

NET Institute*

www.NETinst.org

Working Paper #18-14

October 2018

**Agency Pricing and Bargaining:
Evidence from the E-Book Market**

Babur De los Santos
Clemson University

Daniel P. O'Brien
Compass Lexecon

Matthijs R. Wildenbeest
Indiana University, Kelley School of Business

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

AGENCY PRICING AND BARGAINING: EVIDENCE FROM THE E-BOOK MARKET*

Babur De los Santos[†]

Daniel P. O'Brien[‡]

Matthijs R. Wildenbeest[§]

October 2018

Abstract

This paper examines the relationship between agency contracts and retail prices when vertical contracts are determined through bargaining. Our model shows that whether agency contracts lead to higher or lower retail prices than traditional wholesale contracts depends on the distribution of bargaining power across upstream and downstream firms. We propose a methodology to structurally estimate a demand and supply model that allows for both vertical contracting models, and uses the Nash-in-Nash bargaining solution to capture competition between upstream and downstream firms. We estimate our model using data from the market for e-books, which is an industry that has experienced several transitions between agency and wholesale contracts. We analyze the latest switch from wholesale to agency after the expiration of a two-year ban on agency pricing following a lawsuit by the US Department of Justice against major publishers in the industry. This ban allowed us to observe the results of a switch to agency contracts after a rarely seen widespread restart of firm's bilateral bargaining over vertical contracts. Using a unique dataset of e-book prices both before and after this recent change in selling method, we show that prices went up substantially at Amazon following the shift but remained relatively flat at Barnes & Noble. Structural estimates of the bargaining model show that our bargaining model gives a better fit to the data than a model with take-it-or-leave-it input contracts. Counterfactual simulations indicate that reinstatement of most favored nation clauses, which were banned in 2012 for a period of five years, would lead to price increases of close to nine percent for non-fiction books.

Keywords: e-books, agency agreements, vertical restraints, bargaining

JEL Classification: C14, D83, L13

*We thank Ginger Jin, Huaxia Rui, and Takanori Adachi for their useful comments and suggestions. This paper has benefitted from presentations at the University of Mannheim, Berlin IO Day at Humboldt University, The Ohio State University, Einaudi Institute for Economics and Finance in Rome, European Commission (DG COMP), Wake Forest, 2017 CES North America Conference, Tuscaloosa AL, 2017 IIOC meeting in Boston, 2017 HAL White Antitrust Conference in Washington DC, and 2017 EARIE Conference in Maastricht. We gratefully acknowledge financial support by the NET Institute, <http://www.NETinst.org>.

[†]Clemson University, E-mail: babur@clemson.edu.

[‡]Compass Lexecon, E-mail: dobrien@compasslexecon.com.

[§]Indiana University, Kelley School of Business, E-mail: mwildenb@indiana.edu.

1 Introduction

In most retail markets, suppliers and retailers use what is commonly known as the *wholesale* model, in which a supplier sells a product to a retailer at a per-unit wholesale price, who then resells the good to a consumer for a price that is set by the retailer. Another vertical contractual approach is the *agency* model, in which retail prices are determined directly by the upstream supplier and the downstream retailer receives a royalty or sales commission for every unit sold. Although agency agreements have been used in some conventional markets (e.g., newspapers sold at kiosks), they are especially prevalent in online markets. An example is Amazon Marketplace in which the Marketplace seller (the upstream firm) sets the price and Amazon (the downstream firm) gets a percentage of the revenue for every item sold.¹ An important question is which of the two selling methods—the wholesale or the agency model—will lead to higher retail prices.

Recent theoretical papers that compare pricing under agency and wholesale models provide conflicting answers. For instance, Johnson (2018) finds that even though the agency model raises prices initially, prices are lower in the long run when wholesale contracts are used. Johnson’s model focuses on consumer lock-in in two-sided markets: retailers in the wholesale model find it optimal to set low price initially to attract consumers, but once sufficiently many consumers have adopted the platform, they find it optimal to raise prices. If upstream firms deliver to multiple platforms, they initially do not face such an incentive in the agency model. In contrast, Foros, Kind, and Shaffer (2017) find that if competition among retailers is stronger than among manufacturers, upstream firms will set higher prices in the agency model than retailers would set in the wholesale model. Both these theoretical models, as well as other papers in the literature on agency versus wholesale pricing, assume take-it-or-leave-it contracts, which implies the upstream firm has all the bargaining power in the wholesale model and the downstream firm has all the bargaining power in the agency model.² However, in many markets bargaining plays an important role in determining the outcome of vertical contracts. Bargaining between upstream and downstream firms in the context of the wholesale model has been extensively studied in the theoretical literature (see, for instance, Horn and Wolinsky, 1988; Dobson and Waterson, 1997) and is considered to be a fundamental economic factor that determines how firms divide the surplus of an economic relationship.

¹Other examples include eBay Buy It Now, the Apple App Store, and the e-book market.

²Other papers in this literature include Gaudin and White (2014) and Johnson (2017). Gaudin and White (2014) focus on the market for reading devices that are complementary to e-books and show that prices are higher in the wholesale model than in the agency model if demand satisfies Marshall’s Second Law of demand and e-books can be read without having to use an essential reading device such as the Kindle for Kindle books, as is the case since 2010 when Amazon developed an app for the iPad. Johnson (2017) focuses on most-favored nation clauses in agency contracts and finds that revenue-sharing contracts typically result in lower retail prices.

In this paper we examine the relationship between agency contracts and retail prices when vertical contracts are determined through bargaining. In Section 2 we extend the bilateral monopoly model in Johnson (2017) to allow for bilateral bargaining between an upstream firm and a downstream retailer. Johnson’s model allows for both wholesale and agency contracts, but assumes input contracts are take-it-or-leave-it, which is a special case of our model in which the upstream firm has all the bargaining power in the wholesale model and the downstream firm has all the bargaining power in the agency model. We show that whether agency or wholesale contracts result in higher retail prices depends on the relative bargaining power of the upstream and downstream firm. For instance, if the downstream firm has most of the bargaining power and prices are set according to the wholesale model, wholesale prices will be low, and retail prices will be low as well. However, in the agency model, when the downstream firm has most of the bargaining power this firm will negotiate a relatively large royalty, and to maximize profits, the upstream firm has to set a relatively high price as well. The opposite is true when the upstream firm has most of the bargaining power: wholesale prices will be relatively high in the wholesale model, while royalties will be low in the agency model, resulting in lower retail prices under agency than under wholesale contracts. Note that the relationship between bargaining power and retail prices under wholesale and agency pricing plays an important role in our identification strategy when estimating bargaining power in our supply side model, as we will explain in more detail below.

In Section 3 we adapt the theoretical model to make it more amenable for estimation. To capture that consumers may be more likely to substitute among competing products within the same retailer than across retailers, we use a nested logit demand structure. We also allow for competition among both upstreams firms and downstream firms. Following Crawford and Yurukoglu (2012) we use the Nash-in-Nash bargaining solution of Horn and Wolinsky (1988). In this framework, each pair of firms reaches the Nash bargaining solution while taking the Nash solutions of the other pairs as given; the Nash equilibrium is then the Nash solution for each pair such that no firm has an incentive to deviate. This approach fits into a broader literature that estimates bilateral bargaining models in oligopolistic markets (Draganska, Klapper, and Villas-Boas, 2010; Crawford and Yurukoglu, 2012; Grennan, 2013; Gowrisankaran, Nevo, and Town, 2015).³ We extend this literature, which has focused on wholesale pricing contracts, to allow for agency contracts between

³Draganska, Klapper, and Villas-Boas (2010) use the Nash bargaining solution to estimate upstream and downstream margins in the German retail market for coffee. Crawford and Yurukoglu (2012) estimate an input-market bargaining model for cable television and use the estimate to calculate the welfare effects of bundling. Grennan (2013) focuses on bargaining and price discrimination in the market for medical devices. Gowrisankaran, Nevo, and Town (2015) study the effects of hospital mergers using the estimates from a model in which hospitals bargain with managed care organizations.

upstream and downstream firms.

We apply our model to the market for e-books. This industry is uniquely suited to study the effects of bargaining under wholesale and agency contracts because the industry has experienced various transitions between these selling methods since the introduction of the Kindle e-reader in 2007. In Section 4 we describe the changes in contracts between publishers and book retailers in the e-book industry that happened over time and how these changes have affected retail prices. E-books, like printed books, initially were sold using the traditional wholesale model where publishers sell the e-books to retailers for 30 to 50 percent of the list price and booksellers set retail prices. In this period, Amazon was committed to a pricing strategy of \$9.99 for e-books, particularly for new releases. As De los Santos and Wildenbeest (2017) document, publishers were against this pricing scheme as they believed that it cannibalized hardcover sales and eroded consumers' perception of the value of a book which would eventually lead to lower wholesale prices.⁴

With the introduction of the iPad in 2010, major publishers negotiated agency contracts with Apple, which allowed them to sell e-books in Apple's new iBookstore. The conditions of the agency contracts with Apple, particularly the most favored nation (MFN) clause which required publishers to match lower retail prices at other retailers, prompted five of the major publishers known as the Big Six to collude to compel the adoption of agency contracts on Amazon. As a result of the adoption of agency contracts in the industry, publishers imposed higher retail prices, and in 2012 the Department of Justice sued Apple and five of the Big Six publishers for conspiring to raise e-book prices. All of the five publishers settled the lawsuit and agreed to a two-year ban on publisher-set prices which effectively meant the return to traditional wholesale contracts. De los Santos and Wildenbeest (2017) analyze the transition from agency to wholesale contracts following the ban and find that retail prices decreased by 18 percent at Amazon and 8 percent at Barnes & Noble as a result of the switch.

The expiration of the two-year ban on agency pricing meant that by the end of 2014, publishers could negotiate agency contracts to control retail prices directly. Bargaining between publishers and retailers played an important role in the renegotiation of existing contracts. In Section 4 we describe some aspects of the bargaining dispute between Amazon and Hachette, which included inventory reductions and price increases of Hachette titles. These negotiations took over six months, and were extensively covered by the media and involved public pressure by some of Hachette's bestseller authors. Despite the lengthy bargaining period, by the end of 2015 all of the major publishers

⁴Related work by Li (2015) estimates a structural model in which consumers choose how many books to buy, which format, and on what platform, and finds that about half of e-book sales come from cannibalization of print book sales. Reimers and Waldfoegel (2017) find that e-book are priced below *static* profit maximizing levels.

had returned to agency contracts with publisher-set prices. In this section we also investigate the effect on retail prices following this latest shift towards agency contracts. Since the new agency contracts between publishers and retailers were implemented at different points in time, we exploit this variation to estimate the change in retail prices as a result of the switch using a difference-in-differences approach. Our findings indicate that prices went up on average 16 percent on Amazon whereas prices at Barnes & Noble went down by an average of 1.5 percent. Our estimates also show that there is substantial heterogeneity in the price effect across publishers. Although these findings are difficult to explain using most existing theoretical models, they are consistent with the retailers in our model having different bargaining power in their negotiations with the publishers.⁵

In Section 5 we discuss how to structurally estimate the empirical model developed in Section 3 and present demand and supply side estimates of the bargaining model. The extent to which prices change following a change in selling method is related to the relative bargaining power of the firms involved. To fully exploit this mechanism for identification, we use data from both before and after the latest switch to agency to structurally estimate the model. The estimation of the model consists of two stages. In the first stage we estimate the demand side using price and sales rank data. Estimation of the demand side of the model does not depend on the selling method, so we can combine data from both periods to obtain the demand side parameters. In the second stage, we estimate the parameters of the supply side model, taking the demand side estimates as given. The nested logit demand structure allow us to account for the fact that consumers lock-in to an e-book reader platform, and thus consumers are more likely to substitute within platform (e.g., Amazon's Kindle device or Kindle app or Barnes & Noble's Nook) than across platforms. The objective is to estimate a relative bargaining power parameter and agency royalty parameter for each publisher-retailer pair, as well as the parameters of a marginal cost equation. Although the supply side model changes when moving from wholesale contract to agency contracts, we assume the relative bargaining power parameters do not change when switching, allowing us to connect the two periods. Specifically, for a given set of agency royalty parameters we can use the pricing and bargaining first-order conditions in the agency model to solve for the bargaining power parameters, which we can use to derive the margins of the upstream firms in the wholesale period. The margins for the upstream and downstream firms are functions of these royalties as well as wholesale prices, and can be solved for marginal costs. By relating the marginal cost estimate for the upstream firms

⁵An important difference with the analysis in De los Santos and Wildenbeest (2017) is that our results compare prices under wholesale to prices under agency without collusion playing a role, whereas De los Santos and Wildenbeest analyze changes in retail prices after the switch to agency that resulted from a collusive agreement of competing publishers in the industry.

to marginal cost observables, the goal of the estimation procedure is then to find the marginal cost and bargaining parameters that minimize the sum of squared residuals of the marginal cost equation for the upstream firms.

