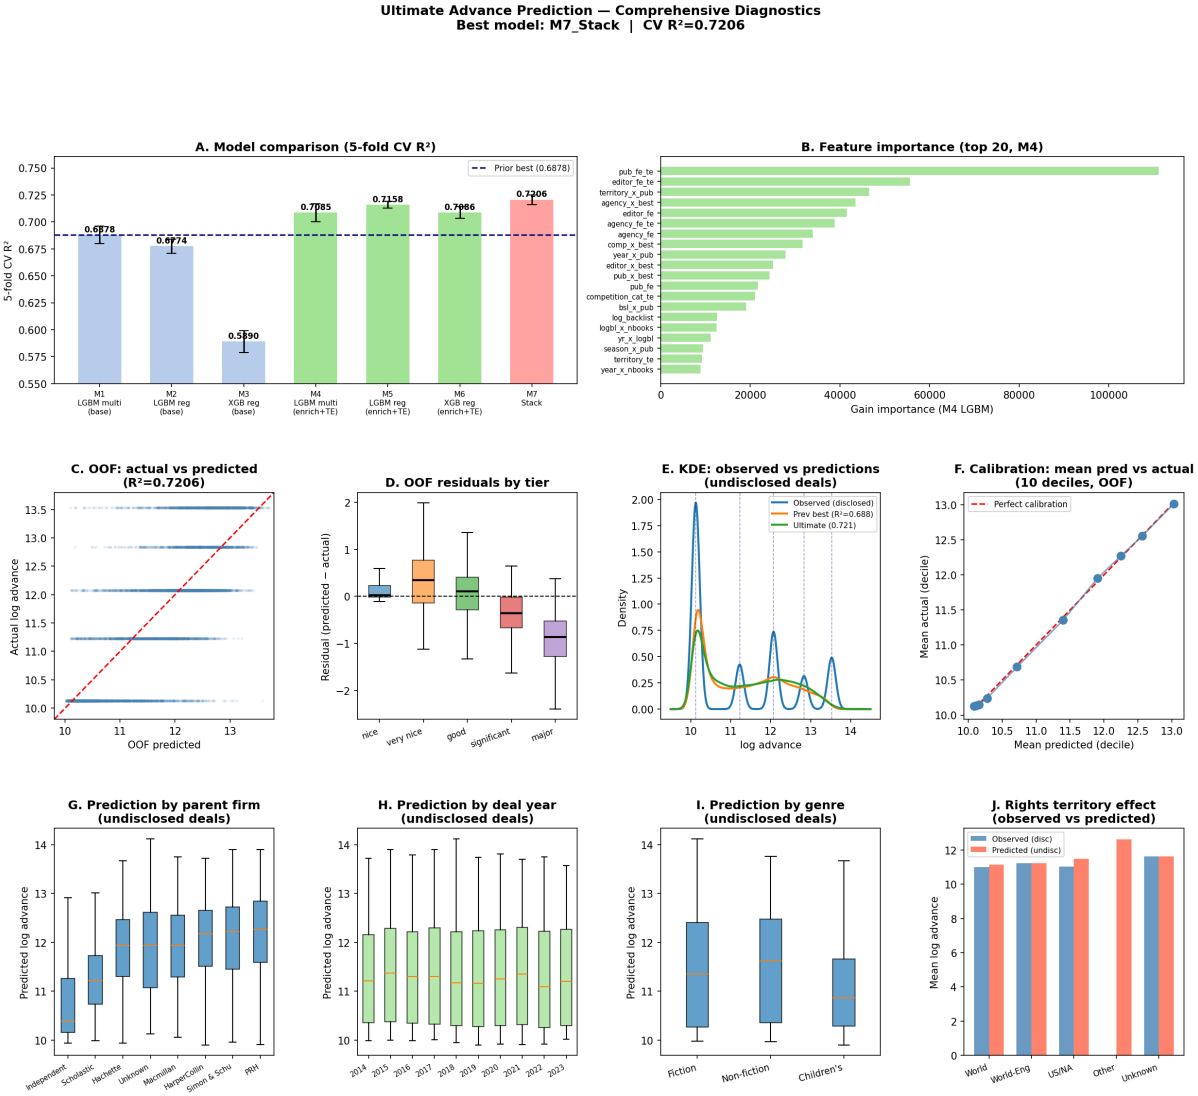
Table D3: Out-of-Fold Fit by Advance Category

Category	Range ($)	Midpoint ($)	N (training)	Hit rate	Mean residual (log points)
Nice	[1–49,999]	25,000	5,580	0.936	+0.254
Very nice	[50,000–99,999]	75,000	1,204	0.257	+0.275
Good	[100,000–249,999]	175,000	2,090	0.600	−0.087
Significant	[250,000–499,999]	375,000	890	0.103	−0.536
Major	[500,000+]	750,000	1,390	0.506	−1.091

Notes: Out-of-fold diagnostics for the LightGBM multiclass classifier (M4) on the training sample, $N = 11{,}154$ post-2014 disclosed deals. Hit rate is the share of disclosed deals whose highest-probability predicted category equals the disclosed category (overall hit rate = 0.680). Mean residual is predicted minus actual log advance, using the probability-weighted midpoint prediction $\sum_c \hat{p}_{jc}\, m_c$.

The model is well calibrated across the bulk of the distribution (Figure D1, panel F), but residuals show a downward bias for significant and major deals, and the middle categories are frequently classified into an adjacent tier (hit rates of 0.26 for very nice and 0.10 for significant). This is a known limitation: these tiers are rare in the training data, and the most informative deal-specific signals (e.g., the number of competing publishers, the author's royalty history, the acquirer's strategic priorities) are unobserved.

Figure D1: Advance Imputation: Calibration and Distributional Diagnostics



Notes: Panel A compares five-fold cross-validated R^2 across the seven models. Panel B reports feature importance (gain) for the top 20 features of the multi-class classifier (M4). Panel C plots out-of-fold predicted (stacking ensemble, M7) against actual log advances on the training sample, and panel D the corresponding out-of-fold residuals by category. Panel E compares the density of disclosed log advances with the density of imputed log advances for undisclosed deals. Panel F reports calibration: mean actual against mean predicted log advance by decile of the prediction. Panels G–I show the distribution of imputed advances for undisclosed deals by parent firm, deal year, and genre. Panel J compares mean imputed (undisclosed) and mean disclosed advances by rights territory. Training sample: 11,154 disclosed post-2014 deals. Prediction sample: 39,399 undisclosed post-2014 deals.

E Alternative HHI

In Section 4 I use differences in concentration between the best-seller and non-best-seller segments to look for heterogeneous changes after the PRH merger in 2013. There, market shares are based on each publisher's count of deals in each segment and year, and the best-seller segment is restricted to publishers with at least 50 deals with best-seller authors over the sample period. Here I report changes in concentration under alternative thresholds for defining the publishers in the best-seller segment, as well as an HHI based on advances paid.

Tables E2 and E1 lower the best-seller threshold to 30 and 20 deals, respectively, while Table E3 weights deals by the advance paid rather than counting them, using the midpoint of each Publishers Marketplace advance tier. The pattern from the main text holds in all cases: concentration in the best-seller segment is substantially higher than in the non-best-seller segment, and its ΔHHI is above 450 in every specification. Although the dollar-weighted measure yields a ΔHHI of about 240 in the non-best-seller segment, both its pre- and post-merger levels (933 and 1,173) remain well below the 1,800 threshold for a highly concentrated market in the 2023 Merger Guidelines, whereas the best-seller segment is above that threshold after the merger in the baseline specifications.

Table E1: HHI and ΔHHI After the PRH Merger (Best-Seller Threshold: 20 Deals)

Market	HHI pre-merger	HHI post-merger	ΔHHI
Best Seller	1419.5	1894.0	474.5
Non Best Seller	358.5	447.5	88.9

Notes: Herfindahl–Hirschman indices computed from publisher shares of deals signed in each segment and year, using Publishers Marketplace data. The unit is the deal; multi-book contracts count once. Pre- and post-merger values are averages of yearly indices over 2008–2013 and 2014–2022 (2023, a partial year, is excluded). Penguin and Random House are treated as separate publishers before the merger. The best-seller segment includes publishers with at least 20 deals with best-seller authors; the non-best-seller segment includes all firms with deals with non-best-seller authors. The index measures concentration of completed acquisitions, not bidder participation.

Table E2: HHI and ΔHHI After the PRH Merger (Best-Seller Threshold: 30 Deals)

Market	HHI pre-merger	HHI post-merger	ΔHHI
Best Seller	1475.9	2024.4	548.5
Non Best Seller	358.5	447.5	88.9

Notes: Herfindahl–Hirschman indices computed from publisher shares of deals signed in each segment and year, using Publishers Marketplace data. The unit is the deal; multi-book contracts count once. Pre- and post-merger values are averages of yearly indices over 2008–2013 and 2014–2022 (2023, a partial year, is excluded). Penguin and Random House are treated as separate publishers before the merger. The best-seller segment includes publishers with at least 30 deals with best-seller authors; the non-best-seller segment includes all firms with deals with non-best-seller authors. The index measures concentration of completed acquisitions, not bidder participation.

Table E3: HHI and ΔHHI After the PRH Merger (Dollar-Weighted)

Market	HHI pre-merger	HHI post-merger	ΔHHI
Best Seller	1899.6	2614.8	715.2
Non Best Seller	932.7	1172.5	239.8

Notes: Herfindahl–Hirschman indices computed from publisher shares of total advance dollars paid in deals signed in each segment and year, using Publishers Marketplace data. Advances are measured at the midpoint of each Publishers Marketplace tier (nice $25k, very nice $75k, good $175k, significant $375k, major $750k); only deals with a disclosed advance tier are included. Pre- and post-merger values are averages of yearly indices over 2008–2013 and 2014–2022 (2023, a partial year, is excluded). Penguin and Random House are treated as separate publishers before the merger. The best-seller segment includes publishers with at least 50 deals with best-seller authors; the non-best-seller segment includes all firms with deals with non-best-seller authors. The index measures concentration of advance spending on completed acquisitions, not bidder participation.

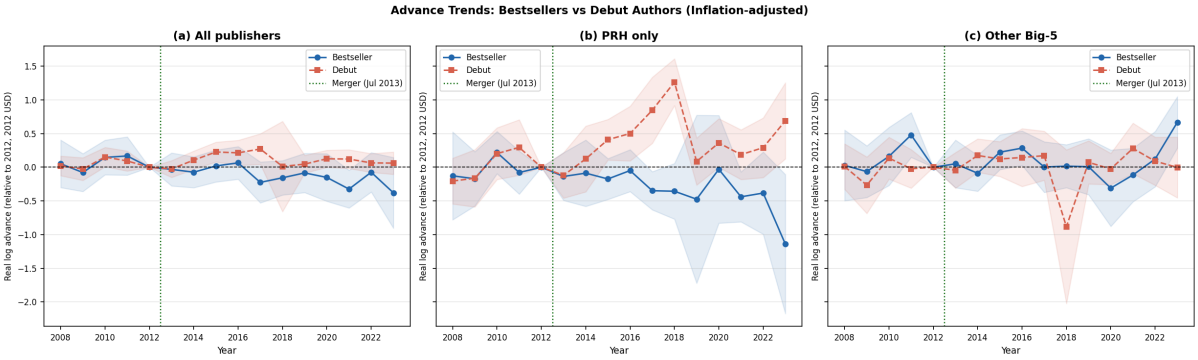
F Descriptive Evidence: Robustness

This appendix collects supplementary evidence for the retrospective analysis of Section 4: robustness of the event-study estimates to alternative comparison groups.