The estimates indicate that demand is relatively elastic—the estimated price coefficient for the nested logit specification implies a median own-price elasticity of -2.7 . The supply-side estimates point out that the bargaining power is close to equally distributed between publishers and retailers on average, although there are substantial differences between the different retailer-publisher pairs. The estimated bargaining power parameters imply an agency royalty of 30.7 percent on average, which is slightly higher than the thirty percent that was commonly used during the first agency period. Furthermore, we compare the fit of the bargaining model to estimates of alternative vertical models. In particular, we estimate the model assuming take-it-or-leave-it contracts (which is the standard assumption in the theoretical literature on agency versus wholesale), in which the publishers set the wholesale price in the wholesale model and the retailers determine the agency royalty in the agency model, as well as models in which either the publishers or the retailers have all the bargaining power. We find that the bargaining model gives a better fit to the data than these alternative specifications.

In Section 6 we also discuss the results of a counterfactual analysis in which we use the estimates of the bargaining model to simulate the effect of MFN clauses on retail prices. MFN clauses guarantee that the same title is sold at the same price everywhere and were part of the contracts that were used during the first agency period in the e-book industry. The settlements between the U.S. Department of Justice (DOJ) and the publishers banned the use of MFN clauses for a period of five years as they were considered to have played a crucial role in the industry initial shift towards agency agreements and the ensuing conspiracy to increase prices. The role of MFN has been explored by Johnson (2017) who finds that price-parity restrictions raise retail prices. In line with these theoretical findings, our counterfactual simulations indicate that prices will increase an additional three percent on average if the agency contracts are replaced with contracts that include MFN clauses, with higher price increases for non-fiction books than fiction books.

2 Vertical Bargaining Model

In this section we extend the bilateral monopoly model of Johnson (2017) to allow for bargaining over input prices. Let there be two firms, an upstream firm U and a downstream firm D , that produce and sell a product to consumers for a price p . Consumer demand is given by a strictly

decreasing demand function $Q(p)$. Marginal cost is $c^U > 0$ for the upstream firm and $c^D > 0$ for the downstream firm. We consider two pricing structures: (1) the “wholesale model,” where firms first agree to a wholesale contract for the product and the downstream firm sets the retail price; and (2) the “agency model,” where firms first agree to a royalty contract for the product and then the upstream firm sets the retail price.

2.1 Wholesale Pricing

In the wholesale model, the downstream firm sets the price, whereas we assume the upstream and downstream firm bargain over the wholesale price w . First consider the downstream firm D . This firm takes the wholesale price as given when determining the price p that maximizes

$$\pi^D = (p - w - c^D)Q(p).$$

The first-order condition is given by

$$p - w - c^D = \phi(p), \tag{1}$$

where $\phi(p)$ is a measure of the sensitivity of demand:

$$\phi(p) = -\frac{Q(p)}{Q'(p)}.$$

Following Johnson (2017), we assume $\phi'(p) < 1$ to ensure the profit function is well behaved.⁶

Profits of the upstream firm are given by

$$\pi^U = (w - c^U)Q(p).$$

The wholesale price w is determined through Nash bargaining (Nash, 1950) between the upstream and downstream firm. Specifically, assuming zero disagreement payoff, the (asymmetric) Nash bargaining solution is determined by maximizing

$$(\pi^U)^\lambda (\pi^D)^{1-\lambda},$$

⁶The sign of $\phi'(p)$ determines whether the demand function is log-concave ($\phi'(p) < 0$), log-convex ($\phi'(p) > 0$), or log-linear ($\phi'(p) = 0$).

where $\lambda \in [0, 1]$ is a bargaining parameter, which takes on value 0 if the downstream firm has all the bargaining power and 1 if the upstream firm has all the bargaining power (which corresponds to the take-it-or-leave-it case). If $\lambda = 0.5$ then the bargaining power is evenly distributed between the upstream firm and the downstream firm. Taking the first-order condition and rearranging gives

$$\lambda \pi^D \pi'^U + (1 - \lambda) \pi^U \pi'^D = 0. \quad (2)$$

Following Johnson (2017), rewriting the retail pricing first-order condition in equation (1) as $w = p - c^D - \phi(p)$ is useful because it allows us to write the profits of the firms in terms of prices only, i.e., $\pi^U = (p - \phi(p) - c^U - c^D)Q(p)$ and $\pi^D = \phi(p)Q(p)$. Using this to replace the profits in the bargaining first-order condition in equation (2) gives

$$\begin{aligned} & \lambda \phi(p)Q(p) [(1 - \phi'(p))Q(p) + (p - \phi(p) - c^U - c^D)Q'(p)] \\ & + (1 - \lambda) (p - \phi(p) - c^U - c^D)Q(p) [\phi'(p)Q(p) + \phi(p)Q'(p)] = 0. \end{aligned}$$

Dividing both sides by $Q'(p)$ and rearranging gives

$$(p - \phi(p) - c^U - c^D) \left[1 + \left(\frac{1 - \lambda}{\lambda} \right) \frac{\phi(p)(1 - \phi'(p))}{\phi(p)} \right] = \phi(p)(1 - \phi'(p)),$$

which can be further rearranged to obtain an expression for the markup as a function of $\phi(p)$, $\phi'(p)$, and λ :

$$p - c^U - c^D = \phi(p) \left(1 + \frac{\lambda(1 - \phi'(p))}{\lambda + (1 - \lambda)(1 - \phi'(p))} \right). \quad (3)$$

2.2 Agency Pricing

In the agency model, the upstream firm sets the price, whereas the two firms bargain over the agency royalty r . First consider the upstream firm U . This firm takes the royalty as given when determining the price p that maximizes

$$\pi^U = ((1 - r)p - c^U) Q(p).$$

The first-order condition is given by

$$(1 - r)p - c^U = (1 - r)\phi(p). \quad (4)$$

Profits of the downstream firm are given by

$$\pi^D = (rp - c^D)Q(p).$$

We can rewrite the retail pricing first-order condition in equation (4) as

$$r = 1 - \frac{c^U}{p - \phi(p)}.$$

Using this to write the profits in terms of prices only gives $\pi^U = \frac{c^U}{p - \phi(p)}\phi(p)Q(p)$ for the upstream firm and $\pi^D = \left(p - \phi(p)\frac{c^U}{p - \phi(p)} - c^U - c^D\right)Q(p)$ for the downstream firm. The derivative of the upstream profits is

$$\pi'^U = \frac{c^U [\phi'(p)Q(p) + \phi(p)Q'(p)] (p - \phi(p)) - c^U \phi(p)Q(p)(1 - \phi'(p))}{(p - \phi(p))^2},$$

whereas the derivative of the downstream profits is

$$\pi'^D = Q(p) + (p - c^U - c^D)Q'(p) - \pi'^U.$$

Using this in the bargaining first-order condition (see equation (2)) gives

$$\lambda \left[\left(p - \phi(p)\frac{c^U}{p - \phi(p)} - c^U - c^D \right) Q(p) \right] \pi'^U + (1 - \lambda) \left(\frac{c^U}{p - \phi(p)} \phi(p)Q(p) \right) \times [Q(p) + (p - c^U - c^D)Q'(p) - \pi'^U] = 0.$$

Notice that

$$\pi'^U/Q'(p) = \phi(p) \left[\frac{c^U p(1 - \phi'(p))}{(p - \phi(p))^2} \right],$$

so dividing both sides of the bargaining first-order condition by $Q'(p)$ and rearranging gives

$$p - c^U - c^D = \phi(p) \left(1 - \lambda + c^U \frac{p(1 - \phi'(p))}{(p - \phi(p))^2} \right) \cdot \frac{1}{1 - \lambda + \lambda \cdot \frac{p(1 - \phi'(p))}{p - \phi(p)}}. \quad (5)$$

2.3 Comparison of Selling Models

Proposition 1 shows that it depends on the relative bargaining power of the two firms whether prices are higher or lower under agency in comparison to wholesale pricing.

Proposition 1 *Retail prices are higher under agency than under wholesale pricing if the down-*

stream firm has all the bargaining power and lower if the upstream firm has all the bargaining power.

The proof for this proposition can be found in Appendix A. Figures 1 and 2 illustrate the results stated in Proposition 1 for the demand function $Q(p) = p^{-1/\kappa}$. For this constant-elasticity demand function the equilibrium price in the wholesale model is $p^w = (c^U + c^D)(1 + \sigma(\lambda - 1))/(\kappa - 1)^2$ and $p^a = 2(c^U + c^D(1 - \kappa))/((1 - \kappa) \cdot (1 + \kappa(\lambda - 1)))$ in the agency model. In Figure 1(a) we set $\kappa = 0.5$ and $c^U = c^D = 0.1$ and plot the equilibrium price as a function of the bargaining power parameter λ .

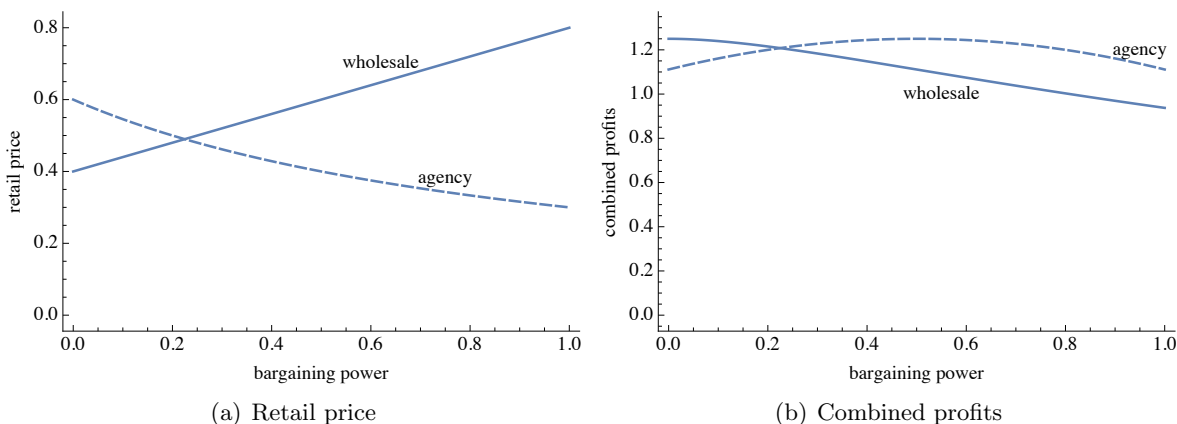


Figure 1: Retail prices and combined profit as a function of bargaining power

Retail prices are increasing in λ in the wholesale model—the more bargaining power the upstream firm has, the higher the wholesale price it will set, with higher retail prices as a result. On the other hand, retail prices are decreasing in λ in the agency model, because the upstream firm gets to keep more of the retail price when it has more bargaining power, which makes it optimal for the upstream firm to set lower retail prices. Figure 1(a) also illustrates that it depends on the exact value of the bargaining parameter whether retail prices are higher or lower under agency. In this example, prices are higher under agency than under wholesale for bargaining power parameters that are less than 0.23 and lower otherwise. Also note that in the case of take-it-or-leave-it offers, which corresponds to $\lambda = 1$ for the wholesale model and $\lambda = 0$ for the agency model, prices under wholesale are higher than prices under agency. This result is consistent with Lemma 2 of Johnson (2017), which gives conditions under which the retail price under agency pricing is strictly less than under wholesale pricing.

Figure 1(b) gives the combined profits of the upstream and downstream firm as a function of the bargaining power parameter for each of the two selling models. For this particular example, the

joint firm profits are maximized under agency pricing when the firms share equal bargaining power. However, under wholesale pricing joint profits are maximized when the downstream firm has all the bargaining power. The latter happens because when the downstream firm has all the bargaining power it will demand a wholesale price that equals the marginal cost of the upstream firm, which will eliminate the double marginalization problem and thus maximizes joint firm profits.

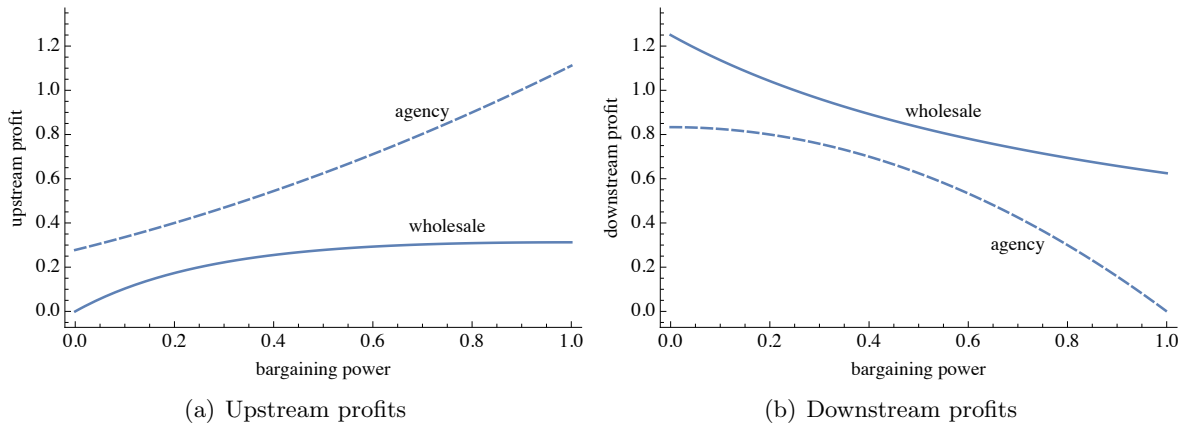


Figure 2: Upstream and downstream profits as a function of bargaining power

Figure 2(a) compares the upstream firm’s profit under the two selling methods, whereas 2(b) makes the same comparison for the downstream firm. For a given value of the bargaining power parameter, the upstream firm always prefers agency pricing whereas the downstream firm prefers wholesale pricing. This may explain why the publishers wanted to move to agency pricing in 2010, whereas retailers such as Amazon were initially unwilling to move to agency pricing. Notice that we get the opposite result when firms are using take-it-or-leave-it contracts. In that case $\lambda = 1$ when using wholesale pricing while $\lambda = 0$ when using agency pricing which means that moving to agency means higher profits for the downstream firms and lower profits for the upstream firms (see also Proposition 3 of Johnson, 2017).

3 Empirical Strategy

To make the model amendable for estimation, in this section we extend the model to allow for multiple upstream and downstream firms, as well as multi-product firms. In addition, we model demand using a nested logit discrete choice framework. Below we solve for the equilibrium conditions, which we take to the data in Section 5.

3.1 Demand

We consider a market in which there are N different products sold by n^U “upstream” publishers through n^D “downstream” retailers. Each product is a book-retailer pair. That is, a product is a specific book sold by a specific retailer. This means that product j sold by one retailer may be the same book as product k sold by another retailer—the idea is that different retailers represent different points in product space. The utility consumer i derives from product j is given by

$$u_{ij} = \alpha p_j + x_j' \beta + \xi_j + \bar{\varepsilon}_{ij}, \quad (6)$$

where p_j is the price of product j , x_j and ξ_j are observed and unobserved characteristics of book j , α and β are corresponding demand parameters, and $\bar{\varepsilon}_{ij}$ is a consumer-product specific utility shock. We allow for an outside option, which has utility $u_{i0} = \bar{\varepsilon}_{i0}$. To allow for correlation between utility shocks of related products, we assume the utility shock is distributed as in the nested logit model. Specifically, we assume that each product belongs to a mutually exclusive group $g = 0, \dots, G$. The utility shock can then be written as

$$\bar{\varepsilon}_{ij} = \zeta_g + (1 - \sigma)\varepsilon_{ij},$$

where ε_{ij} is extreme value distributed, σ is a nesting parameter, and ζ_g is distributed such that $\bar{\varepsilon}_{ij}$ is extreme value distributed as well. Note that for $\sigma = 0$, within-group correlation is zero, and the model is equivalent to a conditional logit model.

Letting $\delta_j = \alpha p_j + x_j' \beta + \xi_j$, the within-group market share of product j is

$$s_{j|g} = \frac{\exp[\delta_j/(1 - \sigma)]}{D_g}.$$

where $D_g = \sum_{j \in \mathcal{J}_g} \exp[\delta_j/(1 - \sigma)]$ and \mathcal{J}_g is the set of products in group g . The probability of buying a product from group g is

$$s_g = \frac{D_g^{(1-\sigma)}}{\sum_g D_g^{(1-\sigma)}}.$$

The market share of product j is then $s_j = s_{j|g}s_g$, or

$$s_j = \frac{\exp[\delta_j/(1 - \sigma)]}{D_g^\sigma \left[\sum_g D_g^{(1-\sigma)} \right]}.$$

Group 0 represents the outside option; its market share is given by

$$s_0 = \frac{1}{\sum_g D_g^{(1-\sigma)}}.$$

3.2 Wholesale Model

We model wholesale and retail pricing as a two stage game. In stage one, the publisher and retailer of each product j agree to a wholesale contract in which the retailer pays the publisher a wholesale price w_j for book j .⁷ All contracts are determined simultaneously in stage one. In stage two, retailers simultaneously choose retail prices given the wholesale terms established in stage one. Normalizing the size of the market to one, the variable profit of the retailer from selling product j is given by

$$\pi_j^D(p) = (p_j - w_j - c_j^D) s_j(p), \quad (7)$$

where p_j is the price of product j , w_j is the wholesale price, and c_j^D is the retailer's marginal cost for product j . The profit equation of the publisher for product j is then

$$\pi_j^U(p) = (w_j - c_j^U) s_j(p), \quad (8)$$

where c_j^U is the publisher's marginal cost for product j . The variable joint profit of the publisher and retailer associated with product j is $\pi_j^J = (p_j - c_j^D - c_j^U) s_j(p)$.

Downstream Market

Overall profits of the retailer are given by

$$\pi^D = \sum_{j \in \Omega^D} (p_j - w_j - c_j^D) s_j(p),$$

⁷The actual contracts used are typically called agency contracts since the retailer keeps a fraction r_j of the recommended price ρ_j for every book sold and the publisher receives the remainder. However, the retailer is free to set a discount, which means that these type of contracts are equivalent to wholesale agreements. This follows from the variable profits of the retailer from selling product j , which is given by

$$\begin{aligned} \pi_j^D(p) &= (r_j \rho_j - (\rho_j - p_j) - c_j^D) s_j(p); \\ &= (p_j - (1 - r_j) \rho_j - c_j^D) s_j(p), \end{aligned}$$

where $\rho_j - p_j$ reflects the discount the retailer may set. Note that the term $(1 - r_j) \rho_j$ can be treated as a per-book wholesale price w_j that is paid to the publisher and that is what we will do in this section.

where Ω^D is the set of products sold by the retailer. We assume a pure-strategy Nash equilibrium in retail prices, which means the first-order condition for product j is given by

$$s_j + \sum_{k \in \Omega^D} m_k^D \frac{\partial s_k}{\partial p_j} = 0, \quad (9)$$

where $m_k^D = p_k - w_k - c_k^D$ is the downstream margin on product k . The marketshare derivative with respect to price is given by

$$\frac{\partial s_k}{\partial p_j} = \begin{cases} \alpha s_k (1 - s_k + \gamma(1 - s_{k|g})) & \text{if } k = j; \\ -\alpha s_j (s_k + \gamma s_{k|g}) & \text{if } k \neq j \text{ and in same nest;} \\ -\alpha s_j s_k & \text{if } k \neq j \text{ and not in same nest,} \end{cases} \quad (10)$$

where $\gamma = \sigma/(1 - \sigma)$.

Upstream Market

We next describe how prices are set in the upstream market, assuming wholesale prices w are the outcome of a bilateral bargaining process between publishers and retailers. We assume that separate wholesale prices are chosen for each product, i.e., w is allowed to vary across retailer-book pairs.

Overall profits of the upstream firm are given by

$$\pi^U = \sum_{j \in \Omega^U} m_j^U s_j(p),$$

where $m_j^U = w_j - c_j^U$ is the upstream margin for product j , Ω^U is the set of products produced by the upstream firm, and c_j^U is the upstream firm's marginal cost for product j .

We assume the wholesale price is determined through Nash bargaining. The Nash bargaining solution for each publisher-retailer combination du is a vector of input costs w_{du} that maximizes the Nash product

$$NP_{du}(w_{du}; w_{-du}) = (\pi^U - d_{du}^U)^\lambda (\pi^D - d_{du}^D)^{1-\lambda}, \quad (11)$$

where d_{du}^U and d_{du}^D are disagreement payoffs and $\lambda \in [0, 1]$ is a bargaining power parameter between a retailer and publisher pair. Although we do not index λ to keep the notation simple, in our empirical application we allow for a different λ for each publisher-retailer pair. Because of downstream competition, the bargaining payoffs depend on the outcome of bilateral nego-

tiations of other parties. This problem has been extensively studied in the theoretical literature (e.g., Horn and Wolinsky, 1988; McAfee and Schwartz, 1994) and has been used in empirical work to model bilateral bargaining over input prices in oligopolistic markets (see Crawford and Yurukoglu, 2012; Grennan, 2013; Gowrisankaran, Nevo, and Town, 2015; Collard-Wexler, Gowrisankaran, and Lee, 2017; Ho and Lee, 2017). As in Crawford and Yurukoglu (2012), we are looking for a Nash-in-Nash equilibrium, i.e., the input prices that maximize the Nash product, taking as given input prices determined between other publisher-retailer pairs (cf. Horn and Wolinsky, 1988), i.e.,

$$w_{du} \text{ maximizes } NP_{du}(w_{du}; w_{-du}) \text{ given } w_{-du} \quad \forall d, \forall u.$$

The Nash-in-Nash equilibrium concept implies that each publisher and retailer meet separately and simultaneously, so w_{-du} are the conjectured input prices of the other pairs. We assume that if negotiations break down between a publisher and retailer, then all the titles of the publisher will not be carried by the retailer. However, unlike the bilateral monopoly case discussed in Section 2, consumers may buy the book at another retailer or shift to books published by other publishers, which means disagreement payoffs will generally not be zero. Specifically, we assume disagreement payoffs for publisher-retailer combination du are given by

$$\begin{aligned} d_{du}^U &= \sum_{k \in \Omega^U \setminus \{j \in du\}} m_k^U \Delta s_k^{-du}; \\ d_{du}^D &= \sum_{k \in \Omega^D \setminus \{j \in du\}} m_k^D \Delta s_k^{-du}. \end{aligned}$$

In these expressions, Δs_k^{-du} is defined as the additional market share for product k when books of the publisher-retailer pair du are not offered, i.e.,

$$\Delta s_k^{-du} = \frac{\exp[\delta_k/(1-\sigma)]}{D_g'^{\sigma} \left[\sum_g D_g'^{(1-\sigma)} \right]} - \frac{\exp[\delta_k/(1-\sigma)]}{D_g^{\sigma} \left[\sum_g D_g^{(1-\sigma)} \right]}, \quad (12)$$

where $D_g' = \sum_{l \in \mathcal{J}_g \setminus \{j \in du\}} \exp \delta_l/(1-\sigma)$. So the disagreement payoff for the pair du consists of the additional profits on all other books that are sold or published by either the upstream or downstream firm, taking into account that all books of publisher u are not available at retailer d and that the demand for those other books may have increased as a result of the products du not being sold.

The bargaining first-order condition is found by setting the derivative of equation (11) with

respect to w_{du} equal to zero for all products that belong to the set of products offered by the publisher-retailer pair du . Let j be such a product. The first-order condition for product j is

$$\lambda (\pi^U - d_{du}^U)^{\lambda-1} (\pi^D - d_{du}^D)^{1-\lambda} \frac{\partial \pi^U}{\partial w_j} + (1 - \lambda) (\pi^U - d_{du}^U)^\lambda (\pi^D - d_{du}^D)^{-\lambda} \frac{\partial \pi^D}{\partial w_j} = 0. \quad (13)$$

Equation (13) be simplified to

$$\lambda (\pi^D - d_{du}^D) \frac{\partial \pi^U}{\partial w_j} + (1 - \lambda) (\pi^U - d_{du}^U) \frac{\partial \pi^D}{\partial w_j} = 0. \quad (14)$$

The partial derivatives $\partial \pi^U / \partial w_j$ and $\partial \pi^D / \partial w_j$ are given by

$$\frac{\partial \pi^U}{\partial w_j} = \sum_{k \in \Omega^U} \frac{d\pi_k^U}{dw_j} \quad \text{and} \quad \frac{\partial \pi^D}{\partial w_j} = \sum_{k \in \Omega^D} \frac{d\pi_k^D}{dw_j},$$

where $d\pi_k^U / dw_j$ and $d\pi_k^D / dw_j$ are the *total* derivatives of π_k^U and π_k^D with respect to w_j . These total derivatives include the direct effect of w_j on the profits as well as an indirect effect that comes through changes in equilibrium prices $p^*(w)$ and are derived in Appendix B.⁸ Condition (14) together with condition (9) yield the equilibrium wholesale input prices and retail prices.