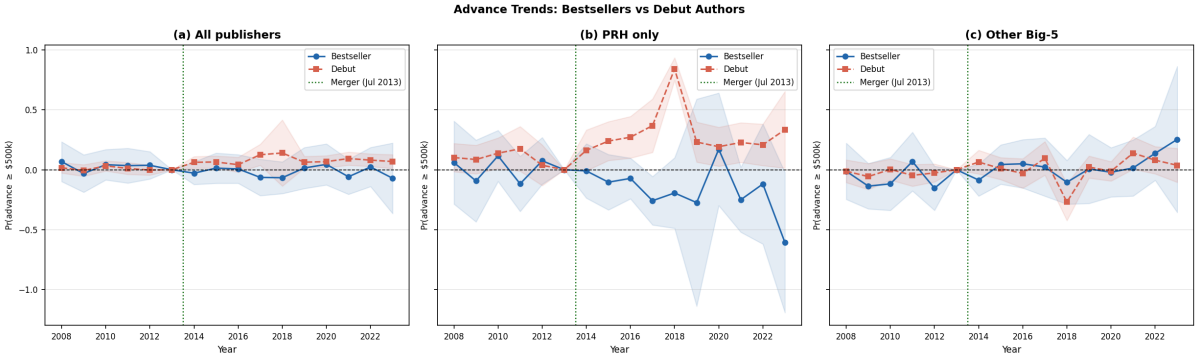
Figure F1 reports the event-study estimates for best-seller deals relative to debut deals only, rather than the pooled non-best-seller group of Figure 2. Debut authors are the comparison group where the merging firms should have lower competition overlap, so the narrower design guards against the possibility that mid-list deals, which the channel also lifts, contaminate the baseline comparison. The dynamics are unchanged: the best-seller differential is flat before the announcement and emerges after 2013, as in the pooled design.

Figure F1: Event-study: Advances Paid to Best-seller Authors Relative to Debut Authors

(a) Log advances



(b) Probability of an advance above $500,000



Notes: Event-study estimates for titles by best-seller authors relative to debut authors only, around the 2013 merger. Panel (a) reports the change in log advances (inflation-adjusted). Panel (b) reports the change in the probability that a deal exceeds $500,000. Within each panel, the three columns show, from left to right: all publishers, the merging parties, and the remaining Big Five. Standard errors are clustered at the imprint level. Estimates are inverse-probability-weighted to address selection into the disclosed-advance subsample.

G A Bargaining Alternative for the Downstream Market

The main text models publishers as setting retail prices in a multiproduct Bertrand–Nash game, which is how e-book prices are set under agency contracts. This appendix develops the leading alternative for the print channel: publishers negotiate wholesale prices with the retail platform, which then sets retail prices. Three results are established. First, with retail prices and quantities but no data on wholesale terms, contracts, or negotiation breakdowns, the bargaining weights are not identified: any weight vector is observationally equivalent to the recovered equilibrium (Proposition 2). Second, and more consequentially for this paper, publisher marginal costs, and consequently publisher variable profits, are unrecoverable under the bargaining regime for any weights (Proposition 4): the alternative is not a different specification of the paper's exercise but a regime under which the exercise is infeasible. Third, across the conduct assumptions for the print channel, the merger's downstream effect on the merging parties operates in the same direction as under the pricing model whenever the platform's retail response satisfies a sign condition (Assumption 2), and the one regime under which the effect vanishes is difficult to reconcile with the changes in advances observed after the 2013 merger (Lemma 3, Proposition 5).

G.1 Environment

Publisher f owns catalog $\mathcal{J}_f$ of titles and negotiates the terms under which a single retail platform (e.g., Amazon) sells them. The platform sets retail prices and earns a margin on every copy sold, consumers choose among titles according to the demand system of Section 5.1. Timing: in stage one, each publisher and the platform bargain bilaterally over the wholesale price w_j of every title j, following Horn and Wolinsky (1988). In stage two, the platform observes $\{w_j\}$ and sets retail prices to maximize total retail profit. Consumers then purchase. The equilibrium concept is subgame perfection. Two conventions follow the applied bargaining literature: negotiations are title-by-title, and disagreement payoffs are normalized to zero, consistent with the platform's dominance of the retail channel for the sample titles.[45]

G.2 Retail Stage and Margin Recovery

Taking wholesale prices as given, the platform chooses the retail price vector $\mathbf{p}$ to maximize

$$\Pi^R(\mathbf{p}) \;=\; \sum_j \left(p_j - w_j - c_j^R\right) q_j(\mathbf{p}), \tag{41}$$

where $q_j = M \cdot s_j(\mathbf{p})$ and $c_j^R \geq 0$ is the platform's own marginal cost of retailing title j (e.g., fulfillment, shipping). The recovery below imposes $c_j^R = 0$, an assumption whose role Proposition 2 makes explicit. Stacking the first-order conditions, with the share Jacobian $\Delta_{jk} = \partial s_j / \partial p_k$,

$$\mathbf{s}(\mathbf{p}) + \Delta(\mathbf{p})' \left(\mathbf{p} - \mathbf{w} - \mathbf{c}^R\right) = \mathbf{0}, \tag{42}$$

[45] In practice, publisher–platform negotiations are catalog-level periodic contracts, as the 2014 Hachette–Amazon dispute demonstrates publicly. Title-level bargaining with zero disagreement payoffs is a tractability device.

so the platform's per-unit margin over its perceived cost is pinned by demand alone:

$$\mathbf{m} \equiv \mathbf{p} - \mathbf{w} - \mathbf{c}^R = -\left[\Delta(\mathbf{p})'\right]^{-1} \mathbf{s}(\mathbf{p}). \tag{43}$$

Evaluated at observed prices with the estimated demand system, every element of the right-hand side of (43) is computable. This is the model's key property and its key limitation: prices are observed and margins are computed, so the total $w_j + c_j^R = p_j - m_j$ is known exactly, but nothing in the retail stage separates the wholesale price w_j from the platform's own retail cost c_j^R within it.

G.3 Bargaining Stage

The Nash problem for title j, with publisher weight $b_j \in (0,1)$ and zero disagreement payoffs, is

$$\max_{w_j} \left[w_j \, q_j \right]^{b_j} \left[\left(p_j - w_j - c_j^R \right) q_j \right]^{1-b_j}, \tag{44}$$

whose first-order condition at retail price p_j yields the closed form

$$w_j = b_j \left(p_j - c_j^R \right). \tag{45}$$

Under the conventions $c_j^R = 0$ and zero disagreement payoffs, combining (43) and (45) recovers a candidate weight product by product:

$$\hat{b}_j = \frac{p_j - m_j}{p_j} = \frac{w_j}{p_j} \tag{46}$$

that is, the ratio of the wholesale to the retail price. Mechanically, titles on which the platform's optimal margin is larger are assigned lower publisher weights; Section G.4 discusses what this recovered weight does and does not measure.

G.4 What the Model Recovers

Under the two conventions of Section G.3 the model delivers a clean recovery: combining the margin condition (43) with the Nash split (45) yields a unique weight per title, $\hat{b}_j = (p_j - m_j)/p_j$, computable from observed prices and the estimated demand system alone. The recovery is internally consistent and, as a maintained construction, unobjectionable: it delivers the wholesale-to-retail ratio implied by the platform's optimal pricing, title by title.

Its content, however, should be read for what it is. The recovered $\hat{b}_j$ equals one minus the platform's equilibrium relative markdown into which no information about the division of surplus ever enters: cross-title variation in $\hat{b}_j$ reflects heterogeneity in the platform's markdown incentives, not differences in publishers' negotiating strength. And the conventions that produce it are untestable with the available data:

Proposition 2 (Observational equivalence)**.** *Let the observables be retail prices, quantities, and product characteristics* $\{\mathbf{p}_t, \mathbf{q}_t, \mathbf{x}_t\}$*, and let the demand system be given. Fix any candidate weight vector* $\tilde{\mathbf{b}} \in (0,1)^J$*. Then there exist retail costs* $\mathbf{c}^R \geq \mathbf{0}$ *and disagreement payoffs* $\{(d_j^P, d_j^R)\}$*,*

feasible for the bargaining problem, such that the model with parameters $(\tilde{\mathbf{b}}, \mathbf{c}^R, \mathbf{d})$ *generates the same observables.*

Proof. The platform's optimality at the observed prices requires $\mathbf{p}^{\text{obs}} - \mathbf{w} - \mathbf{c}^R = \mathbf{m}$ with $\mathbf{m}$ given by (43). The data thus determine the sum $w_j + c_j^R = p_j^{\text{obs}} - m_j$ for every title and place no further restriction on the supply side.

Step 1 (retail costs absorb the wholesale price). For any target $\tilde{w}_j \in (0,\, p_j^{\text{obs}} - m_j]$, set $c_j^R = p_j^{\text{obs}} - m_j - \tilde{w}_j \geq 0$. The platform's perceived cost, and hence its pricing problem and equilibrium prices, are unchanged. All observables are unchanged.

Step 2 (disagreement payoffs absorb the weight). Fix any $\tilde{b}_j \in (0, 1)$. With general disagreement payoffs, both parties' surpluses are linear in w_j at fixed retail price, so the Nash first-order condition is

$$\tilde{b}_j \left[\left(p_j^{\text{obs}} - \tilde{w}_j - c_j^R \right) q_j - d_j^R \right] = \left(1 - \tilde{b}_j \right) \left[\tilde{w}_j q_j - d_j^P \right],$$

one linear equation in (d_j^P, d_j^R). A feasible solution exists: for small $\varepsilon > 0$, set $d_j^P = \tilde{w}_j q_j - \varepsilon$. The equation then gives $d_j^R = \left(p_j^{\text{obs}} - \tilde{w}_j - c_j^R \right) q_j - \frac{1-\tilde{b}_j}{\tilde{b}_j}\, \varepsilon$, and the two parties' gains from trade equal ε and $\frac{1-\tilde{b}_j}{\tilde{b}_j}\, \varepsilon$, both strictly positive. The positivity of both parties' profits at the candidate is what leaves room to place each disagreement payoff strictly between zero and the corresponding profit, so that the observed split can be attributed to threat points rather than to the weight. Since both stages hold at the observed equilibrium under $(\tilde{\mathbf{b}}, \mathbf{c}^R, \mathbf{d})$, the observables are unchanged. As $\tilde{\mathbf{b}}$ was arbitrary, the identified set for the weights is $(0, 1)^J$. □

Identification in the empirical bargaining literature is obtained by observing precisely what this construction shows to be indispensable: negotiated prices themselves (Gowrisankaran et al., 2015), or disagreement events that reveal threat points. Neither is available here, with wholesale terms being confidential, and the single public negotiation breakdown (Hachette–Amazon, 2014) predates the sample and operated through availability degradation rather than delisting. None of this prevents adopting the conventions and the recovery as a maintained construction. The binding problem lies elsewhere.

G.5 The Counterfactual Wall

The difficulty with the bargaining alternative is not in the baseline but in the fact that no merger counterfactual can be built on it with the available data.

Proposition 3 (Merger neutrality under the recovery conventions)**.** *Under the conventions required for the recovery of the bargaining weights* ($\mathbf{d} = \mathbf{0}$, $\mathbf{c}^R = \mathbf{0}$)*, a publisher merger has no effect on any negotiated wholesale price, any retail price, or any welfare object.*

Proof. The Nash solution (45) for title j depends only on (b_j, p_j), and the platform's pricing problem (42) depends only on wholesale costs and demand, neither object is a function of catalog ownership. The pre-merger equilibrium therefore remains an equilibrium, unchanged, after any reassignment of titles across publishers. □

The proposition formalizes a simple point: the channel through which a merger operates in a bargaining model, i.e., the improvement of the merged entity's disagreement position, since

a breakdown now threatens the platform with the loss of both catalogs, is removed by the same normalization that makes the baseline recoverable. The two halves of the exercise are incompatible: the assumption needed to estimate the model deletes the object the counterfactual would move.

Nor can the disagreement payoffs be reinstated empirically. Modelling them requires each publisher's payoff in the event of a breakdown with the platform and the corresponding profit the platform would retain without the publisher's catalog. Sales outside Amazon are not observed in these data, so the threat points cannot be constructed, and the convention route back to them is closed by Proposition 2, which shows the baseline data cannot pin them either. The counterfactual is thus doubly blocked: by the normalization, and by the data one would need to relax it. The convention in the literature (Gowrisankaran et al., 2015) is unavailable here for exactly these reasons: the disagreement points were normalized away to recover the baseline, and the outside-channel data required to model them instead do not exist.

The only remaining lever is to move the recovered weights directly: impose $b_j^{\text{post}} > \hat{b}_j$ on the merged entity's titles and re-solve the retail equilibrium. But this lever does not deliver the comparative static one might expect. At fixed retail prices, an increase in b_j is a pure transfer from the platform to the publisher, the harm enters through the platform's re-pricing response. A higher wholesale share raises the platform's perceived cost, retail prices rise, and quantities fall so the publisher's profit on title j, $(b_j p_j - mc_j)\, q_j$, moves ambiguously: the share effect raises it, the quantity effect lowers it, and the net depends on how much share there is left to gain. For a title whose $\hat{b}_j$ is already high, a further increase buys little share and costs real quantity, and profit can fall. Only titles with low recovered shares are safely improved. A counterfactual built on this lever must therefore be designed around its intended conclusion: the floor rule

$$b_j^{\text{post}} \;=\; \max\Big\{\hat{b}_j,\; \bar{b}\Big\} \quad \text{for } j \in \mathcal{J}_{\text{PRH}} \cup \mathcal{J}_{\text{SS}}, \qquad b_j^{\text{post}} = \hat{b}_j \;\; \text{otherwise}, \tag{47}$$

raises shares precisely and only where they are low enough for the share effect to dominate. I accordingly treat constructions of this form as illustrations of the model's direction, not as measurements of the merger: the selection of the treated set is part of the construction, not an implication of the data.

G.6 Publisher Costs and Profits Are Unrecoverable

A further limitation is independent of the counterfactual and bears directly on this paper's exercise.

Proposition 4 (Non-recoverability of publisher profits)**.** *Under the bargaining regime, publisher marginal costs mc_j, and hence publisher variable profits, are not recoverable from retail prices, quantities, and characteristics, for any bargaining weights.*

Proof. mc_j appears in no equilibrium condition of the model: the publisher takes no pricing action whose optimality would reveal it, and with disagreement payoffs normalized to zero it does not enter the Nash split (45). The observables therefore place no restriction on mc_j beyond $mc_j \geq 0$.[46] □

[46]With general disagreement payoffs, the publisher's threat point involves its profit in outside channels and

The consequence is structural rather than statistical. The paper's mechanism runs through publishers' expected variable profits: $\mathbb{E}_\xi[\Delta_{fj}]$ enters the valuation equation (16), disciplines bids, and determines the division of surplus between publishers and authors in Section 5.3. Under the bargaining regime none of these objects exists: the model delivers publisher revenues but not profits, so the upstream auction cannot be estimated, the downstream-to-upstream feedback cannot be computed, and the distributional accounting between publishers and authors cannot be performed. The bargaining alternative is thus not an alternative specification of the paper's exercise, but a conduct regime under which the exercise is infeasible with retail data of any richness. Under the pricing regime of the main text, by contrast, the publishers' own first-order conditions identify mc_{jt} (equation (8)), a recovery that is, like all markup estimation, conditional on conduct, and whose conduct assumption Section 2 supports institutionally.

One distinction within this negative result is used below: while the level of publisher profit is unidentified, a merger-induced profit change satisfies $\Delta[(w_j - mc_j)q_j] = \Delta(w_j q_j) - mc_j\,\Delta q_j$ under production costs unchanged by the merger, so for titles whose quantities weakly fall, the wholesale-revenue change is a lower bound on the profit change. Signs of changes are recoverable where revenue changes are, but levels are not, and the paper's quantitative exercise requires both.

G.7 Direction of the Merger's Effect: The Regimes Agree

This final subsection shows that where the bargaining model's implications can be characterized at all, they run in the same direction as the pricing model of the main text under a sign condition on the platform's re-pricing that I verify at the estimates (Assumption 2), so conduct uncertainty affects the magnitude of the paper's downstream channel but not its sign.

Three conduct regimes exhaust the possibilities for the print channel: (i) publishers set retail prices; (ii) the platform sets retail prices over wholesale terms that do not respond to the merger; and (iii) the platform sets retail prices over wholesale terms that improve for the merged entity. By Proposition 3, the recoverable version of the bargaining model is regime (ii): the merger is exactly neutral downstream, and the full model collapses to the upstream-only counterfactual of Section 7.3. Regime (ii) is therefore the polar bound on the paper's mechanism, and it has a testable implication for advances. If the merger is neutral downstream, publishers' valuations are unchanged, so advances can move only in auctions that both merging parties contested and are exactly unchanged elsewhere, as in the upstream-only exercise of Section7.3. The evidence from the 2013 merger is difficult to reconcile with this prediction: advances for non-best-seller and debut titles, where the merging parties rarely competed with each other, rose after the merger, and rose most at the merged firm (Section 4). Regime (ii) also applies at most to the print channel, since publishers set e-book prices directly, so even under it the downstream channel would be attenuated rather than eliminated. Under regime (iii), the effects can be signed for the platform and the merging parties.

Lemma 3 (The platform loses)**.** *The platform's maximized retail profit is strictly decreasing in each wholesale price: $\partial\Pi^{R*}/\partial w_j = -q_j(\mathbf{p}^*) < 0$. Any merger outcome that weakly raises the merged entity's negotiated wholesale shares weakly lowers platform profit.*

hence mc_j, but this adds a further unobserved object to the split rather than an equation that identifies costs.

The lemma is an envelope-theorem statement of a simple observation: the platform prices optimally before the merger, so it cannot benefit from the merger through re-pricing alone given that any re-pricing gains available afterward were available before. Within this regime, none of the merger's gain accrues to the retailer.

Assumption 2 (Sign conditions for the retail response)**.** *Let p^{pre} and p^{post} denote the platform's optimal retail prices before and after the floor rule (47), let $J^{\dagger} \equiv \{j \in J_{\text{PRH}} \cup J_{\text{SS}} : \hat{b}_j < \bar{b}\}$ be the set of treated titles, assumed nonempty, and write $q_j^{\text{pre}} \equiv q_j(p^{\text{pre}})$, $q_j^{\text{post}} \equiv q_j(p^{\text{post}})$. Then*

1. $q_j^{\text{post}} \leq q_j^{\text{pre}}$ *for every $j \in J^{\dagger}$; and*
2. $\sum_{j \in J^{\dagger}} \left(\bar{b}\, p_j^{\text{post}} q_j^{\text{post}} - \hat{b}_j\, p_j^{\text{pre}} q_j^{\text{pre}}\right) > 0.$

Neither condition follows from substitutability and the platform's second-order conditions. For a marginal change in the weights, the price response is $\partial p / \partial b = (-H)^{-1} G$, where $H \equiv \nabla_p^2 \Pi^R$ and $G_{jk} \equiv \partial^2 \Pi^R / \partial p_j \partial b_k = -\partial R_k / \partial p_j$. The second-order condition makes $-H$ positive definite, hence invertible, but places no sign restriction on the entries of $(-H)^{-1}$. Under substitutability, moreover, G has a nonnegative diagonal at the optimum and nonpositive off-diagonal entries: a higher publisher share on title k raises the platform's incentive to price k up but lowers its incentive to price other titles up, since diverting demand toward k is now worth less to it. The sign of each price and quantity response is therefore an empirical matter.Assumption 2 is therefore maintained rather than derived: it is the condition under which the bargaining model's implications for the merged entity can be signed.

Proposition 5 (Direction, merged entity)**.** *Let Assumption 2 hold and let production costs mc_j be unchanged by the merger. Then:*

(i) Decomposition. *For every $j \in J^{\dagger}$, the change in the merged entity's wholesale revenue on title j is*

$$\Delta(w_j q_j) = \underbrace{(\bar{b} - \hat{b}_j)\, R_j^{\text{post}}}_{\textit{share effect} > 0} + \underbrace{\hat{b}_j\, \Delta R_j}_{\textit{re-pricing effect}}, \qquad R_j \equiv p_j q_j. \tag{48}$$

(ii) Profit bound. *For every $j \in J^{\dagger}$, the change in variable profit satisfies $\Delta\big[(w_j - mc_j) q_j\big] \geq \Delta(w_j q_j)$.*

(iii) Direction. *The merged entity's variable profit on the treated titles strictly rises: $\sum_{j \in J^{\dagger}} \Delta\big[(w_j - mc_j) q_j\big] > 0$.*

Consequently, the downstream value the merged entity attaches to a marginal manuscript in the treated set rises, and the downstream channel of Section 5 operates under bargaining in the same direction as under the pricing model of the main text.

Proof. With $c^R = 0$ and $w_j = b_j p_j$, the merged entity's wholesale revenue on title j is $b_j R_j$, with $b_j = \hat{b}_j$ before the merger and $b_j = \bar{b}$ after it for $j \in J^{\dagger}$.

(i) Add and subtract $\hat{b}_j R_j^{\text{post}}$: $\bar{b} R_j^{\text{post}} - \hat{b}_j R_j^{\text{pre}} = (\bar{b} - \hat{b}_j) R_j^{\text{post}} + \hat{b}_j (R_j^{\text{post}} - R_j^{\text{pre}})$. The share effect is positive because $\bar{b} > \hat{b}_j$ on $J^{\dagger}$ and $R_j^{\text{post}} > 0$.

(ii) With mc_j unchanged, $\Delta\big[(w_j - mc_j) q_j\big] = \Delta(w_j q_j) - mc_j\, \Delta q_j$. Since $mc_j \geq 0$ and $\Delta q_j \leq 0$ by Assumption 2(a), the second term is nonnegative.

(iii) Sum (ii) over $J^\dagger$: $\sum_{j\in J^\dagger}\Delta\big[(w_j - mc_j)q_j\big] \geq \sum_{j\in J^\dagger}\Delta(w_j q_j) > 0$, where the strict inequality is Assumption 2(b). The level of mc_j, which Proposition 5 shows is not identified, is not needed for the sign. □

Corollary 2 (Non-merging publishers: ambiguity)**.** *Under the conditions of Proposition 5, the effect of the merger on a non-merging publisher's variable profit is ambiguous, and its sign cannot be determined from the observables. The rival's weight b_j is unchanged by (47), so its per-unit wholesale price $w_j = b_j p_j$ moves with its retail price, and its wholesale revenue with $p_j q_j$. Neither has a sign: the platform's response on rival titles combines a direct incentive to lower their prices, since diverting demand toward treated titles is worth less to it, with the indirect effect of higher prices on treated titles. The rival's profit change, $\Delta\pi_j = b_j\,\Delta(p_j q_j) - mc_j\,\Delta q_j$, therefore depends on the sign of its revenue change and on the level of mc_j, which Proposition 4 shows is not recoverable.*

Proof. Under equation (47), $b_j = \hat{b}_j$ for every rival title, so $\Delta\pi_j = b_j\,\Delta(p_j q_j) - mc_j\,\Delta q_j$ with mc_j unchanged. By the argument preceding Assumption 2, the signs of Δp_j and Δq_j are not determined by substitutability and the platform's second-order conditions, so $\Delta(p_j q_j)$ can take either sign. The second term is positive when $\Delta q_j < 0$ and negative when $\Delta q_j > 0$, and in both cases it scales with the unidentified mc_j. The observables therefore place no restriction on the sign of $\Delta\pi_j$. □

For non-merging publishers, the effect is characterized rather than signed. A rival's negotiated terms are unchanged, what changes is its environment. Higher retail prices on treated titles shift demand toward substitutes, raising rival titles' quantities at given prices. The platform then re-optimizes retail prices over its entire catalog, and because it internalizes substitution among all titles, the sign of its response on rival titles is not general. A rival publisher's revenue changes by its fixed wholesale share of the change in its titles' retail revenue, $b_j\,\Delta(p_j q_j)$: positive whenever the substitution-driven quantity gain dominates the platform's re-pricing response, and attenuated relative to the Bertrand case, in which rivals both re-optimize their own prices and retain their full margins.

The welfare accounting differs across regimes in one respect worth recording. Under the pricing model of the main text, the merger's downstream effects are divided among publishers, authors (through advances), and consumers. Under the bargaining regime a further agent enters: the platform, which by Lemma 3 strictly loses from any strengthening of the merged entity's terms. Under Assumption 2, the merged entity's downstream gain keeps its sign, while the rivals' gain is no longer signed (Corollary 2). Part of the merging parties' gain is extracted from the platform's margin rather than from consumers alone, and rivals' gains, where they occur, reach their bids attenuated by the fixed wholesale share. The bargaining alternative thus reallocates the merger's burden across the vertical chain without reversing the merged entity's downstream gain.

Summing up: the recoverable version of the bargaining model predicts that the merger does nothing downstream. The versions under which the merger does something cannot be estimated, and and where their implications can be signed, which for the merged entity requires Assumption 2, they reproduce the direction of the main text's downstream channel, with part of

the incidence shifted onto the platform. Conduct uncertainty therefore moves the size of the revaluation channel, with the upstream-only benchmark of Section 7.3 as its lower bound, without threatening the incidence result for jointly contested manuscripts, which operates through the competition channel under every regime.