3.3 Agency Model

In the agency model, retail prices are set by the publishers, while the retailers obtain a royalty r_j . The variable profit of the retailer from selling product j is

$$\pi_j^D(p) = (r_j p_j - c_j^D) s_j(p).$$

The variable profit of the publisher from selling product j is

$$\pi_j^U(p) = ((1 - r_j) p_j - c_j^U) s_j(p).$$

The variable joint profit of the publisher and retailer associated with product j is $\pi_j^J = (p_j - c_j^D - c_j^U) s_j(p)$.

⁸An alternative approach, which is used in Draganska, Klapper, and Villas-Boas (2010) and Ho and Lee (2017), assumes retail prices and input prices are simultaneously determined, which allows one to treat retail prices as fixed.

Upstream Market

In the agency model, the publisher determines the retail price p_j . Overall profits of the publisher are given by

$$\pi^U = \sum_{j \in \Omega^U} ((1 - r_j)p_j - c_j^U) s_j(p).$$

As in the wholesale model, we assume a pure-strategy Nash equilibrium in retail prices so that the first-order condition for product j is

$$(1 - r_j)s_j + \sum_{k \in \Omega^U} m_k^U \frac{\partial s_k}{\partial p_j} = 0. \quad (15)$$

where $m_k^U = (1 - r_j)p_j - c_j^U$ is the upstream margin on product k and the marketshare derivative with respect to price is given by equation (10).

Downstream Market

The Nash bargaining solution is a vector of royalties r that maximizes the Nash product:

$$NP_{du}(r_{du}; r_{-du}) = (\pi^U - d_{du}^U)^\lambda (\pi^D - d_{du}^D)^{1-\lambda},$$

The bargaining first-order condition for product j is found by setting the derivative of the Nash product with respect to r_j equal to zero, and can be simplified to

$$\lambda (\pi^D - d_{du}^D) \frac{\partial \pi^U}{\partial r_j} + (1 - \lambda) (\pi^U - d_{du}^U) \frac{\partial \pi^D}{\partial r_j} = 0. \quad (16)$$

The partial derivatives $\partial \pi^U / \partial r_j$ and $\partial \pi^D / \partial r_j$ are given by

$$\frac{\partial \pi^U}{\partial r_j} = \sum_{k \in \Omega^U} \frac{d\pi_k^U}{dr_j} \quad \text{and} \quad \frac{\partial \pi^D}{\partial r_j} = \sum_{k \in \Omega^D} \frac{d\pi_k^D}{dr_j},$$

where $d\pi_k^U / dr_j$ and $d\pi_k^D / dr_j$ are the *total* derivatives of π_k^U and π_k^D with respect to r_j . These total derivatives include the direct effect of r_j on the profits as well as an indirect effect that comes through changes in equilibrium prices $p^*(r)$ and are derived in Appendix B. Condition (16) together with condition (15) yield the equilibrium agency royalties and retail prices.

4 Vertical Contracts in the E-Book Industry

We apply the empirical model developed in the previous section to the market for e-books. In this section, we first provide a description of important changes in the contracts between book publishers, the upstream firms in market for e-books, and the downstream book retailers. We then use a large dataset on retail prices to show how retail prices changed as a result of the switch from wholesale to agency contracts between publishers and bookstores in the period 2014-2015, which happened after a period of intense bargaining between retailers and publishers. The goal of this analysis is to motivate the use of the bargaining model developed in Sections 2 and 3 to study this market.

4.1 Background

Initially e-books were sold using wholesale contracts: the publishers would set a list price for the e-book and would sell the book to a retailer for roughly half the list price, which would then set a retail price for which the book is sold to consumers. This changed in 2010, when Apple together with five of the (then) Big Six publishers developed the agency model. The two most important differences between the agency and the wholesale model are that the retail price is set by the publisher and that the retailer receives a royalty share of each book sold. The switch from wholesale to agency contracts led to an immediate increase in retail prices.

In 2012 the US Department of Justice sued Apple and the publishers for conspiring to raise the prices of e-books. Three of the publishers settled right away, the other two followed in early 2013. As part of the settlement agreement, the publishers could not set retail prices for a period of two years. Moreover, the most-favored nation clauses that were fundamental for making the switch to the agency model were banned for a period of five years. One motivation of the two-year ban by the U.S. district court was to provide a reset of the bilateral bargaining relationship between retailers and publishers.

Table 1 displays in the first column the effective date of the start of the ban on agency contracts observed in the period after the settlement with the DOJ (De los Santos and Wildenbeest, 2017) for each of the now Big Five publishers. Although Random House was not a conspirator defendant in the DOJ lawsuits, Random House adopted the terms of the settlement after a merger with Penguin in July 2013. The second column of Table 1 displays the dates when Amazon and each of the publishers announced that a bilateral agreement has been reached. The third column displays the dates when the switch to agency can be identified in the data. The table shows that each publisher

Table 1: New contract announcement and switch dates

	Start of the agency ban	New agency agreement announcement	Amazon switch to agency
Simon & Schuster	Dec 17, 2012	Oct 20, 2014	Jan 01, 2015
Hachette	Dec 04, 2012	Nov 13, 2014	Apr 09, 2015
Macmillan	Apr 04, 2013	Dec 18, 2014	Jan 05, 2015
Harper Collins	Sep 10, 2012	Apr 13, 2015	Apr 15, 2015
Penguin Random House	Sep 01, 2013	Jun 18, 2015	Sep 01, 2015

Sources: The contract announcement comes from various news reports and Amazon.com. The switch to the new Agency comes from data from Amazon.com collected for this study. The dates of the start of the agency ban, which corresponds to the switch to wholesale model under the terms on settlement comes from De los Santos and Wildenbeest (2017).

announced an agreement with Amazon prior to the end of the two-year ban. The agreements did not go into effect immediately; the switch to agency varied between January and September of 2015.

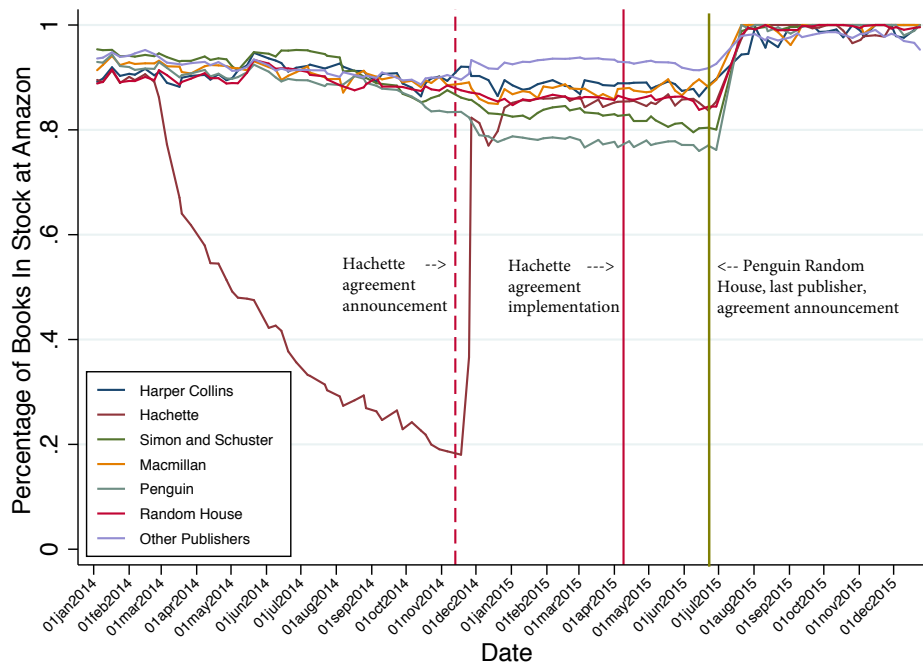


Figure 3: Amazon book inventory decisions by publisher

In the period leading to the expiration of the two-year ban on agency contracts, publishers and retailers engaged in a relatively lengthy period of negotiations over the conditions under which the publishers would regain control of retail prices. Especially the negotiations between Hachette became very public, and included various tactics to put pressure on the other side, including the delayed delivery of books, selling books at list price, removing the pre-order button, and the public

involvement of some bestseller writers on social media. Figure 3 provides some evidence of this by plotting the percentage of books sold at Amazon that are in stock, split out by publisher. The dispute between Amazon and Hachette started in March 2014 when Amazon did not allow customers to pre-order and reduced inventories of newly released Hachette books, which lead to significant shipping delays of Hachette books for Amazon's consumers. The figure shows that the percentage of Hachette books in our sample that were in stock at Amazon declined sharply from levels around 90 percent, which was similar to other publishers, to less than 20 percent shortly before the agreement announcement in November 2014. After the agreement, the percentage of Hachette books in stock immediately returned to around 80 percent, which was again similar to other publishers. The figure also show that there was a gradual reduction of the percentage of books in stock for other Big Five publishers starting from the beginning of the year 2014, particularly for Penguin Random House, which was the latest publisher to reach an agreement with Amazon. After a new agreement with this last publisher was announced on June 18, 2015, the percentage of books in stock increased sharply to levels near 100 percent for all publishers.

4.2 The Effect of the Switch to Agency on Retail Prices

In this section, we use a large dataset on retail prices for e-books to study the price effects of the switch to agency contracts. Our sample runs between early 2014 and the end of 2015 and consists of daily prices for a large number of e-book titles. All titles are new and former New York Times bestseller books. For a specific title we observe its retail price as well as sales rank at both Amazon and Barnes & Noble. Moreover, we observe book characteristics such as list price, publisher, ratings, number of reviews, and number of pages. We also have data on the printed version of the book, including format, list price, and retail price. Table 2 summarizes the variables.

For the analysis in this section we use a similar difference-in-differences (DID) approach as in De los Santos and Wildenbeest (2017), but whereas De los Santos and Wildenbeest study the transition from agency contracts to wholesale contracts that followed the Justice Department's lawsuit against the major publishers and Apple in 2012, we focus on the transition from wholesale to agency that happened after the two-year ban on agency had expired in the period 2014-2015.

As shown in Table 1, new contracts were announced between Amazon and the major publishers at different points in time, which also meant that the actual switching dates at Amazon were staggered. We exploit this cross-publisher variation in the timing of the switch in a difference-in-

Table 2: Summary statistics

	Harper Collins	Hachette	Simon & Schuster	Macmillan	Penguin Random House	Other publishers
<i>Price e-book</i>						
Amazon	9.05 (3.60)	8.70 (2.83)	10.38 (3.04)	9.08 (2.68)	8.83 (2.87)	8.29 (4.07)
Barnes & Noble	10.71 (3.79)	9.28 (2.76)	11.54 (2.77)	9.86 (2.63)	10.98 (2.72)	10.42 (4.08)
Kindle list price	16.95 (7.32)	15.37 (8.09)	16.73 (6.42)	16.44 (7.32)	15.53 (6.94)	16.91 (7.65)
<i>Book characteristics</i>						
Ratings	4.21 (0.38)	4.18 (0.41)	4.25 (0.40)	4.18 (0.42)	4.21 (0.40)	4.35 (0.36)
Number of reviews	956.14 (1433.81)	1537.35 (2043.11)	1187.04 (1859.16)	903.21 (1432.55)	1395.94 (2668.86)	880.05 (2658.83)
Number of years since release	2.26 (2.14)	2.21 (1.83)	2.38 (2.07)	2.29 (2.58)	2.27 (2.16)	2.18 (2.63)
<i>Print book characteristics</i>						
Sales rank	150269.0 (271761.1)	119502.6 (223733.9)	114651.2 (206086.7)	134988.7 (188678.5)	92223.89 (153229.68)	207529.8 (476889.3)
Number of pages	407.22 (134.79)	415.05 (134.18)	401.88 (145.2)	393.91 (120.99)	414.86 (223.77)	353.25 (144.69)
Book weight (oz.)	17.20 (8.43)	17.10 (8.53)	17.15 (9.13)	16.81 (8.17)	17.50 (9.48)	15.51 (9.95)
List price	22.49 (7.73)	20.49 (7.72)	21.78 (7.53)	22.13 (7.49)	21.25 (8.80)	19.06 (8.27)
Number of titles	347	361	351	223	1,124	2,945
Number of observations	70,199	84,483	83,870	52,624	280,007	507,618

Notes: The table presents the means of each variable, standard deviation in parentheses.

differences setup. Specifically, the baseline specification we estimate is

$$\ln(\text{price}_{jt}) = \gamma \cdot (\text{agency}_{jt} \times \text{bigfive}_j) + \beta \cdot X_j + \lambda_p + \lambda_w + \varepsilon_{jt}, \quad (17)$$

where price_{jt} is the e-book price of title j at time t , agency_{jt} is an indicator for whether at time t title j was sold using an agency contract, X_j are book characteristics, and λ_p and λ_w are publisher and week fixed effects. The interaction $\text{agency}_{jt} \times \text{bigfive}_j$ is the difference-in-differences estimator and captures the effect of the switch to the agency model.