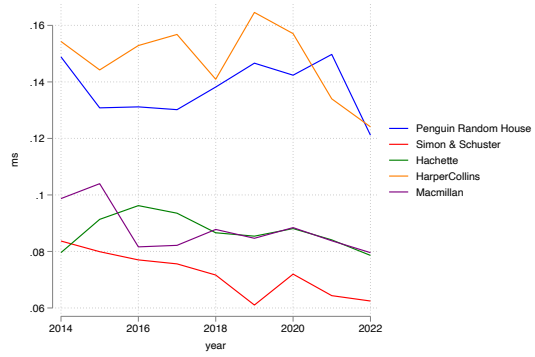
H Acquisition and Publication Trends

In This appendix I show that firms do not tend to vary their acquisition or publication trends dramatically. I use this to justify the idea that, when acquiring titles, any firm may have a good idea of what share of titles would be acquired by any of the other firms. First, I show the yearly upstream market share of each of the Big Five firms and for a set of the biggest firms outside the Big Five in Figures H1a and H1b respectively.[47]

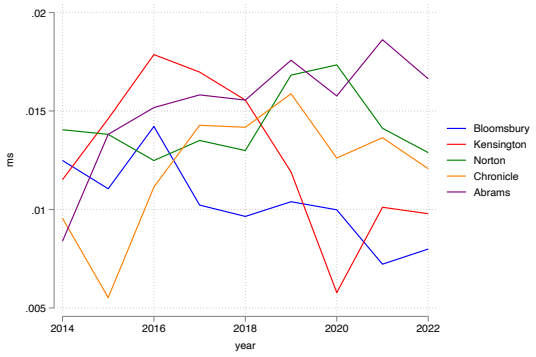
In Figure H1a I see that the yearly share of PRH's acquisition varies between 12% and 15% and that of HarperCollins between 13% and 16%. Below those we find Hachette and Macmillan with a share between 8% and 10% and Simon & Schuster with the lowest rates, between 6% and 8%. On the other hand, looking at other big firms outside the Big Five in Figure H1b I see that, even though their rates have a more erratic pattern, each firm usually stays between 1% and 2%.

Figure H1: Firms' Upstream Market Share.

(a) Yearly Upstream Market Share of Big Five Firms



(b) Yearly Upstream Market Share of Other Big Firms.

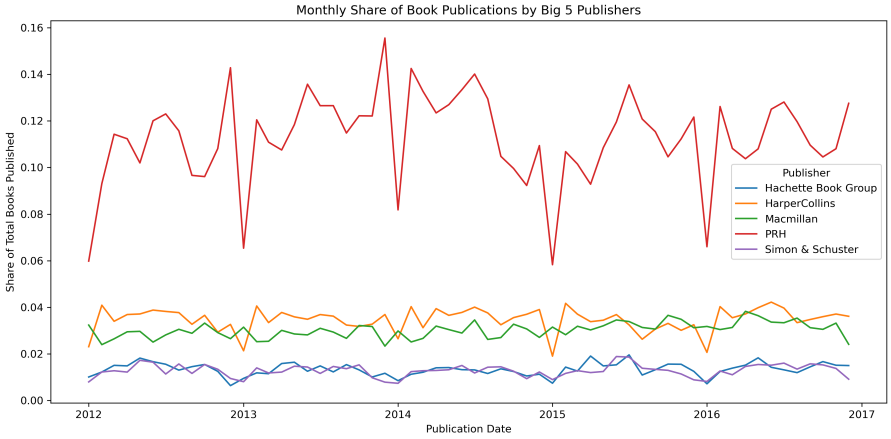


Notes: Deals data from Publishers Marketplace.

In Figures H2 and H3 I do an analogous exercise using publishing data from Goodreads. The trends show a similar pattern here, with the Big Five having a constant share through the years. Even though the variance within the year seems high for the other set of big firms, it tends to be fairly constant through the period. For instance, Norton publishes between 0.1% and 0.2% of the total titles in a month. Notice that here the shares are much lower compared to the ones presented before, since this includes all the self-published titles and titles that are published with very small firms, which do not go through the process of rights sale described before.

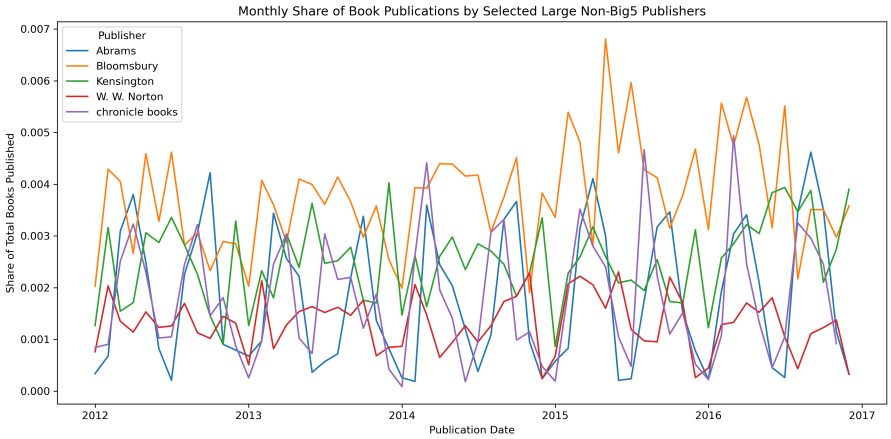
[47] Here market share is defined as the number of purchased titles in a given year, divided the total number of titles purchased by all firms.

Figure H2: Monthly Share of Published Titles by Big Five Firms.



Notes: Goodreads data from Wan and McAuley (2018) and Wan et al. (2019).

Figure H3: Monthly Share of Published Titles by Other Big Firms.



Notes: Goodreads data from Wan and McAuley (2018) and Wan et al. (2019).

I Descriptive Evidence: Changes in Supply

In this Appendix I provide evidence supporting the idea that, after a merger there is no significant change in the supply of titles, in order to validate the assumption of exogenous set of manuscripts in the structural model. Under the idea that every book that authors write end up published (either with a publisher or self-published), I show that after the Penguin Random House merger, the merged entity did not change how many titles it publishes.

The measure of title output comes from the Open Library bibliographic database, an openly redistributable catalog of edition-level records maintained by the Internet Archive. I use the July 2026 vintage of its complete editions snapshot, which contains approximately 56 million records.[48] Because records continue to accumulate retroactively, the vintage of the snapshot matters and is reported here so that the extract can be reconstructed.

I retain records whose publication-date field resolves to a year between 2010 and 2017. Each record is credited to its first-listed publisher, which keeps the panel additive in the sense that the sum across publishers equals the number of records. I count distinct works rather than distinct editions, so that hardcover, paperback, electronic and audio releases of the same title are counted once. This matters because format proliferation is uneven across firms, and counting editions mechanically understates established houses relative to self-publishing platforms. The resulting panel contains 7,427,723 works attributed to 439,824 distinct publishers, after collapsing 1,508,597 format duplicates.

In Figure I1 I show the number of titles in the industry per month between 2010 and 2017. Here, it can be seen that after the merger of Penguin and Random House there is not a clear change in the number of titles produced. In Figure I2 I show the evolution for each of the Big Five firms, in order to check whether the change of PRH titles is different to that of its closest rivals. Overall, each Big Five experiences a decrease in the published titles, most likely due to the expansion of self-publishing during the time. In Table I1 I show the relative changes by firm, where PRH suffers the highest change, closely followed by HarperCollins and Simon & schuster. Potential differences in levels can also be attributed to the risk of being replaced by self-publishing.

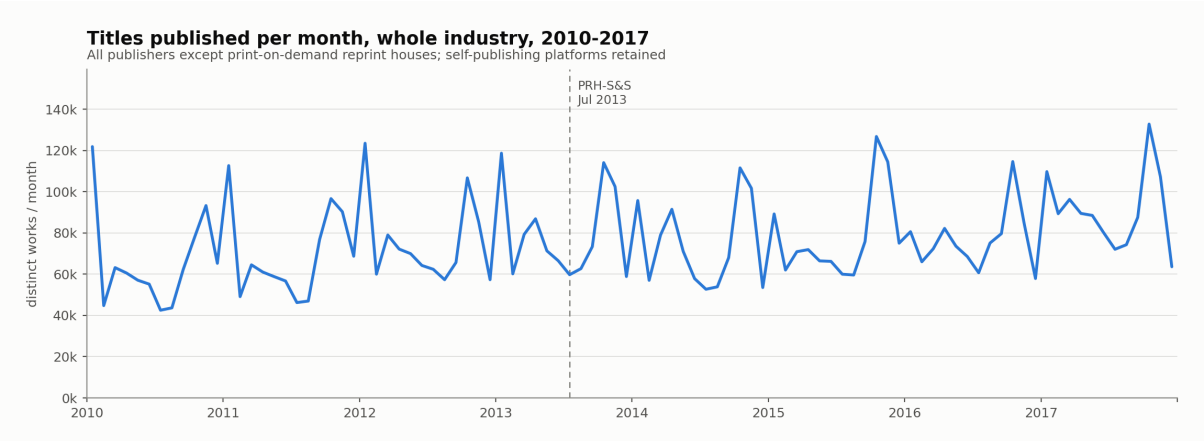
As a robustness to this last exercise, I do the same using the Goodreads records from Wan and McAuley (2018) and show the trends by firm in Figure I3. As observed with the data from Open Library, trends for all of the Big Five firms are similar.[49]

This evidence rule out the two previously outlined concerns, i.e., that the merger can affect the overall supply of titles, and that the new merged entity can change its supply of products.

[48] Open Library is the most comprehensive freely available source with meaningful trade coverage: it is strongest on English-language, ISBN-bearing books published in the United States and the United Kingdom, and thinner on non-English, non-ISBN and small-press output.

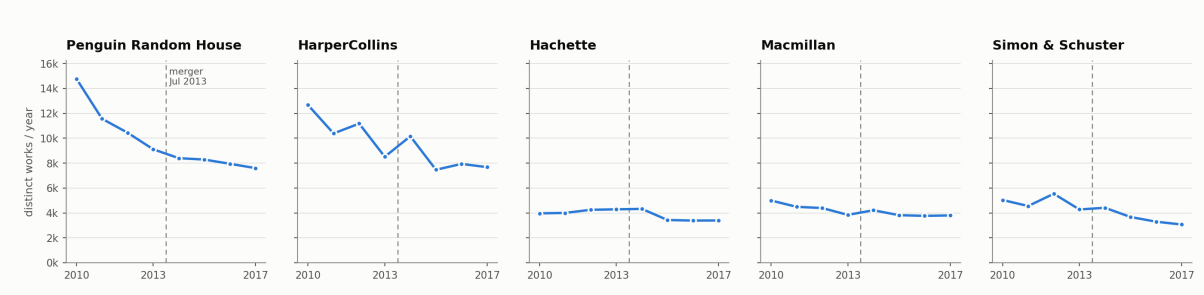
[49] I do not show the overall number of titles from this data due to a selection issue. Titles published early (e.g., January 2012) are more likely to be registered in these data than titles published later (e.g., December 2017), given that the first ones have more time to read and get reviewed in the platform. This biases the data towards more old titles.

Figure I1: Monthly Titles Published Between 2010 and 2017



Notes: Data from Open Library. Each observation is year-specific, summing all titles by traditional publishers and self-publishing firms, for both fiction and non-fiction market.

Figure I2: Monthly Titles Published by Each Big Five Between 2010 and 2017



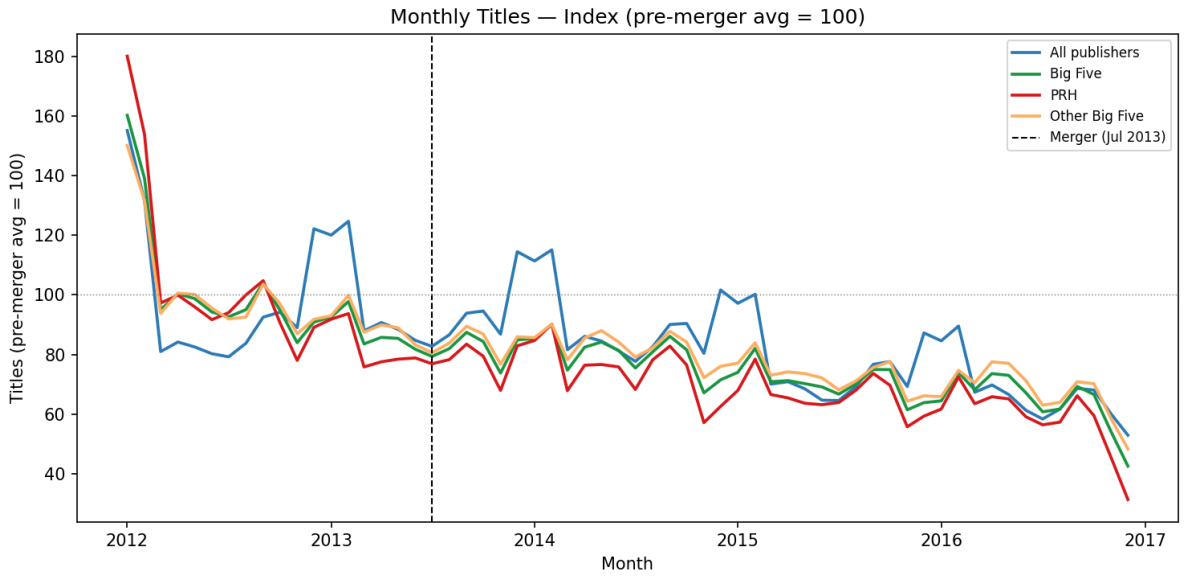
Notes: Data from Open Library. Each panel represents one of the Big Five firms, with Penguin Random House counting titles of the merging parts before the merger in June 2013.

Table I1: Change in title output by firm, 2010–2017

	Titles published		Change
Firm	2010	2017	(%)
Penguin Random House	11,559	7,596	−34.3
HarperCollins	10,383	7,674	−26.1
Hachette	4,001	3,394	−15.2
Macmillan	4,490	3,799	−15.4
Simon & Schuster	4,555	3,068	−32.6
Big Five	34,988	25,531	−27.0

Data from Open Library.

Figure I3: Monthly Titles Published by Each Big Five Between 2012 and 2017



Notes: Goodreads data from Wan and McAuley (2018). Trends are normalized to the pre-merger average by firm.

J Auction Model With Correlation in Private Shocks

The auction model in Section 5.2.2 assumes that, conditional on the observable characteristics of the title and of the bidders, publishers' valuations are independent. This appendix relaxes that assumption by letting the cost shocks share a title-level component. It first explains what the extension means and how it changes the model, then how it is estimated, why the degree of correlation is imposed rather than estimated, and how the estimates and the merger's competitive effect change as correlation rises.

J.1 Model

In the main specification, the cost shock ε_{fj} in (18) is drawn independently across the publishers bidding for title j. Here, I replace it with

$$\varepsilon_{fj} = \sigma^{\varepsilon}_{k(f),h(j)} \left(\sqrt{\kappa}\, \eta_j + \sqrt{1-\kappa}\, u_{fj} \right), \qquad \eta_j,\, u_{fj} \overset{iid}{\sim} \mathcal{N}(0,1), \qquad \kappa \in [0,1). \tag{49}$$

The component η_j is common to every publisher bidding for j, and u_{fj} is specific to publisher f. Three properties of this specification are useful. First, the marginal distribution of each valuation is unchanged: $V_{fj} \mid X_j \sim \mathcal{N}(\bar{V}_{fj}, \sigma^2_{k(f),h(j)})$, exactly as in (19). The parameter κ reallocates a fixed amount of dispersion between a common and an idiosyncratic part, so $\sigma^{\varepsilon}_{k(f),h(j)}$, $\phi_{k(f)}$, $\psi_{h(j)}$, $\chi_{g(j)}$ and $r_{h(j)}$ have the same interpretation for every κ, and $\kappa = 0$ is the main specification.

Second, any two bidders' valuations for the same title have correlation κ:

$$\operatorname{Corr}\left(V_{fj}, V_{f'j} \mid X_j\right) = \kappa, \qquad f \neq f' \in \mathcal{N}_j. \tag{50}$$

For $\kappa \geq 0$ this joint normal distribution is affiliated in the sense of Milgrom and Weber (1982), where a high valuation of one bidder makes high valuations of the others more likely.[50] Conditional on η_j, valuations are independent, so the model belongs to the class of conditionally independent private values (Li et al., 2000).

Third, values remain private. The common component η_j is a title-specific part of the cost of bringing j to market, net of revenue not captured by Δ_{fj}. It is shared across houses but not observed by the econometrician. Examples include how much editorial work the manuscript needs, the author's own publicity platform, and the title's potential in subsidiary rights markets. Because η_j enters each publisher's own payoff and says nothing about any rival's payoff beyond what the rival already knows, there is no winner's curse. Remaining active until the price reaches one's own valuation is still a weakly dominant strategy, so the auction outcome is still (20). It does not matter whether publishers observe η_j and u_{fj} separately or only their sum ε_{fj}: under private values, the equilibrium of the ascending auction depends only on each bidder's own valuation.

Why correlation matters for the counterfactual. The merger lowers the advance of a title only when the two merging publishers hold the two highest valuations. In that case the advance falls from $V_{(2)j}$ to $V_{(3)j}$. Correlation changes both parts of this calculation. It raises the probability that the merging publishers occupy the top two positions, because it compresses valuations around the common component and makes the ranking depend more on the systematic differences $\bar{V}_{fj} - \bar{V}_{f'j}$ across publishers. It also shrinks the gap between the second- and third-highest valuations. Under independence, observing a high advance says nothing about the losing rivals' valuations, so the price-setting rival's replacement is a fresh draw from its own distribution. Under correlation, a high advance signals a high η_j, which raises every bidder's valuation, including the rival that would set the price after the merger. The first force increases the number of titles affected by the merger, and the second reduces the price effect on each of them. The sign of the net effect is an empirical question, which Table J1 addresses.

J.2 Estimation

Conditional on η_j, valuations are independent with

$$V_{fj} \mid X_j, \eta_j \sim \mathcal{N}\Big(\bar{V}_{fj} - \sigma^{\varepsilon}_{k(f),h(j)}\sqrt{\kappa}\,\eta_j,\ (1-\kappa)\,\sigma^2_{k(f),h(j)}\Big), \tag{51}$$

so the likelihood of the main specification applies conditionally on η_j. Let $L_j^{\mathrm{IPV}}(\theta \mid \eta)$ denote the likelihood of the observed pair (w_j, a_j) derived in Section 6.2, evaluated at these conditional means and standard deviations. The likelihood under correlation integrates out the common component:

$$L_j(\theta;\kappa) = \int L_j^{\mathrm{IPV}}(\theta \mid \eta)\,\phi(\eta)\,d\eta \;\approx\; \sum_{q=1}^{Q} \omega_q\, L_j^{\mathrm{IPV}}(\theta \mid \eta_q), \tag{52}$$

[50] A multivariate normal vector is affiliated if and only if the off-diagonal elements of its precision matrix are non-positive. With covariance matrix DRD, where D is the diagonal matrix of the $\sigma^{\varepsilon}_{k(f),h(j)}$ and R has ones on the diagonal and κ elsewhere, the off-diagonal elements of R^{-1} equal $-\kappa/[(1-\kappa)(1+(N_j-1)\kappa)] \leq 0$.

where ϕ is the standard normal density and $\{(\eta_q, \omega_q)\}$ are Gauss–Hermite nodes and weights.

For each value of κ, θ is estimated by maximizing $\sum_j \log L_j(\theta; \kappa)$. The bidder sets $\mathcal{N}_j$ are the same draws used in the main specification, so the estimates differ across κ only because of the correlation. At $\kappa = 0$, a single node suffices, and the log-likelihood reproduces the main specification's log-likelihood exactly at any θ. Each κ is started from the estimate at the previous value of the grid, beginning from the main estimates at $\kappa = 0$.

J.3 Non-identification of κ

The data record one transaction price and the winner's identity per auction. With this information, the dependence between bidders' valuations is not identified.

Nonparametrically, Athey and Haile (2002) show that in ascending auctions the transaction price and the winner's identity identify an asymmetric independent private values model. The same data do not identify models with correlated private values or common values, and they cannot distinguish those models from independence. The reason is dimensional: a joint distribution of N_j valuations cannot be recovered from one order statistic per auction without restricting the dependence. Aradillas-López et al. (2013) show that, with transaction prices alone, correlated private values models are at best partially identified.

The component η_j is also a form of unobserved auction heterogeneity. Separating such heterogeneity from bidder-specific shocks requires at least two bids per auction (Krasnokutskaya, 2011), because it is the co-movement of bids within an auction that reveals the common component. An ascending auction with a single recorded price provides no within-auction co-movement to exploit.

In the parametric model (49), the likelihood does depend on κ. A larger common component raises the dispersion of advances across titles with the same observables and narrows the gaps between order statistics within a title. This variation, however, identifies κ only through the normality assumption and through the bidder sets, which are themselves imposed (see Section 5.2.2). The same patterns in advances can be generated by different values of $\sigma^{\varepsilon}_{k(f),h(j)}$, by a different number of bidders, or by a different composition of rival sets. An estimate of κ would therefore reflect these auxiliary assumptions rather than information in the data. I instead fix κ on a grid, $\kappa \in \{0, 0.1, 0.25, 0.5, 0.75\}$, re-estimate all other parameters at each value, and treat the results as a sensitivity analysis. For completeness, Table J1 reports the maximized log-likelihood at each κ.

J.4 Results

Table J1 reports the estimates for fiction titles with $N_j = 3$.

Table J1: Auction estimates under correlated cost shocks (fiction, $N_j = 3$)

	$\kappa = 0$ (main)	$\kappa = 0.1$	$\kappa = 0.25$	$\kappa = 0.5$	$\kappa = 0.75$
A. Fit					
Log-likelihood	−5,947.8	−5,954.0	−6,007.3	−6,191.9	−6,608.1
Change relative to $\kappa = 0$	0.0	−6.2	−59.5	−244.1	−660.3
Simulated / observed mean advance	1.07	1.24	1.28	1.33	1.38
B. Dispersion of the winner's valuation					
Mean $\sigma_{k(f),h(j)}$	413	403	387	366	350
Share of variance: observables (%)	20.8	21.8	23.1	24.4	16.6
Share of variance: common shock (%)	0.0	7.8	19.2	37.8	62.5
Share of variance: publisher shock (%)	67.6	70.4	57.7	37.8	20.8
C. Publisher-group effects relative to non-Big Five					
Penguin Random House	379	374	356	313	288
Simon & Schuster	281	277	268	260	251
D. Loss of competition between PRH and S&S					
Pr(PRH and S&S hold the top two valuations) (%)	46.7	48.3	51.5	56.6	65.3
Change in advance (%)	−35.4	−35.5	−35.3	−33.3	−30.4

Notes: Each column re-estimates the auction model with the correlation of cost shocks fixed at κ; the column $\kappa = 0$ is the main specification. The likelihood is integrated over the common shock with Gauss–Hermite quadrature. Panel A statistics are computed from 500 simulated auctions per title. Panel B reports the mean standard deviation of the winner's cost shock and decomposes the variance of the winner's valuation into the variance of $\bar{V}_{wj}$ and the common ($\kappa\sigma^2$) and publisher-specific ($(1-\kappa)\sigma^2$) parts of the cost shock. Panel C reports the difference in the valuation intercept, $-(\phi_k - \phi_{\text{non-Big Five}})$. Panel D holds all estimated parameters fixed and, in the titles whose bidder set contains both PRH and Simon & Schuster, replaces the two publishers with a single bidder whose valuation is the larger of theirs. It isolates the upstream competition channel of the merger, as in the competition-only counterfactual of Section 8, and abstracts from the changes in expected profits and group effects that enter the full merger counterfactual.

The log-likelihood declines at every step of the grid: modestly up to $\kappa = 0.1$, and by 244 and 660 points at $\kappa = 0.5$ and $\kappa = 0.75$. Fit to advances deteriorates in the same direction. The model overpredicts the mean advance by 20% under independence and by 38% at $\kappa = 0.75$.

Panel D shows the two opposing forces described above. As κ rises from 0 to 0.75, the probability that PRH and Simon & Schuster hold the two highest valuations for a jointly contested title rises from 47% to 65%. Combining the two bidders then lowers the advance of these titles by 35% under independence and by 30% at $\kappa = 0.75$. The effect is essentially unchanged up to $\kappa = 0.25$ and falls by five percentage points at the highest correlation considered.

This exercise measures only the loss of upstream competition. It is therefore comparable to the competition-only counterfactual of Section 8, not to the full merger counterfactual, which adds the change in the merged firm's expected profits and group effect. Under independence, the 35% reduction is close to the 36% reduction the competition-only counterfactual predicts for jointly contested fiction titles. The two exercises differ: the counterfactual conditions on the observed advance and winner of the 2021 auctions, whereas Panel D simulates valuations for all auctions in the estimation sample. The upstream competitive effect of the merger is therefore not an artifact of the independence assumption: allowing a substantial common component in publishers' valuations moderates it only slightly.