Table 3 gives the main results for the difference-in-differences analysis. We estimate equation (17) separately for Amazon and Barnes & Noble. For each retailer, we estimate a specification that allows for publisher fixed effects and a specification that has book fixed effects. As can be seen from Table 3, the difference-in-differences estimator is very similar across the two specifications. For Amazon, the estimates imply that prices went up by approximately 16 percent as a result of the switch from wholesale to agency; for Barnes & Noble prices went down by approximately 2

percent.

Table 3: Main results difference-in-differences analysis

	Amazon		Barnes & Noble	
	Publisher fixed effects	Book fixed effects	Publisher fixed effects	Book fixed effects
<i>Difference-in-differences estimator</i>				
Agency × Big Five	0.150*** (0.010)	0.154*** (0.009)	-0.015** (0.007)	-0.022*** (0.007)
<i>Other controls</i>				
ln(sales rank)	0.003 (0.003)	-0.013*** (0.003)	-0.007** (0.003)	-0.003 (0.002)
Rating on Amazon	-0.032*** (0.009)	-0.012 (0.010)	-0.006 (0.026)	-0.012* (0.007)
Number of reviews on Amazon	-0.000*** (0.000)	-0.000 (0.000)	-0.000*** (0.000)	-0.000 (0.000)
Years since release	-0.008*** (0.002)	0.034 (0.033)	-0.000 (0.002)	0.016 (0.027)
Kindle list price	0.024*** (0.001)		0.023*** (0.001)	
Number of pages in the book	-0.000* (0.000)		-0.000 (0.000)	
Weight of the book (ounces)	0.004*** (0.001)		0.003*** (0.001)	
Constant	1.864*** (0.061)	2.192*** (0.084)	2.077*** (0.103)	2.364*** (0.074)
<i>R</i> -squared	0.416	0.681	0.260	0.819
Number of observations	476,899	476,899	486,165	486,165

Notes: Dependent variable is $\ln(\text{price})$. All specifications include week fixed effects. Standard errors (clustered by book) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 4 shows the results for a specification in which we split out the effect by publisher. Consider first Amazon. The results for the baseline specification show that the effect is not the same across publishers: while the effect for Hachette is not significantly different from zero, the estimates for Penguin Random House indicate that prices went up by 32 percent following the switch. The findings for books sold at Barnes & Noble are also mixed: for books published by Macmillan and Harper Collins prices went down following the switch, while prices for books published by Penguin Random House prices went up by approximately 5 percent.

A possible explanation for the absence of a price effect for Hachette books sold at Amazon is that Amazon was selling many Hachette books at list price during their contract negotiations to put pressure on Hachette. Figure 4 shows the average price changes at Amazon by publisher around the agreement announcement (dashed vertical line) and the date of the switch to agency (solid vertical line). Although prior to the dispute Amazon's prices for Hachette books were consistently lower in comparison to Barnes & Noble, Amazon increased prices for Hachette books to levels similar to

Table 4: Effect of the switch to agency by publisher

	Amazon			Barnes & Noble		
	Baseline	30-day window	7-day window	Baseline	30-day window	7-day window
Agency × Harper Collins	0.188*** (0.020)	0.287*** (0.025)	0.176*** (0.031)	-0.051*** (0.015)	0.007 (0.016)	-0.034* (0.021)
Agency × Hachette	0.017 (0.010)	-0.02*** (0.010)	-0.014* (0.008)	0.011 (0.008)	0.037*** (0.010)	0.020 (0.020)
Agency × Simon & Schuster	0.201*** (0.015)	0.142*** (0.019)	0.085*** (0.019)	-0.007 (0.010)	-0.030** (0.013)	-0.048*** (0.010)
Agency × Macmillan	0.142*** (0.022)	0.092*** (0.016)	0.063*** (0.015)	-0.084*** (0.017)	-0.043*** (0.015)	-0.061*** (0.013)
Agency × Penguin Random House	0.279*** (0.020)	0.340*** (0.022)	0.270*** (0.022)	0.049*** (0.009)	0.044*** (0.012)	0.055*** (0.011)
<i>R</i> -squared	0.424			0.820		
Number of observations	476,924			486,455		

Notes: The table presents difference-in-differences coefficient estimates by publisher. The baseline specification includes switching interaction coefficients for each publisher. As publishers switched at various dates, the windowed coefficients are obtained from separate regressions using observations around the time of the switch of each publisher. All specifications include controls as in the main specification. Dependent variable is $\ln(\text{price})$. Week fixed effects included. Standard errors (clustered by book) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Barnes & Noble in the months before an agreement was reached, which could have been motivated to hurt sales of Hachette titles during this period. Figure 4(a) shows that Amazon reduced prices for Hachette books to pre-dispute levels after the agreement announcement. The evolution of prices for the other publishers, which are shown in Figures 4(b)-4(e), do not exhibit similar patterns. For these publishers, the price gap between retail prices at Barnes & Noble and Amazon is stable, and even increased after the announcement as in the case of Penguin Random House. After the agency agreements were in place, each publisher increased prices to similar levels as Barnes & Noble. Notice that another price pattern that is not captured by the windowed DID regressions above, but is evident from Figures 4(a) to 4(e), is the increase in prices for every publisher at both Amazon and Barnes & Noble after the announcement of the agreement between Amazon and Penguin Random House (as indicated by the dashed black vertical lines), the last Big Five publisher to reach an agreement (on June 18, 2015). These price increases could be interpreted as the publishers' strategic responses to anticipated price increases of Penguin Random House book titles at Amazon, as happened after September 1, 2015.

In Appendix D we show that the above results are robust to different ways of aggregating the data (weekly and monthly), using non-Big Five publishers as an additional control group, and including print-book prices in the analysis. In Appendix D we also discuss the results of a Placebo test in which we replicate the analysis using print-book prices instead of e-book prices—since printed

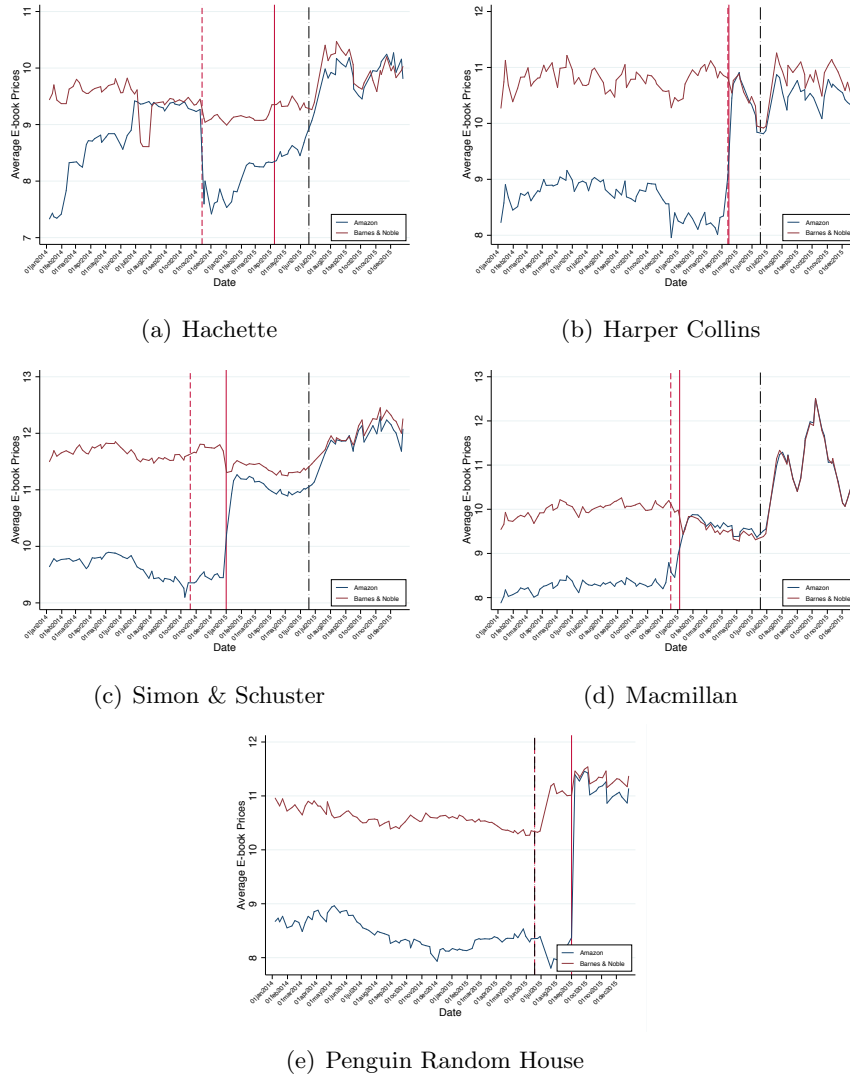


Figure 4: Retail prices at Amazon by publisher

books did not see a change in selling method, one would not expect to find any significant results for the DID estimator.

The results from the difference-in-differences analysis point to two important observations: (1) the effect of the switch to agency agreements was different for Amazon than for Barnes & Noble; and (2) there is heterogeneity in the magnitude of the effect across publishers. These results are consistent with our theoretical and empirical framework discussed in Sections 2 and 3, for a situation in which different retailers and publishers have different relative bargaining power parameters, and therefore respond differently to a move from wholesale contracts to agency contracts.

5 Estimation of the Bargaining Model

In this section we provide estimates of the bargaining model. We estimate the model in two steps. First, we estimate the demand side parameters. Second, we estimate the marginal cost and bargaining parameters of the supply side model under both wholesale and agency contracts, taking the demand side parameters as given.

5.1 Data

To estimate the structural model we use a subset of the data discussed in the previous section. For the estimation of the wholesale model we use four weeks of data from July 2014, which corresponds to a period in which all Big Five publishers were using wholesale agreements. For the estimation of the agency model we use four weeks of data from September 2015. The last publisher to make the switch was Penguin Random House on September 1st, 2015, which means that all Big Five publishers were using agency contracts by this time. We exclude books sold by the other (smaller) publishers, since these books were not sold using agency agreements during the second half of our sample.

As we lack quantity data for each book title, we use the observed sales rank data at both Amazon and Barnes & Noble to infer sales at each of these retailers. Following Chevalier and Goolsbee (2003), we assume that books sales quantities follow a Pareto distribution, i.e., the probability that an observation τ exceeds a level Sales is

$$\Pr(\tau > \text{Sales}) = \left(\frac{k}{\text{Sales}} \right)^\theta, \quad (18)$$

where k and θ are the scale and shape parameter of the Pareto distribution. Since the fraction of books that has more sales than a particular title is $(\text{Rank} - 1)/(\text{Total number of books})$, we can write equation (18) as

$$\frac{\text{Rank} - 1}{\text{Total number of books}} = \left(\frac{k}{\text{Sales}} \right)^\theta.$$

Solving for Sales gives

$$\text{Sales} = k \cdot \left(\frac{\text{Rank} - 1}{\text{Total number of books}} \right)^{-1/\theta}.$$

Taking logs gives

$$\ln(\text{Sales}) = \gamma_0 + \gamma_1 \ln(\text{Rank} - 1), \quad (19)$$

where $\gamma_0 = \ln(k) + \frac{1}{\theta}(\text{Total number of books})$ and $\gamma_1 = -\frac{1}{\theta}$. Using various sources of sales data,

Chevalier and Goolsbee find θ to be in the range 0.9 to 1.3 and use 1.2 as the basic estimate of θ in their analysis. To obtain an estimate of the shape parameter of the Pareto distribution that fits our sales rank data, we use an online sales rank calculator that transforms Kindle sales ranks to Kindle sales data.⁹ Estimating equation (19) using this data by OLS leads to an estimate for θ of 1.19, which is very close to the estimate of 1.2 that Chevalier and Goolsbee (2003) use throughout their paper.¹⁰ We use the fitted sales to transform Kindle sales rank data into quantity data. To obtain Barnes & Noble sales we use the same estimated equation, but shift the intercept to reflect that overall e-book sales at Barnes & Noble are approximately one-quarter of those at Amazon.¹¹

Table 5 provides summary statistics of the main variables we use by publisher. Panel A of the table gives summary statistics for the period in which wholesale agreements were used (July 2014) whereas panel B gives summary statistics for the period in which agency agreements were used by the Big Five publishers (September 2015). We aggregate the data into weekly observations. For both periods we use 77 different titles, which results in a total of 1,146 weekly observations (616 for the wholesale period and 530 for the agency period). The largest Big Five publisher, Penguin Random House, also represent most observations in our sample. Macmillan is the smallest with overall 9 titles.