K First-stage IV Results

This section presents the first-stage results addressing the endogeneity of retail prices. I instrument for price using two sets of instruments: instruments that exploit variation in own and rival product characteristics within genres (Gandhi and Houde, 2019), and input cost indices interacted with page length. In both markets, the F-statistics are large, ruling out weak-instrument concerns. The GH instruments are the primary drivers of the first stage; cost shifters contribute little. This is consistent with books being printed in bulk: the cost of the initial print run largely determines the retail price. What moves prices is mostly competition — entry of new titles into the same genre, greater rival density, or higher-quality competitors.

Table K1: First-Stage Regression: RCNL Demand Estimation

	Fiction		Nonfiction	
	Estimate	(SE)	Estimate	(SE)
A. GH product differentiation				
Age	0.4362	(0.0055)	0.6150	(0.0074)
Page length	0.0067	(0.0021)	−0.0170	(0.0034)
Rating	−0.0155	(0.0017)	−0.0113	(0.0037)
B. Cost instruments				
Paper price × page length	0.0000	(0.0000)	0.0000	(0.0000)
Ink price × page length	−0.0000	(0.0000)	−0.0001	(0.0000)
Publisher wages × page length	0.0001	(0.0001)	0.0002	(0.0001)
C. GH competitor characteristics (within genre)				
No. competitors (N)	0.0004	(0.0013)	0.0060	(0.0040)
Age	−1.3844	(0.3145)	1.1878	(0.7284)
Page length	0.2606	(0.1328)	−0.5130	(0.1206)
Rating	0.0449	(0.0911)	−0.1113	(0.0729)
D. GH interaction terms				
$N\times$ Age	0.0000	(0.0000)	−0.0000	(0.0000)
$N\times$ Page length	−0.0000	(0.0000)	−0.0000	(0.0000)
$N\times$ Rating	0.0000	(0.0000)	−0.0000	(0.0000)
Age × Page length	0.0007	(0.0001)	−0.0005	(0.0002)
Age × Rating	0.0032	(0.0006)	−0.0025	(0.0014)
Page length × Rating	−0.0005	(0.0003)	0.0011	(0.0002)
Product FE	Yes		Yes	
Genre × Market FE	Yes		Yes	
F-statistic (HC1)	451.1		439.2	
Partial R^2	0.0897		0.1236	
Observations	94,070		67,245	
Instruments	16		16	

Notes: Instruments: 13 Gandhi–Houde (GH) instruments (panels A, C, D) and 3 cost shifters (panel B). Panel A characteristics are own-minus-rival averages squared within genre; panel C reports rival averages. Cost instruments (panel B) are input price indices interacted with page length to generate product-level variation. N denotes the number of competing titles in the same genre. Standard errors are HC1 heteroskedasticity-robust.

L Alternative Demand Models

L.1 Setting

This appendix presents demand estimates from a nested logit model with two levels of nesting, in which consumers first choose a format (ebook or print), then a genre within that format, and finally a title. The nesting structure captures the idea that titles sharing both format and genre are closer substitutes than titles sharing only a format, which in turn are closer substitutes than titles from different formats.

Consumer i's indirect utility from purchasing title j in market t is

$$u_{ijt} = \delta_{jt} + \zeta_{ift} + (1 - \sigma_2)\tilde{\zeta}_{igft} + (1 - \sigma_1)\,\varepsilon_{ijt}, \tag{53}$$

where δ_{jt} is mean utility as defined in (3), ζ_{ift} is a consumer-format taste shock common to all titles in format f, $\tilde{\zeta}_{igft}$ is a consumer-genre taste shock common to all titles in genre g within format f, and ε_{ijt} is an i.i.d. idiosyncratic shock. The joint distribution of $(\zeta_{ift}, \tilde{\zeta}_{igft}, \varepsilon_{ijt})$ is generalized extreme value (GEV), which induces two within-group correlations: $\sigma_1 \in [0, 1)$ across titles within the same format-genre cell, and $\sigma_2 \in [0, \sigma_1)$ across titles sharing a format but belonging to different genres. A consistent nesting structure requires $0 < \sigma_2 < \sigma_1 < 1$. When $\sigma_1 = \sigma_2 = 0$ the model collapses to the standard logit. When $\sigma_2 = 0$ it reduces to the one-level format nested logit. As $\sigma_1 \to 1$, titles within the same format-genre cell become perfect substitutes.

Following Berry (1994), the GEV assumption yields closed-form aggregate share equations that can be algebraically inverted to produce the linear estimating equation

$$\ln s_{jt} - \ln s_{0t} = \delta_{jt} + \sigma_1 \ln s_{j|fgt} + \sigma_2 \ln s_{g|ft}, \tag{54}$$

where s_{0t} is the outside-good share, $s_{j|fgt}$ is the share of title j within its format-genre cell, and $s_{g|ft}$ is the share of genre g within format f. Three regressors are endogenous: price p_{jt}, the within-cell share $\ln s_{j|fgt}$, and the genre-within-format share $\ln s_{g|ft}$. I estimate (54) by IV-2SLS, absorbing product and market fixed effects prior to estimation.

For price I use the differentiation instruments of Gandhi and Houde (2019), constructed from the distribution of product characteristics (see Section 5.1) across competing titles. For the within-cell share $\ln s_{j|fgt}$ I use BLP-style instruments at the format-genre level: the number of other titles in the same format-genre cell and differences between the focal title's characteristics and the cell-level average of competing titles' characteristics. For the genre-within-format share $\ln s_{g|ft}$ I instrument with the number of other genres present in the same format and the average characteristics of titles belonging to those competing genres. Together these instruments generate plausibly exogenous variation in price, within-cell crowding, and genre-level crowding within format.

This model follows the idea of Reimers (2019), where they defines nests based on the different formats available for titles (i.e., paperback, hardcover, ebook). In my case, I still group physical versions together (paperback and hardcover) due to absence of multiple versions across my

sample, but keep those separated from electronic versions.[51] In addition to the format nest, I introduce a second level of nesting based on the genre groups defined in Section 5.1, similar to Genakos et al. (2025).

L.2 Results

This model yields own-price elasticities close to those presented in Section 7, though fiction titles are slightly more inelastic. The nesting parameters are large and statistically significant, highlighting close substitution within groups, and satisfy the ordering $1 \geq \sigma_1 \geq \sigma_2 \geq 0$ required for consistency with random utility maximization (McFadden, 1972).

Given that the results are qualitatively similar to those in the main specification, and that the latter allows for richer substitution patterns across book quality, I retain it as my preferred specification. In Appendix R, I present welfare effects from simulating the PRH–SS merger under this model and show that they do not differ significantly from those reported in Section 8.

Table L1: Demand Estimates: Format×Genre Nested Logit

	Fiction		Nonfiction	
	Estimate	(SE)	Estimate	(SE)
A. Mean utility				
Price (α)	−0.160	(0.008)	−0.116	(0.005)
B. Nesting parameters				
Within format×genre (σ_1)	0.698	(0.006)	0.760	(0.006)
Within format (σ_2)	0.641	(0.024)	0.677	(0.014)
Product FE	Yes		Yes	
Market FE	Yes		Yes	
Instruments	27		27	
Median own-price elasticity	−6.13		−7.26	
Observations	126,523		89,249	
Products	6,972		4,328	
Markets	44		44	
Genres (nests)	6		7	
Format×genre cells	12		14	

Notes: Robust standard errors in parentheses. Two-level nested logit estimated via IV-2SLS with product and market fixed effects. The proper 3-level NL requires $0 \leq \sigma_2 \leq \sigma_1 \leq 1$ (satisfied in both markets). The number of observations is larger than the one presented in Section 7 given that here we reduced the grouping of titles across versions, separating electronic and physical ones.

[51] The separation between paperback and hardcover versions can bring an additional issue. Usually titles are first sold in hardcover versions and the paperback version only comes several months after. Then, if the decision to release paperback versions depend on the performance of hardcover ones, we only observe the former when titles are above some quality threshold. The multiplicity of editions will then be correlated with ξ_j, generating an endogeneity problem.

M Demand Rescaling

The raw sales data have two limitations when the objective is to analyze the performance of the book in the US market. The first limitation relates to the fact that I only observe Amazon data, while books are sold in other outlets (bookstores, chains, etc.). To account for this I use the findings in Bondi and Cabral (2021), where authors show that PRH's Amazon sales represent a third of total sales made by the firm. Assuming this is homogeneous across firms, units are rescaled by a factor of three.

The second limitation is that books enter the dataset at different times and are observed for heterogeneous lengths of time. Some titles appear near the beginning of the sample and are followed for almost their entire commercial lifecycle, whereas others enter only toward the end and are observed for a handful of months. Because a substantial share of a book's total sales occurs near release, naively summing observed monthly sales underestimates lifetime sales disproportionately for late-entering titles. A uniform multiplicative scaling factor would therefore introduce systematic bias, inflating sales for well-observed books and undercorrecting for poorly observed ones. To address this issue, I construct book-specific lifecycle adjustment factors based on the empirical evolution of sales over time.

Let τ denote months since publication. Using a subset of "well-observed" books —those that (i) are first observed within the first three months since publication and (ii) are followed for at least 24 months— I estimate the average sales profile:

$$\bar{q}_\tau = \frac{1}{N_\tau} \sum_{j \in \mathcal{J}_{\text{life}}} q_{j\tau} \tag{55}$$

where $q_{j\tau}$ are observed sales for title j at age τ, and $\mathcal{J}_{\text{life}}$ is the set of "well-observed" titles. I then normalize this to get a measure of lifecycle share:

$$s_\tau = \frac{\bar{q}_\tau}{\sum_t \bar{q}_t}, \quad \text{s.t.} \quad \sum_\tau s_\tau = 1 \tag{56}$$

where s_τ measures the expected fraction of a book's lifetime sales that occurs in month τ.

Let $T_j \subseteq \{\tau\}$ denote the set of months in which title j is observed. I then compute the coverage of each title as:

$$c_j = \sum_{\tau \in T_j} s_\tau \tag{57}$$

with $c_j \in (0, 1]$ measuring the fraction of expected lifetime sales that the observed period captures from total sales of title j. Titles that show up late or are observed for a short window will have a small c_j. On the other hand, titles followed over most of their life-cycle have a coverage close to one. In Table M1 I show the distribution of $\frac{1}{c_j}$, the factor that multiplies the sales of each product. On average I rescale fiction and non-fiction titles by a factor of 3.3 and 2.8 respectively.

A natural concern is whether the counterfactual profit effects are driven by titles that require a large adjustment, i.e., books observed for only a short window and therefore assigned a high $\frac{1}{c_j}$. To assess this, Table M2 splits titles in the demand estimation sample into those above and below the segment-specific median of $\frac{1}{c_j}$. Titles above the median are observed over a smaller fraction

Table M1: Distribution of Lifecycle Rescaling Weights ($\frac{1}{c_j}$)

	Fiction	Nonfiction
N products	4,882	3,046
Mean	3.259	2.828
Std. dev.	3.848	3.388
p1	1.000	1.000
p5	1.060	1.008
p10	1.187	1.083
p25	1.491	1.288
p50	2.200	1.913
p75	3.959	3.720
p90	6.095	4.873
p95	8.240	6.380
p99	14.131	12.229
Max	122.926	122.926

Notes: Factor is $1/c_j$ where $c_j = \sum_{\tau \in T_j} s_\tau$ is the sum of the estimated lifecycle profile shares over the ages at which title j is observed in the data. Values are computed at the product (title) level for the demand estimation sample.

of their predicted lifecycle and therefore rely more heavily on the extrapolation. In fiction the median adjustment factor is 2.20 and in nonfiction it is 1.91. If anything, one would expect heavily-extrapolated titles to account for a disproportionately large share of total profits, since a large $\frac{1}{c_j}$ directly scales up their lifetime sales. The table shows the opposite: above-median titles represent exactly half the sample by construction but account for only 46.6% of baseline variable profits in fiction and 39.7% in nonfiction, below their population share in both segments. This indicates that profits are not inflated by the titles with the largest adjustments, with the results driven by the better-observed half of the distribution.

Using this and the rescaling from Amazon sales, I obtain a measure of total sales made by title j as:

$$\tilde{Q}_j = \frac{3Q_j^{obs}}{c_j} \tag{58}$$

where Q_j^{obs} is the observed sales of title j in the data.

The main purpose of obtaining this rescaled measure of sales $\tilde{Q}_j$ is to calculate the expected profits that firms have for each title, which is a key element in defining firms' valuations. This is, when firms consider how much to bid for each title, they consider all future sales they can make, not only the sales from Amazon nor the sales obtained in a couple of months.

Table M2: Baseline Variable Profits by Lifecycle Extrapolation Group

	N titles	% titles	% profits
Fiction (median $\frac{1}{c} = 2.2$)			
Above median	2,441	50.0%	46.6%
Below or equal median	2,441	50.0%	53.4%
Total	4,882	100.0%	100.0%
Nonfiction (median $\frac{1}{c} = 1.9$)			
Above median	1,523	50.0%	39.7%
Below or equal median	1,523	50.0%	60.3%
Total	3,046	100.0%	100.0%
Combined			
Above segment median	3,964	50.0%	43.2%
Below or equal median	3,964	50.0%	56.8%
Total	7,928	100.0%	100.0%

Notes: Each segment uses its own median rescaling factor as the threshold: fiction median $\frac{1}{c} = 2.200$; nonfiction median $\frac{1}{c} = 1.913$. Profits are baseline (pre-merger) variable profits summed across all market-months.A higher value indicates the title is observed over a smaller fraction of its predicted sales lifetime.

N Cost Function

Table N1: Cost Function Estimates

	Fiction	Non-fiction
Dep. variable: $\widehat{mc}_{jt}$ *(\$)*		
Page length	-4.4990***	0.2396
	(0.2553)	(0.3954)
Page length^2	0.0000***	0.0000***
	(0.0000)	(0.0000)
Months since pub.	-0.1954***	-0.3687***
	(0.0063)	(0.0081)
Months since pub.^{2}	0.0045***	0.0077***
	(0.0001)	(0.0001)
Year published	-0.0639	0.7073***
	(0.0677)	(0.0905)
Pages × paper price	-0.0000***	-0.0000***
	(0.0000)	(0.0000)
Coverage	-1.2423***	-1.3314***
	(0.0919)	(0.1391)
Goodreads bestseller author	-216.5888***	1145.5045***
	(72.3081)	(171.7903)
Debut author	0.2733***	0.1352***
	(0.0403)	(0.0430)
Has hardcover	1.3036***	-3.4285***
	(0.1211)	(0.1637)
Share hardcover	4.1951***	5.1850***
	(0.0766)	(0.1159)
Has e-book	-0.2450**	0.3256***
	(0.0998)	(0.1185)
Share e-book	-3.1865***	-1.6032***
	(0.0704)	(0.0844)
Year pub. × pages	0.0022***	-0.0001
	(0.0001)	(0.0002)
GR bestseller × year pub.	0.1075***	-0.5676***
	(0.0358)	(0.0851)
Hardcover × pages	-0.0013***	0.0103***
	(0.0003)	(0.0004)
E-book × pages	0.0010***	-0.0044***
	(0.0002)	(0.0004)
Observations	95,221	67,427
R^2 (within)	0.390	0.357
R^2	0.547	0.588
Publisher × Genre × Year FE	Yes	Yes

Notes: Dependent variable is the marginal cost (in \$) recovered from Nash-Bertrand first-order conditions. Reference genre for pages × genre interactions: Comics & Graphic Novels (Fiction) and Business & Tech (Non-fiction). Standard errors (in parentheses) are iid. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

O Robustness: Number of Bidders

I do not observe the number of bidders in each auction, and with data on winners and transaction prices alone it cannot be identified (Athey and Haile, 2002). The baseline therefore fixes $N_j = 3$ by assumption. Rather than defend a single value, I ask whether the paper's central finding

survives across the plausible alternatives, by re-estimating the auction model and re-simulating the merger under two further assumptions and reporting the effects on advances in Tables O1 and O2. The bidder count enters the likelihood in equation (28), setting how many order statistics generate each observed (winner, price) pair, so each assumption yields its own set of valuation-parameter estimates rather than a different counterfactual layered on the baseline estimates. The demand estimates, which do not depend on N_j, are held fixed throughout.

The first specification sets $N_j = 2$ in every auction, while the second lets the bidder count increase with the size of the advance, sorting auctions into quartiles of the advance distribution and assigning $N_j = 2$ in the lowest quartile, one more bidder per quartile up to $N_j = 5$ in the highest. Both retain the two-channel structure of the baseline. The main message of the exercise is that the number of bidders governs the magnitude of the concentrated harm on jointly contested titles, but not the sign of the aggregate: across every specification I consider, total author compensation does not fall, not even under the assumption most damaging to authors.

That case is $N_j = 2$, reported in Table O1. Here the merger leaves every jointly contested auction with a single remaining bidder. With no second-highest valuation to pin the price, the advance is no longer disciplined by competition, and I set it to \$25,000, the midpoint of the lowest advance category, representing the author's outside option absent any rival bid. This is the most adversarial bidder assumption in the paper, and the harm it inflicts where the merging firms met head-to-head is severe: advances on jointly contested titles fall by 87.5%, against 19.9% in the baseline. Yet the aggregate does not follow. Because jointly contested titles are a small share of the market and the downstream channel keeps lifting bids across the far larger set of auctions the merging parties did not contest (advances there rise by 7.4% to 9.7%), total aggregate advances is essentially flat, edging up by 0.8%. The contrast with a one-sided reading is the point. The decline the DOJ anticipated is what the competition channel alone would produce. Once the downstream channel is allowed to operate, the gains it generates across the untouched majority of the market are just large enough to absorb the entire harm concentrated on jointly contested authors, and the aggregate stops falling.

The second specification makes the number of bidders increase with the observed advance, responding to the idea that competition increases the price paid by firms (U.S. Department of Justice, 2022a). Table O2 reports the results. Two features of the design push the aggregate above even the $N_j = 3$ baseline. First, the titles the merging parties actually contest are concentrated in the upper half of the advance distribution, where N_j is now 4 or 5, so losing one bidder is only a marginal reduction in competition. Second, the severe $N_j = 2$ case is confined to the lowest advance quartile, where PRH and SS are rarely present, so the single-bidder floor almost never binds on a contested title. Together these are strong enough that the competition channel no longer dominates even head-to-head: advances on jointly contested titles rise by 18.0%, rather than falling. Those rise by less than advances elsewhere, 36.8% where one merging party participated and 15.2% where neither did. However, with enough surviving competition, the downstream channel lifts advances across the entire market. Total aggregate advances rise by 26.2%, the largest gain of the three specifications.

Taken together, the three specifications agree on the overall effect. The number of bidders is not identified, and it does move the results: the harm on the jointly contested segment declines

monotonically as bidders are added, from -87.5% under $N_j = 2$ to -19.9% under $N_j = 3$ and to a net 18.0% under the advance-varying schedule. But the sign of the aggregate does not move with it. Aggregate advances are 0.8%, 10.8%, and 26.2% higher across the three cases, non-negative throughout, and turning on the downstream channel rather than the bidder count. Whatever the true number of bidders, the two-sided model does not reproduce the monopsony-only prediction of falling aggregate advances on which the case was brought: in the worst case the downstream channel exactly offsets the concentrated harm, and with more surviving competition it more than offsets it.

Table O1: Merger Counterfactual: Upstream Market (N_j=2)

Panel A: By baseline auction winner			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
PRH	-\$13,945	-4.2	-6.4%
Simon & Schuster	-\$83,857	-9.1	-48.6%
Hachette	\$27,016	2.6	22.5%
HarperCollins	\$18,188	3.7	10.3%
Macmillan	\$27,637	3.6	14.9%
Non-Big5	\$12,956	5.0	11.4%
Overall	\$1,304	1.6	0.8%
Panel B: By merging firms participation in baseline			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
Both PRH & SS	-\$174,280	-14.5	-87.5%
One of PRH/SS	\$18,176	10.6	9.7%
Neither	\$9,763	5.5	7.4%
Overall	\$1,304	1.6	0.8%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

P Clustered Bidders

P.1 Estimation

This appendix presents a robustness variant that changes only how the rival bidders for each title are drawn. In the baseline, the upstream market is segmented by publisher type: a title won by a Big Five publisher faces Big Five rivals, and a title won by any other publisher faces rivals from outside the Big Five. Within the winner's segment, rivals are drawn without replacement with probability proportional to each publisher's share of titles acquired in that segment, so the largest publishers are the most frequent rivals regardless of the title's characteristics. This abstracts from the fact that many publishers, especially smaller ones, are highly specialized: the natural competitors for a graphic-novel acquisition are comics and manga houses, not an arbitrary small press.

Table O2: Merger Counterfactual: Upstream Market (N_j Increasing with Advance)

Panel A: By baseline auction winner			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
PRH	\$18,093	5.5	8.3%
Simon & Schuster	-\$1,024	-0.1	-0.6%
Hachette	\$77,079	7.3	64.3%
HarperCollins	\$89,209	18.4	50.4%
Macmillan	\$70,961	9.2	38.3%
Non-Big5	\$28,735	10.7	27.7%
Overall	\$42,009	51.0	26.2%
Panel B: By merging firms participation in baseline			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
Both PRH & SS	\$51,064	8.8	18.0%
One of PRH/SS	\$63,982	34.0	36.8%
Neither	\$15,914	8.1	15.2%
Overall	\$42,009	51.0	26.2%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

To assess how sensitive the results are to this assumption, I keep the segmentation but change the drawing probabilities. Within the winner's segment, rivals are drawn without replacement with probability proportional to the number of titles each publisher has previously published in the title's tier–genre cell, counting only titles published before the auction. A publisher active in a genre and tier is therefore more likely to be drawn as a competitor for titles in that cell. This mainly changes the composition of rival sets outside the Big Five, where specialization is greatest. Every other component of the model is left unchanged, so any differences from the baseline come only from the composition of the bidder sets. Tables P1 and P2 report the parameter estimates for fiction and non-fiction titles.

Table P1: Structural Auction Model: Estimated Parameters for Fiction Titles (Clustered Bidders)

		Estimate	Std. Error
A. Tier intercepts			
Midlist		0.078	(0.020)
Debut		-0.088	(0.015)
Bestseller		-0.211	(0.065)
B. Firm fixed effects (ref: non-Big5)			
Hachette		0.528	(0.057)
HarperCollins		2.001	(0.045)
Macmillan		0.398	(0.066)
Penguin Random House		0.681	(0.020)
Simon & Schuster		0.130	(0.052)
C. Observable covariates			
β_2: X2=Var (midlist)		-0.415	(0.000)
β_2: X2=Var (debut)		-0.600	(0.014)
β_2: X2=Var (bestseller)		-1.521	(0.131)
D. Signal SDs ($\sigma_{f\tau}$)			
	Midlist ($\sigma_{\text{nb5,mid}}$)	1.393	(0.030)
non-Big5	Debut ($\sigma_{\text{nb5,deb}}$)	1.169	(0.009)
	Bestseller ($\sigma_{\text{nb5,best}}$)	0.731	(0.088)
	Midlist ($\sigma_{\text{Hac,mid}}$)	0.958	(0.008)
Hachette	Debut ($\sigma_{\text{Hac,deb}}$)	1.076	(0.096)
	Bestseller ($\sigma_{\text{Hac,best}}$)	1.035	(0.148)
	Midlist ($\sigma_{\text{HC,mid}}$)	2.132	(0.015)
HarperCollins	Debut ($\sigma_{\text{HC,deb}}$)	1.513	(0.011)
	Bestseller ($\sigma_{\text{HC,best}}$)	1.326	(0.077)
	Midlist ($\sigma_{\text{Mac,mid}}$)	0.913	(0.040)
Macmillan	Debut ($\sigma_{\text{Mac,deb}}$)	0.832	(0.062)
	Bestseller ($\sigma_{\text{Mac,best}}$)	0.499	(0.128)
	Midlist ($\sigma_{\text{PRH,mid}}$)	0.997	(0.000)
Penguin Random House	Debut ($\sigma_{\text{PRH,deb}}$)	0.700	(0.036)
	Bestseller ($\sigma_{\text{PRH,best}}$)	0.585	(0.058)
	Midlist ($\sigma_{\text{S\&S,mid}}$)	0.666	(0.042)
Simon & Schuster	Debut ($\sigma_{\text{S\&S,deb}}$)	0.566	(0.066)
	Bestseller ($\sigma_{\text{S\&S,best}}$)	0.352	(0.043)
Auctions		4,328	
Log-likelihood		-8,077.3	
Pred./actual price ratio		1.075	

Notes: Standard errors from the numerical Hessian (observed Fisher information matrix).