A comparison of the average prices under the two selling regimes indicates that even though average prices were about 10 percent higher at Barnes & Noble than at Amazon during the wholesale period, average prices were basically the same across the retailers during the agency period, despite the five-year ban on the use of MFN clauses that followed the settlements between the Big Five publishers and DOJ in 2012.

5.2 Demand Estimates

We estimate the demand side of the model using data on markets shares, prices, and product characteristics. Denote θ the parameters of the demand model that need to be estimated. The predicted market share $s_j(\delta(\theta), \theta)$ should match the observed market shares \hat{s}_j . Since we are assuming the model has a nested logit structure, we estimated the following specification:

$$\log(s_j) - \log(s_0) = X_j\beta + \alpha p_j + \sigma \log(s_{j|g}) + \xi_j.$$

⁹See <https://kindlepreneur.com/amazon-kdp-sales-rank-calculator>.

¹⁰Using 3,720 daily sales rank observations for e-books sold at Amazon as well as sales data obtained using the sales rank calculator, the estimated equation is $\ln(\widehat{\text{Sales}}) = 10.572 - 0.843 \ln(\text{Rank} - 1)$, with $R^2 = 0.967$.
(0.015) (0.003)

¹¹Although precise figures are not available, according to their sales data the Digital Book Publisher Vook (re-branded as Pronoun in 2015) estimates Amazon's market share to be 60 percent, while Barnes & Noble's market share is 15 percent. See https://vook.com/static/media/media/Ebook_Marketplaces_and_Royalties_guide.pdf.

Table 5: Summary statistics

	Hachette	Harper Collins	Macmillan	Penguin Random House	Simon & Schuster
Panel A. Wholesale period					
<i>Price e-book</i>					
Amazon	10.28 (2.79)	9.39 (3.47)	6.18 (3.29)	8.18 (3.01)	10.12 (3.14)
Barnes & Noble	11.17 (2.45)	10.25 (3.40)	8.79 (3.05)	9.51 (2.53)	11.37 (2.02)
<i>Weekly sales</i>					
Amazon	10,440 (14,599)	3,815 (5,657)	1,177 (570)	5,190 (10,863)	5,841 (6,770)
Barnes & Noble	1,664 (1,335)	741 (884)	243 (202)	1,058 (2,224)	1,422 (1,938)
<i>Book characteristics</i>					
Ratings	4.26 (0.35)	4.35 (0.27)	4.31 (0.16)	4.22 (0.44)	4.12 (0.53)
Number of reviews	4,077 (3,713)	3,261 (3,355)	2,033 (2,840)	3,142 (2,940)	2,489 (2,546)
Number of years since release	1.34 (2.07)	3.61 (6.40)	6.35 (10.58)	3.68 (6.25)	0.68 (0.93)
List price	22.70 (7.27)	19.66 (6.00)	13.16 (6.82)	20.59 (7.41)	23.88 (6.70)
Number of titles	10	9	6	43	9
Number of observations	80	72	48	344	72
Panel B. Agency period					
<i>Price e-book</i>					
Amazon	12.68 (2.84)	10.47 (4.46)	10.33 (3.81)	10.93 (2.16)	11.91 (1.96)
Barnes & Noble	12.62 (3.04)	10.69 (4.61)	9.19 (4.42)	10.59 (1.96)	12.03 (1.86)
<i>Weekly sales</i>					
Amazon	4,067 (6,617)	1,749 (4,883)	6,034 (7,934)	2,115 (1,990)	1,920 (2,013)
Barnes & Noble	2,054 (3,014)	481 (464)	1,715 (1,837)	772 (789)	782 (722)
<i>Book characteristics</i>					
Ratings	4.26 (0.27)	4.17 (0.35)	4.46 (0.28)	4.36 (0.30)	4.41 (0.22)
Number of reviews	4,608 (6,895)	3,829 (5,164)	4,834 (5,947)	8,626 (9,490)	3,001 (2,713)
Number of years since release	0.72 (0.74)	4.55 (7.25)	1.05 (0.75)	4.34 (5.59)	9.09 (23.59)
List price	22.31 (7.88)	18.82 (9.18)	22.50 (5.74)	17.18 (5.48)	21.17 (7.10)
Number of titles	9	10	3	44	11
Number of observations	71	70	19	290	80

Notes: The table presents the means of each variable, standard deviation in parentheses.

Since both p_j and $\log(s_{j|g})$ are likely to be correlated with ξ_j , we estimate the model using two-stage least squares. Note that the demand side specification should not depend on the selling method, so we can pool data from both periods.

Table 6 gives the demand-side estimates. In all specifications we estimate both periods together and allow for product-store fixed effects as well week fixed effects. In specification (1) we do not control for price endogeneity, which means we can estimate the demand side using OLS. Although the price coefficient is highly significant, demand is estimated to be relatively inelastic—the median own-price elasticity is -1.614 and the proportion of products for which demand is estimated to be inelastic is 0.142 .

The unobserved characteristic ξ_j in equation (6) captures unobserved quality, which is likely to be correlated with a book’s price. Since all our specifications include product-store fixed effects, the product-store-specific variation in unobserved quality that does not vary over time is captured by the product-store dummies. However, the product-store fixed effects will not pick up variation in prices due to differences in unobserved quality over time. For instance, a favorable review in Oprah’s Book Club may lead to a sudden increase in demand, which may have an effect on prices as well. To deal with any unobserved quality differences over time, we therefore estimate the model by two-stage least squares. Unfortunately, the BLP-type instruments that are typically used when estimating demand (see Berry, Levinsohn, and Pakes, 1995) are difficult to apply in this context, since e-book attributes do not explain much of the variance in sales and demand. Hausmann-type instruments are not suitable either, since there is no regional price variation in this market. However, an instrument that it is easy to apply in this context is the lagged price (see Villas-Boas and Winer, 1999). This instrument has been used in other markets in which it is difficult to use traditional instruments such as the market for console video games (Nair, 2007; Shiller, 2013). As can be seen in column (2) of Table 6, the estimated price coefficient increases in magnitude when using the lagged price as an instrument for price. The absolute value of the median own-price elasticity increases as a result, with less products facing inelastic demand. Note that we lose about a quarter of observations to create the lagged prices instrument.

Specification (3) gives demand estimates for the nested logit model estimated by two-stage least squares. Arguably, consumers are more likely to switch within retailer than across retailer because of the platform nature of the e-book market (Kindle versus Nook), and to capture this we assume Amazon and Barnes & Noble are part of different nests. In addition, we assume fiction and non-fiction books are part of different nests, which means we have four different nests in total: Amazon fiction, Amazon non-fiction, Barnes & Noble fiction, and Barnes & Noble non-fiction. As in specification (2) we use lagged prices to instrument for price. To instrument for the within-nest market shares $\log(s_{j|g})$ we use the number of within-nest products. Both the estimated price coefficient and the log of the within-nest market shares are highly significant. Demand is more

Table 6: Demand-side estimates

Variable	(1) Logit OLS		(2) Logit 2SLS		(3) Nested logit 2SLS	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Price	-0.162	(0.014)***	-0.257	(0.054)***	-0.163	(0.044)***
$\log(s_{j g})$	—		—		0.405	(0.141)***
<i>R</i> -squared	0.954					
Median own-price elasticity	-1.614		-2.563		-2.738	
Proportion inelastic demand	0.142		0.029		0.030	
Number of observations	1,146		853		853	

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. All specifications include product-store fixed effects, as well as week fixed effects. Both the logit 2SLS and nested logit 2SLS specifications use the lagged price as an instrument for price, while the nested logit specification also uses the number of within-nest products as an instrument for $\log(s_{j|g})$.

elastic than in the logit case, although the difference in median own-price elasticity is small.

5.3 Supply Estimates

Taking the estimated demand parameters as given, we next describe how to estimate the supply side. Our approach is to use the equilibrium conditions of the model to derive an expression for the marginal cost of the upstream firms and then use observed product characteristics that affect these marginal costs to get an estimate of the marginal cost parameters as well as the vector of relative bargaining power parameters λ . Throughout the analysis, we assume the bargaining parameters do not change throughout the sample, i.e., we estimate one bargaining parameter for each publisher-retailer combination (which does not depend on the selling method). In addition to assuming the bargaining parameters do not change as a result of change in selling method, we assume that the demand parameters as well as the marginal cost parameters remain the same across the two period, which means that the change in prices only reflects the change in selling method. Identification of the bargaining parameters therefore comes from the change in prices as a result of switching from the wholesale model to the agency model.

The marginal cost of the upstream firm can be obtained as the difference between the input price and the upstream margin where the input price in the wholesale model is the wholesale price and in the agency model the share of the price that goes to the upstream firm, i.e.,

$$mc_j^U = \begin{cases} w_j - m_j^U, & \text{if wholesale;} \\ (1 - r_j)p_j - m_j^U, & \text{if agency.} \end{cases} \quad (20)$$

The wholesale price w_j in the wholesale model is a fraction of the recommended price (see footnote 7)—since recommended prices are observed as part of our data, and the fraction of the recommended price going to the publishers was 70 percent during the wholesale period, we can treat w_j in equation (20) as known when estimating the model.¹² We are therefore left with estimating the royalty r during the agency period, as well as the upstream margins under both wholesale and agency.

The upstream margin when agency contracts are used, can be found by solving equation (15) for m^U , which gives, using matrix notation,

$$m^U = - (T^U \Delta)^{-1} (1 - r)s, \quad (21)$$

where T^U is an ownership matrix whose (j, k) th element is 1 if products j and k are published by the same publisher and zero otherwise and Δ is a matrix of market share derivatives with respect to price whose (j, k) th element is given by $\partial s_k / \partial p_j$ in equation (10).

The upstream margin when wholesale contracts are used can be found by solving the equilibrium condition in equation (14) for m^U —we show in Appendix C that, using matrix notation, this can be written as

$$m^U = - \left(T^U Z^w + E^D (T^U S) \frac{1 - \lambda}{\lambda} \right)^{-1} s, \quad (22)$$

where E^D is a vector whose j th element is given by $E_j^D = (\pi^D - d^D)^{-1} (\partial \pi^D / \partial w_j)$, Z^w is a matrix that captures how market shares change through changes in equilibrium prices, and whose (j, k) th element is given by

$$Z_{jk}^w = \alpha s_k (1 - s_k + \gamma(1 - s_{k|g})) \frac{\partial p_k^*}{\partial w_j} - \sum_{\substack{l \neq k \\ \text{same nest}}} \alpha s_l (s_k + \gamma s_{k|g}) \frac{\partial p_l^*}{\partial w_j} - \sum_{\substack{l \neq k \\ \text{not same nest}}} \alpha s_l s_k \frac{\partial p_l^*}{\partial w_j},$$

and where S is a matrix that has the market shares on the diagonal and the differences in market shares when product j is not offered in equation (12) as off-diagonal elements, i.e.,

$$S = \begin{bmatrix} s_1 & -\Delta s_2^{-1} & \dots & -\Delta s_N^{-1} \\ -\Delta s_1^{-2} & s_2 & \dots & -\Delta s_N^{-2} \\ \vdots & \vdots & \ddots & \vdots \\ -\Delta s_1^{-N} & -\Delta s_2^{-N} & \dots & s_N \end{bmatrix}.$$

Note that E^D will depend on the vector of downstream margins m^D , which, using matrix notation,

¹²According to industry experts (see for instance thepassivevoice.com/an-authors-perspective-on-the-hachette-amazon-battle) this fraction did not change from the 70 percent that was used during the first agency period.

can be written as

$$m^D = -(T^D \Delta)^{-1} s, \quad (23)$$

where T^D is an ownership matrix whose (j, k) th element is 1 if products j and k are sold by the same retailer and zero otherwise.

As shown by equation (21), the upstream margin under the agency model depends on the agency royalty r , whereas equation (22) shows that under wholesale the upstream margin depends on the bargaining power parameter λ . Rather than estimating both λ and r directly we take a different approach. Specifically, as shown in Appendix C, we can use the first-order condition of the bargaining game under agency in equation (16) to solve for the vector of relative bargaining parameters λ as a function the vector of royalties r , downstream margins m^D , and upstream margins m^U , i.e.,

$$L = -(A + T^D B^D + (T^D Z^r) m^D)^{-1} (E^U (T^D S)), \quad (24)$$

where $L = \lambda/(1 - \lambda)$, E^U is a vector whose j th element is given by $E_j^U = (\pi^U - d^U)^{-1} (\partial \pi^U / \partial r_j)$, B^D is a matrix whose (j, k) th element is given by $r_k s_k (\partial p_k^* / \partial r_j)$, and Z^r is a matrix that captures how market shares change through changes in equilibrium prices, and whose (j, k) th element is given by

$$Z_{jk}^r = \alpha s_k (1 - s_k + \gamma(1 - s_{k|g})) \frac{\partial p_k^*}{\partial r_j} - \sum_{\substack{l \neq k \\ \text{same nest}}} \alpha s_l (s_k + \gamma s_{k|g}) \frac{\partial p_l^*}{\partial r_j} - \sum_{\substack{l \neq k \\ \text{not same nest}}} \alpha s_l s_k \frac{\partial p_l^*}{\partial r_j}$$

Note that the downstream margin under agency is given by $m^D = rp - v$, which, assuming zero marginal cost v for the downstream retailer, implies that we can directly obtain m^D as a function of r .¹³ So for a given vector of agency royalties r we can directly obtain the downstream margin under agency m^D as well as use equation (21) to obtain m^U , which we can then use to solve for λ using equation (24). The bargaining parameters λ can then be used together with m^D from equation (23) to obtain the upstream margin m^U under wholesale using equation (22). Upstream marginal costs are a function of the upstream margins according to equation (20), which allows us to estimate a linear marginal cost equation in which we let $\log(mc_j^U)$ depend on observed product characteristics z_j and an unobserved characteristic ν_j that serves as an error term:

$$\log(mc_j^U) = z_j' \eta + \nu_j.$$

¹³Alternatively, a retailer's marginal cost v can be estimated alongside the other parameters.

Table 7: Parameter estimates of the bargaining model

Variable	(A) Bargaining model		(B) Take-it-or-leave-it contracts		(C) Retailers have all bargaining power		(D) Publishers have all bargaining power		(E) Bargaining model	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
<i>Marginal cost publishers</i>										
Constant	1.243	(0.635)*	0.922	(1.140)	1.724	(1.196)	0.966	(0.799)	1.224	(0.664)*
ln(number of reviews)	0.077	(0.020)***	0.081	(0.066)	0.024	(0.068)	0.129	(0.032)***	0.077	(0.020)***
ln(number of pages)	0.027	(0.098)	0.013	(0.167)	-0.036	(0.178)	0.024	(0.138)	0.029	(0.102)
ln(weeks since release)	-0.159	(0.021)***	-0.214	(0.042)***	-0.126	(0.036)***	-0.221	(0.033)***	-0.160	(0.021)***
Fiction indicator	-0.255	(0.040)***	-0.256	(0.116)**	-0.114	(0.104)	-0.364	(0.078)***	-0.253	(0.042)***
ln(weight printed book)	0.206	(0.066)***	0.225	(0.118)*	0.197	(0.103)*	0.235	(0.070)***	0.204	(0.066)***
<i>Marginal cost retailers (agency period)</i>										
Barnes & Noble	—		—		—		—		0.308	(2.074)
Amazon	—		—		—		—		0.643	(0.691)
<i>Bargaining parameters Amazon</i>										
Hachette	0.358	(0.138)***			0.000		1.000		0.394	(0.153)***
Harper Collins	0.803	(0.191)***			0.000		1.000		0.792	(0.177)***
Macmillan	0.547	(0.238)**			0.000		1.000		0.558	(0.219)**
Penguin Random House	0.538	(0.069)***			0.000		1.000		0.563	(0.079)***
Simon & Schuster	0.371	(0.117)***			0.000		1.000		0.430	(0.116)***
<i>Bargaining parameters Barnes & Noble</i>										
Hachette	0.425	(0.092)***			0.000		1.000		0.425	(0.108)***
Harper Collins	0.902	(0.157)***			0.000		1.000		0.926	(0.147)***
Macmillan	0.595	(0.146)***			0.000		1.000		0.599	(0.143)***
Penguin Random House	0.536	(0.063)***			0.000		1.000		0.549	(0.071)***
Simon & Schuster	0.256	(0.123)**			0.000		1.000		0.269	(0.131)**
Objective function	0.251		0.294		0.349		0.420		0.249	
Log-likelihood function	-621		-689		-761		-840		-618	
Vuong-stat			2.047		4.451		6.558			
(p value)			(0.020)		(0.000)		(0.000)			
Number of Observations	853		853		853		853		853	

Notes: Bootstrapped standard errors shown in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Marginal cost specifications for the publisher include a constant, publisher fixed effects, and week fixed effects. Estimates are based on the 2SLS estimates for the nested logit model shown in column (3) of Table 6.

We use the two-stage least squares estimates for the nested logit model shown in column (3) of Table 6 to estimate the supply side. Table 7 gives the parameter estimates of the publishers' marginal cost equation for various specifications of the bargaining model. Specification (A) of Table 7 gives the parameter estimates of the publishers' marginal cost equation for the bargaining model as well as the estimated bargaining parameters. As cost shifters we include the logarithm of the number of reviews, the number of pages, weeks since release, and weight of the printed version of the book, as well as a fiction dummy and publisher and week fixed effects; all marginal cost shifters except for the log of the number of pages are significantly different from zero and have the expected signs. The estimated bargaining power parameters for this specification are also shown in Table 7. All estimates of the bargaining parameters are statistically significant. The average of the estimated bargaining parameters is 0.533, which suggests the publishers have slightly more bargaining power than the retailers. However, there is substantial variation in the estimated bargaining parameters across publishers, with Penguin Random House and HarperCollins—the two largest trade publishers by units sold in 2016—having above average bargaining power and Hachette and Simon & Schuster the least.¹⁴

Table 7 also gives marginal cost estimates for several alternative supply side models. In column (B) we estimate the supply side assuming the publishers make take-it-or-leave-it offers in the wholesale model, and the retailers make take-it-or-leave-it offers in the agency model. Although most estimated marginal cost parameters do not differ much different from the ones estimated in column (A), the objective function values for the two specifications suggest the bargaining model outperforms the take-it-or-leave-it model. However, the two models are non-nested: the bargaining model assumes the bargaining parameters are constant across the two selling methods, while with take-it-or-leave-it contracts the publishers have all the bargaining power in the wholesale model, whereas the retailers have all the bargaining power in the agency model. To formally test which model gives the best fit to the data, we use a Vuong test (Vuong, 1989). The p value of the Vuong-stat in column (B) of Table 7 is very close to zero, which means according to this test, the bargaining model provides a better fit than the take-it-or-leave-it model. In column (C) and (D) of the table we report the marginal cost estimates of specifications in which we assume either the retailers (in column (C)) or the publishers (in column (D)) have all the bargaining power. Specification (C) gives a better fit to the data than specification (D). Although the specification with all the bargaining power assigned to the retailers does a reasonable job in fitting the data, as indicated

¹⁴According to purchases made through outlets tracked by NPD BookScan, the ranking of the Big 5 publisher in terms of units sold in 2016 is (1) Penguin Random House; (2) HarperCollins; (3) Simon & Schuster; (4) Hachette; and (5) Macmillan.

by the objective function value, estimating the bargaining power parameters still gives a better fit to data, as shown by the p value of the Vuong-stat.

When estimating specification (A) we have assumed that the marginal cost of the retailers is equal to zero during the agency period. In specification (E) we drop this assumption and estimate the retailers' marginal costs alongside the marginal cost of the publishers. Although not significant, the estimated marginal cost for Amazon is about 64 cents, whereas the marginal cost for Barnes & Noble is 31 cents. Although the difference are small, the estimated bargaining power parameters are uniformly higher than in specification (A). Notice that the fit of the model has slightly improved when estimating the marginal costs of the retailers during the agency period, so we will treat this as our preferred specification in the remainder of the paper.

Table 8 gives the estimated royalty parameters during the agency period for specification (E) of Table 7. The average royalty share across retailers and publishers is 0.307, which is about the same as the 30 percent that was typically used during the first agency period (between 2010 and 2012). Agency royalties are higher at Amazon than Barnes & Noble, which can be explained by Amazon's better bargaining position due to it's size relative to Barnes & Noble. The differences in agency royalties between the publishers are much higher for Barnes & Noble than for Amazon.

Table 8: Royalty parameter estimates

	Amazon		Barnes & Noble		Average
Hachette	0.332	(0.061)***	0.316	(0.087)***	0.324
HarperCollins	0.323	(0.084)***	0.181	(0.128)	0.252
Macmillan	0.290	(0.103)***	0.256	(0.144)*	0.273
Penguin Random House	0.339	(0.045)***	0.317	(0.074)***	0.328
Simon & Schuster	0.341	(0.059)***	0.376	(0.064)***	0.359
Average	0.325		0.289		0.307

Notes: Bootstrapped standard errors shown in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Estimates are for specification (E) in Table 7.

Table 9 reports the implied margins for the bargaining model estimates in column (E) of Table 7. The table reports figures for both the wholesale and agency models. The table shows that the publishers' margins went up for most of the publishers when switching to agency contracts. On the other hand, the retailers' margins went down as a result of the switch, although this mostly reflects the negative marginal cost estimates during the wholesale period. Note that these results are consistent with the theoretical model of Section 2—as illustrated in Figure 2, upstream profits are higher under agency, whereas downstream profits are higher under the wholesale model.

Table 9: Prices, margins and market shares

	Retail price	Wholesale price	Margin	Marginal cost	Market share
Panel A. Wholesale period					
<i>Retailers</i>					
Amazon	8.72	8.62	6.43	-6.33	0.83
Barnes & Noble	9.91	8.62	6.19	-4.89	0.17
<i>Publishers</i>					
Hachette	10.47	9.16	1.67	7.49	0.24
HarperCollins	9.97	9.09	3.32	5.77	0.09
Macmillan	7.53	6.88	2.13	4.75	0.02
Penguin Random House	8.86	8.40	2.83	5.58	0.53
Simon & Schuster	10.79	9.76	1.35	8.41	0.12
Panel B. Agency period					
<i>Retailers</i>					
Amazon	11.35	7.55	3.16	0.64	0.74
Barnes & Noble	11.11	7.71	3.10	0.31	0.26
<i>Publishers</i>					
Hachette	12.63	8.53	2.75	5.79	0.21
HarperCollins	10.56	7.86	2.83	5.03	0.11
Macmillan	10.69	7.73	2.75	4.98	0.10
Penguin Random House	10.89	7.30	3.02	4.23	0.46
Simon & Schuster	11.98	7.70	2.47	5.23	0.12

Notes: Estimates are for specification (E) in Table 7.

The negative marginal cost estimates for the retailers during the wholesale period are a direct result of the relatively large retail margin estimates. Note that these margins are estimated using the demand side only and do not rely on any assumption we have been making regarding the supply side. Because the difference between the retail price and wholesale price is relatively small (i.e., on average 10 cents for Amazon), the only way these large retail margins can be rationalized is by negative downstream marginal costs. Put differently, despite having sufficient market power to raise prices, Amazon was intentionally setting low prices during the wholesale period. Possible explanations that are consistent with these findings are the use of a loss-leader strategy, or a strategy of customer acquisition in Amazon's ecosystem (see Section 5.3 of De los Santos and Wildenbeest, 2017, for a more detailed discussion of Amazon's pricing strategies in the e-book market). Under these interpretations, negative marginal cost estimates represent the value to the retailers of bringing in consumers beyond selling the e-book. Note that if we ignore this part of the margin and purely look at the difference between the retail and wholesale price during the wholesale period and compare this to the margins during the agency period, it is clear that the retailers, like the publishers, did benefit from moving to agency in terms of actual profits. During the initial agency period that started in 2010, the e-book market was much less mature than during

the second transition, which may explain why Amazon was objecting to the first agency transition, while was willing to go along in 2014.

6 Counterfactual Analysis of the Most Favored Nations Clause

The settlements between DOJ and the Big Five publishers in 2012 explicitly banned the use of MFN clauses for a period of five years. In this section we study what happens to agency prices when MFN clauses are reinstated, which starting in 2017 is again a possibility. According to DOJ, the MFN clauses that were used during the initial switch to agency contracts in 2010 were essential for making the entire industry shift towards agency agreements, with the switch from wholesale to agency leading to higher consumer prices. Even though most of the industry is currently again using the agency model, and MFN was not instrumental for making the switch to this second period of agency pricing, this does not mean that MFN is unlikely to have a further impact on pricing once permitted again. The reason for this is that MFN guarantees a retailer who prefers a higher royalty that if it raises the royalty for one publisher, the retail price will remain the same relative to other retailers. This encourages retailers to increase royalties, which results in higher retail prices (see also Johnson, 2017).