Table P2: Structural Auction Model: Estimated Parameters for Non-fiction Titles (Clustered Bidders)

		Estimate	Std. Error
A. Tier intercepts			
Midlist		0.764	(0.078)
Debut		0.458	(0.072)
Bestseller		0.425	(0.090)
B. Firm fixed effects (ref: non-Big5)			
Hachette		0.137	(0.084)
HarperCollins		0.660	(0.152)
Macmillan		-0.129	(0.051)
Penguin Random House		2.070	(0.075)
Simon & Schuster		0.081	(0.097)
C. Observable covariates			
β_2: X2=Var (midlist)		-1.009	(0.000)
β_2: X2=Var (debut)		-1.237	(0.004)
β_2: X2=Var (bestseller)		-0.839	(0.005)
D. Signal SDs ($\sigma_{f\tau}$)			
non-Big5	Midlist ($\sigma_{\text{nb5,mid}}$)	2.192	(0.051)
	Debut ($\sigma_{\text{nb5,deb}}$)	1.277	(0.011)
	Bestseller ($\sigma_{\text{nb5,best}}$)	0.577	(0.069)
Hachette	Midlist ($\sigma_{\text{Hac,mid}}$)	0.687	(0.047)
	Debut ($\sigma_{\text{Hac,deb}}$)	0.653	(0.050)
	Bestseller ($\sigma_{\text{Hac,best}}$)	0.999	(0.143)
HarperCollins	Midlist ($\sigma_{\text{HC,mid}}$)	1.257	(0.075)
	Debut ($\sigma_{\text{HC,deb}}$)	0.836	(0.079)
	Bestseller ($\sigma_{\text{HC,best}}$)	1.623	(0.223)
Macmillan	Midlist ($\sigma_{\text{Mac,mid}}$)	0.646	(0.008)
	Debut ($\sigma_{\text{Mac,deb}}$)	0.336	(0.040)
	Bestseller ($\sigma_{\text{Mac,best}}$)	0.670	(0.008)
Penguin Random House	Midlist ($\sigma_{\text{PRH,mid}}$)	1.499	(0.000)
	Debut ($\sigma_{\text{PRH,deb}}$)	1.255	(0.052)
	Bestseller ($\sigma_{\text{PRH,best}}$)	1.328	(0.101)
Simon & Schuster	Midlist ($\sigma_{\text{S\&S,mid}}$)	0.727	(0.059)
	Debut ($\sigma_{\text{S\&S,deb}}$)	0.547	(0.066)
	Bestseller ($\sigma_{\text{S\&S,best}}$)	2.033	(0.395)
Auctions		2,793	
Log-likelihood		-5,345.3	
Pred./actual price ratio		1.077	

Notes: Standard errors from the numerical Hessian (observed Fisher information matrix).

P.2 Counterfactual: Clustered Bidders

Using the estimates from the previous subsection, I recompute the merger counterfactual. The counterfactual is constructed exactly as in the main text: I simulate the Penguin Random House–Simon & Schuster merger, letting the merged firm internalize both the competition

between the two formerly rival bidders and the cannibalization across their combined catalog, and I aggregate the implied change in author advances across the titles in the sample. The only difference from the baseline counterfactual is that each title's set of potential bidders is now the tier–genre-weighted clustered draw described above, so the competitive pressure a title faces reflects the publishers actually active in its tier and genre rather than a size-stratified random draw. This is where the alternative selection rule can matter most: for a title in a niche genre, the relevant competitors are the specialized presses that publish in it, and whether such firms are present shapes how much the merger relaxes competition for that title. Table P3 reports the resulting change in advances pooled across fiction and non-fiction. I report only the change in advances, since the downstream (demand-side) outcomes are essentially unchanged relative to the baseline specification.

For titles jointly contested by PRH and Simon & Schuster, advances still fall, but by less than in the baseline: -7.9%, against -19.9% in Section 8. For the remaining titles, the increase is much larger: 30.3% where one merging firm participated and 16.2% where neither did, against 15.6% and 11.3% in the baseline. Most of the additional gain is paid by the non-merging publishers. It reflects both the composition of the rival sets, which now favor publishers active in each title's tier and genre, and the re-estimated valuation parameters. Aggregate advances rise by 21.1%, about twice the baseline increase of 10.2%.

Despite these differences in magnitude, the sign and the main conclusion are unchanged: jointly contested titles see a decrease in advances, while the rest benefit from the merger. The downstream channel still dominates the competition channel.

Table P3: Merger counterfactual (clustered bidders) (N_j=3)

Panel A: By baseline auction winner			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
PRH	\$8,457	2.6	4.0%
Simon & Schuster	-\$3,484	-0.4	-2.0%
Hachette	\$81,279	7.6	66.7%
HarperCollins	\$65,401	13.5	37.0%
Macmillan	\$59,735	7.7	32.6%
Non-Big5	\$29,400	11.4	25.4%
Overall	\$34,419	42.4	21.1%
Panel B: By merging firms participation in baseline			
	Mean Δa_{fj} (\$)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ (\$M)	Total Δa_{fj} (%)
Both PRH & SS	-\$14,611	-1.9	-7.9%
One of PRH/SS	\$58,421	33.7	30.3%
Neither	\$20,280	10.7	16.2%
Overall	\$34,419	42.4	21.1%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

Q Robustness: Counterfactuals

Q.1 Counterfactual: Isolated Upstream Markets

This appendix reports the upstream-only counterpart to the robustness exercise of Appendix O. It uses the same re-estimated auction models and simulates the same merger, but shuts the downstream channel off: valuations are held at their baseline values, and the merger's only effect is to replace PRH and SS with a single bidder wherever both participated. This is the isolation that underlies the validation of Section 7.3, and it mirrors the methodology the DOJ's experts applied in the PRH–SS case, in which the input market is evaluated on its own (U.S. Department of Justice, 2022a). Section 7.3 reports the baseline $N_j = 3$ version (Table 10). Here I extend it to $N_j = 2$ and to the advance-varying bidder count, in Tables Q1 and Q2.

By construction the effect is confined to the jointly contested segment: with valuations frozen, advances move only on titles both PRH and SS contested, and every other auction is exactly unchanged. Within the contested set the harm is large and, as in Appendix O, grows as bidders are removed: advances fall by 15.0% under the advance-varying schedule, by 39.0% at $N_j = 3$ (Section 7.3), and by 86.3% at $N_j = 2$. The last case is handled exactly as in Appendix O: under $N_j = 2$ the merger reduces every jointly contested auction to a single remaining bidder, and with no second-highest valuation to set the price the advance is fixed at the lowest-tier floor of \$25,000, the author's outside option. Because the floor binds regardless of the surviving bidder's valuation, the contested decline is nearly the same as in the full model (-86.3% here, -87.5% in Table O1): where the auction collapses to one bidder, the downstream channel cannot move the price. Diluted over each firm's full portfolio (Panel A), the losses fall on PRH (-6.8% to -12.4%) and, more heavily, on SS (-15.4% to -30.5%). The asymmetry of SS being harmed more than twice as much as PRH (because PRH is disproportionately SS's runner-up) persists across every bidder assumption, consistent with the validation against the DOJ's bid-log estimates. In the aggregate, advances fall in every case: by 3.8%, 4.3%, and 7.0% under the advance-varying schedule, $N_j = 3$, and $N_j = 2$ respectively.

This is the monopsony-only prediction, and it is robust: evaluated in isolation, the upstream market delivers an unambiguous reduction in author compensation whatever the number of bidders. The comparison with Appendix O is the paper's argument in miniature. The two appendices use the same estimated model, the same bidder-count specifications, and the same merger. They differ only in whether the downstream channel is allowed to operate. With it switched off, aggregate advances fall by 3.8% to 7.0%. With it switched on, they rise by 0.8% to 26.2% (Tables O1, 11 and O2). The harm the DOJ anticipated is not wrong within its own frame, but it is incomplete, and what it omits is precisely the downstream channel that reverses the sign of the merger's effect on authors as a whole.

Table Q1: Upstream-Only Merger Counterfactual (N_j=2)

Panel A: By baseline auction winner			
	Mean Δa_{fj} ($)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ ($M)	Total Δa_{fj} (%)
PRH	-$26,608	-8.2	-12.4%
Simon & Schuster	-$52,681	-5.7	-30.5%
Hachette	$0	0.0	0.0%
HarperCollins	$0	0.0	0.0%
Macmillan	$0	0.0	0.0%
Non-Big5	$0	0.0	0.0%
Overall	-$11,319	-13.9	-7.0%
Panel B: By merging firms' participation in baseline			
	Mean Δa_{fj} ($)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ ($M)	Total Δa_{fj} (%)
Both PRH & SS	-$158,085	-13.9	-86.3%
One of PRH/SS	$0	0.0	0.0%
Neither	$0	0.0	0.0%
Overall	-$11,319	-13.9	-7.0%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

Table Q2: Upstream-Only Merger Counterfactual (N_j Increasing in the Advance)

Panel A: By baseline auction winner			
	Mean Δa_{fj} ($)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ ($M)	Total Δa_{fj} (%)
PRH	-$14,530	-4.5	-6.8%
Simon & Schuster	-$26,556	-2.9	-15.4%
Hachette	$0	0.0	0.0%
HarperCollins	$0	0.0	0.0%
Macmillan	$0	0.0	0.0%
Non-Big5	$0	0.0	0.0%
Overall	-$6,042	-7.3	-3.8%
Panel B: By merging firms participation in baseline			
	Mean Δa_{fj} ($)	Total $\Delta \sum_{j \in \mathcal{J}_f} a_{fj}$ ($M)	Total Δa_{fj} (%)
Both PRH & SS	-$40,489	-7.3	-15.0%
One of PRH/SS	$0	0.0	0.0%
Neither	$0	0.0	0.0%
Overall	-$6,042	-7.3	-3.8%

Notes: Panel A groups auctions by the baseline winner. Panel B groups by whether Penguin Random House, Simon & Schuster, both, or neither participated in the baseline auction. Total Δa_{fj} (%) is the aggregate change: total price change over total baseline advance.

R Counterfactual Effects Under Alternative Demand

In this Appendix I present welfare effects from simulating the PRH–SS merger using the alternative demand model presented in Appendix L. In terms of firm profits, the overall change is around $20 million, close to the $17 million obtained with the main specification in Section 8. The relative changes differ mostly due to two elements. First, demand is more inelastic for fiction, which implies lower marginal costs for those products and enlarges baseline profits. Second, the model treats electronic and physical books as separate products, and ebooks carry lower marginal costs. What matters for the main argument of the paper is the total increase in profits, since this is what is ultimately transferred to authors.

Regarding consumer surplus, the change under the NL model is −1.5%, smaller in magnitude than the −2.5% found in the main specification.[52] Combined with the profit increase, this yields a total welfare change of −0.4%, compared to −0.3% in the main model.

Under this model the main message of the paper is maintained: the merger leads to a transfer from consumers to firms, with firms obtaining higher profits through higher prices. Hence, the main results of the paper are not sensitive to the specification of the demand side.

Table R1: PRH–Simon & Schuster Merger: Welfare Effects (Format×Genre Nested Logit)

A. Firm Profit Changes		
	Total $\Delta\pi$ ($M)	$\Delta\pi$ (%)
PRH + Simon & Schuster	5.38	1.20%
HarperCollins	4.41	4.10%
Hachette	2.28	4.16%
Macmillan	2.21	4.02%
Other publishers	5.93	4.54%
All firms	20.19	2.54%
B. Consumer Surplus Change		
	ΔCS ($M)	ΔCS (%)
Fiction	−16.48	−1.55%
Nonfiction	−16.90	−1.38%
Combined	−33.38	−1.46%
C. Total Welfare Change		
	ΔW ($M)	ΔW (%)
Fiction	−6.64	−0.44%
Nonfiction	−6.55	−0.41%
Combined	−13.19	−0.43%

Notes: Results from Nash-Bertrand merger counterfactual using the two-level Nested Logit specification from L. Consumer surplus is computed via the NL logsum formula and scaled to full-market, full-lifetime units using the same sales-weighted factor applied to profits (lifecycle coverage weight divided by Amazon market share). Total welfare $\Delta W = \Delta\text{CS} + \Delta\pi$ (unweighted sum).

[52]The absolute change in CS is larger, though smaller in relative terms. CS is derived from the utility consumers obtain from books, and different models capture this differently. For instance, the model in Section 5 allows for heterogeneous preferences, which alters how utility is computed. Additionally, the NL specification treats ebooks and print as separate products, which enlarges the choice set and raises baseline utility, since consumers benefit from having more options.