Table 10: Prices, royalties and margins MFN

	No MFN			MFN		
	Price	Royalty	Margin	Price	Royalty	Margin
<i>Retailers</i>						
Amazon	11.12	0.325	3.16	11.44	0.333	3.36
Barnes & Noble	11.12	0.289	3.06	11.44	0.325	3.59
<i>Publishers</i>						
Hachette	12.99	0.324	2.71	13.47	0.364	2.49
HarperCollins	10.26	0.252	2.75	10.41	0.260	2.76
Macmillan	11.49	0.273	2.74	11.52	0.276	2.74
Penguin Random House	10.60	0.328	2.98	10.92	0.345	3.02
Simon & Schuster	12.05	0.359	2.45	12.39	0.399	2.19

Notes: Estimates are based on the bargaining model estimates. The reported results only include titles which are sold by both stores in the last week of the sample.

To simulate what happens to retail prices when MFN agreements are used, we use the marginal cost estimates from specification (E) of Table 7 to simulate equilibrium prices and royalties, using the restriction that the same title should have the same prices at both Amazon and Barnes & Noble. This restriction will not only affect the pricing FOC, but will also affect the royalty FOC, assuming that contracts between publishers and retailers are renegotiated. The implications of using MFN

on prices, royalties, and margins are shown in Table 10. The table shows that royalties increase across the board, with Barnes & Noble seeing the biggest changes among the retailers and Hachette and Simon & Schuster among the publishers. Even though retail prices go up as well, by about three percent on average, this price increase is not enough to prevent margins from going down for about half the publishers. Moreover, there is a lot variation in price changes. Although 75 percent of price changes is within -1 and 5 percent, non-fiction books saw a sales-weighted average price increase of 8.7 percent, whereas the average sales-weighted price change for fiction books is only 2.2 percent. The market is typically less competitive for non-fiction books, which may explain the large difference in price effects across book categories.

7 Conclusions

In this paper we have studied the effects of the transition from wholesale contracts to agency contracts in the e-book market that occurred in the period 2014-2015. Using a difference-in-differences analysis, we have shown that prices went up by 16 percent following the switch at Amazon, but went down 2 percent at Barnes & Noble. We have shown in a simple theoretical model that if an upstream and downstream firm are bargaining over input prices, it will depend on the relative bargaining power whether prices will increase or decrease when switching from wholesale agreements to agency agreements.

Our structural model extends this theoretical model to allow for competition among publishers and retailers, multi-product firms, and nested logit demand. We have shown how to estimate this model using sales rank data, prices, and book characteristics. Estimates of the bargaining model have shown that the bargaining power is on average about equally distributed between retailers and publishers, although there are large differences between retailer-publisher pairs. Moreover, the bargaining model better fits the data than a model in which input prices are determined using take-it-or-leave-it contracts. The results from a counterfactual analysis in which we reinstate MFN clauses lead to changes in consumer prices of about three percent, with changes of up to nine percent for non-fiction books.

References

- Berry, Steven, James Levinsohn, and Ariel Pakes: “Automobile Prices in Market Equilibrium,” *Econometrica* 63, 841–890, 1995.
- Chevalier, Judith and Austan Goolsbee: “Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com,” *Quantitative Marketing and Economics* 1, 203–222, 2003.
- Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee: ““Nash-in-Nash” Bargaining: A Microfoundation for Applied Work,” Forthcoming in the *Journal of Political Economy*, 2017.
- Crawford, Gregory S. and Ali Yurukoglu: “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review* 102, 643–685, 2012.
- De los Santos, Babur and Matthijs R. Wildenbeest: “E-book Pricing and Vertical Restraints,” *Quantitative Marketing and Economics* 15, 85–122, 2017.
- Dobson, Paul W. and Michael Waterson: “Countervailing Power and Consumer Prices,” *Economic Journal* 107, 418–430, 1997.
- Draganska, Michaela, Daniel Klapper, and Sofia B. Villas-Boas: “A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel,” *Marketing Science* 29, 57–74, 2010.
- Foros, Øystein, Hans Jarle Kind, and Greg Shaffer: “Apple’s Agency Model and the Role of Most-Favored-Nation Clauses,” *RAND Journal of Economics* 48, 673–703, 2017.
- Gaudin, Germain and Alexander White: “On the Antitrust Economics of the Electronic Books Industry,” Mimeo, 2014.
- Gowrisankaran, Gautam, Aviv Nevo, and Robert Town, “Mergers When Prices Are Negotiated: Evidence from the Hospital Industry,” *American Economic Review* 105, 172–203, 2015.
- Grennan, Matthew: “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review* 103, 145–177, 2013.
- Ho, Kate, and Robin S. Lee: “Insurer Competition in Health Care Markets,” *Econometrica* 85, 379–417, 2017.
- Horn, Henrick and Asher Wolinsky: “Bilateral Monopolies and Incentives for Merger,” *RAND Journal of Economics* 19, 408–419, 1988.

- Johnson, Justin: “The Agency and Wholesale Models in Electronic Content Markets,” Mimeo, 2018.
- Johnson, Justin: “The Agency Model and MFN Clauses,” *Review of Economic Studies* 84, 1151–1185, 2017.
- Li, Hui: “Multichannel Management with Digital Products: How E-Books Affect Online and Offline Print Book Sales,” Mimeo, 2015.
- McAfee, R. Preston and Marius Schwartz: “Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity,” *American Economic Review* 84, 210–230, 1994.
- Nair, Harikesh: “Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video-Games,” *Quantitative Marketing and Economics* 5, 239–292, 2007.
- Nash, John F.: “The Bargaining Problem,” *Econometrica* 18, 155–162, 1950.
- Reimers, Imke and Joel Waldfogel: “Throwing the Books at Them: Amazon’s Puzzling Long-Run Pricing Strategy,” *Southern Economic Journal* 83, 869–885, 2017.
- Shiller, Benjamin R.: “Digital Distribution and the Prohibition of Resale Markets for Information Goods,” *Quantitative Marketing and Economics* 11, 403–435, 2013.
- Villas-Boas, J. Miguel and Russell S. Winer: “Endogeneity in Brand Choice Models,” *Management Science* 45, 1324–1338, 1999.
- Vuong, Quang H.: “Likelihood Ratio Tests for Model Selection and non-nested Hypotheses,” *Econometrica* 57, 307–333, 1989.

A Proof of Proposition 1

Proof. First assume the upstream firm has all the bargaining power, i.e., $\lambda = 1$. The first-order condition for the wholesale model is found by plugging in $\lambda = 1$ in equation (3), which gives

$$p - c^U - c^D = \phi(p) (1 + (1 - \phi'(p))).$$

Note that this condition corresponds to the take-it-or-leave-it case analyzed in Johnson (2017). Using equation (5), the first-order condition for the agency model for $\lambda = 1$ is

$$p - c^U - c^D = \phi(p) \left(\frac{c^U}{p - \phi(p)} \right).$$

To show that the agency price p^a must be lower than the wholesale p^w , suppose that $p^a \geq p^w$. This means that

$$p^a - c^U - c^D \geq \lambda(p^a) (1 + (1 - \lambda'(p^a))).$$

According to the first-order condition for the agency model,

$$p^a - c^U - c^D = \lambda(p^a) \left(\frac{c^U}{p^a - \lambda(p^a)} \right).$$

We get a contradiction if

$$\lambda(p^a) (1 + (1 - \lambda'(p^a))) > \lambda(p^a) \left(\frac{c^U}{p^a - \lambda(p^a)} \right),$$

which can be reduced to

$$1 + (1 - \lambda'(p^a)) > \frac{c^U}{p^a - \lambda(p^a)}.$$

According to the pricing first-order condition under agency, $1 - r = c^U / (p - \phi(p))$, which, given that the royalty rate is at least zero, implies that the right-hand side of this inequality is less or equal to one. By assumption $\phi'(p) < 1$, which means the the left-hand side is bigger than one. This establishes a contradiction, which means that $p^a < p^w$ for $\lambda = 1$.

Next, consider the case in which the downstream firm has all the bargaining power, i.e., $\lambda = 0$. The first-order condition for the wholesale model for this case is

$$p - c^U - c^D = \phi(p).$$

Likewise, the first-order condition for the agency model is

$$p - c^U - c^D = \phi(p) \left(1 + c^U \frac{p(1 - \phi'(p))}{(p - \phi(p))^2} \right).$$

Note that this condition corresponds to the agency take-it-or-leave-it case analyze in Johnson (2017). Proceeding again by contradiction, suppose $p^a \leq p^w$. This means that

$$p^a - c^U - c^D \leq \lambda(p^a)$$

and, using the agency first-order condition,

$$p^a - c^U - c^D = \lambda(p^a) \left(1 + c^U \frac{p^a(1 - \lambda'(p^a))}{(p^a - \lambda(p^a))^2} \right).$$

Therefore, we get a contradiction if

$$\lambda(p^a) < \lambda(p^a) \left(1 + c^U \frac{p^a(1 - \lambda'(p^a))}{(p^a - \lambda(p^a))^2} \right).$$

Since $\phi'(p) < 1$, all terms between brackets on the right hand side of this inequality are positive, which means that the term between brackets exceeds one, and the inequality holds. This establishes a contradiction, which means that $p^a > p^w$ for $\lambda = 0$. ■

B Price derivatives

Wholesale model

The total derivative $d\pi_k^U/dw_j$ is given by

$$\frac{d\pi_k^U}{dw_j} = \frac{\partial \pi_k^U}{\partial w_j} + \sum_{k=1}^N \frac{\partial \pi_j^U}{\partial p_k} \frac{\partial p_k^*}{\partial w_j}, \quad (\text{A1})$$

where

$$\frac{\partial \pi_k^U}{\partial w_j} = \begin{cases} s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases}$$

The derivative $\partial\pi_j^U/\partial p_k$ is given by

$$\frac{\partial\pi_j^U}{\partial p_k} = \begin{cases} m_j^U \alpha s_j (1 - s_j + \gamma(1 - s_{j|g})) & \text{if } k = j, \\ -m_j^U \alpha s_k (s_j + \gamma s_{j|g}) & \text{if } k \neq j \text{ and in same nest,} \\ -m_j^U \alpha s_k s_j & \text{if } k \neq j \text{ and not in same nest.} \end{cases}$$

Similarly, the total derivative $d\pi_k^D/dw_j$ is given by

$$\frac{\partial\pi_j^D}{\partial w_j} = \frac{\partial\pi_k^D}{\partial w_j} + \sum_{l=1}^N \frac{\partial\pi_j^D}{\partial p_l} \frac{\partial p_l^*}{\partial w_j}, \quad (\text{A2})$$

where

$$\frac{\partial\pi_k^D}{\partial w_j} = \begin{cases} -s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases}$$

The derivative $\partial\pi_j^D/\partial p_k$ is given by

$$\frac{\partial\pi_j^D}{\partial p_k} = \begin{cases} -s_j + m_j^D \alpha s_j (1 - s_j + \gamma(1 - s_{j|g})) & \text{if } k = j, \\ m_j^D \alpha s_k (s_j + \gamma s_{j|g}) & \text{if } k \neq j \text{ and in same nest,} \\ m_j^D \alpha s_k s_j & \text{if } k \neq j \text{ and not in same nest.} \end{cases}$$

The price derivatives $\partial p_l^*/\partial w_j$ are derived by totally differentiating the retail price-first order conditions in equation (9). The solution is

$$p_{ww}^* = [\pi_{pp}^D]^{-1} [-\pi_{pw}^D].$$

The (k, l) th element of π_{pp}^D is given by

$$\pi_{pp}^D = T^D(k, l) \frac{\partial^2 \pi_j^D}{\partial p_k \partial p_l}.$$

Straightforward calculations yield the following expression for the derivatives on the right-hand

side of this equation:

$$\frac{\partial^2 \pi_j^D}{\partial p_k \partial p_l} = \begin{cases} 2 \frac{\partial s_j}{\partial p_j} + m_j^D \frac{\partial^2 s_j}{\partial p_j \partial p_j}, & \text{if } j = k = l, \\ \frac{\partial s_j}{\partial p_l} + m_j^D \frac{\partial^2 s_j}{\partial p_j \partial p_l}, & \text{if } j = k \neq l, \\ m_j^D \frac{\partial^2 s_j}{\partial p_k \partial p_k}, & \text{if } j \neq k = l, \\ \frac{\partial s_j}{\partial p_k} + m_j^D \frac{\partial^2 s_j}{\partial p_k \partial p_j}, & \text{if } j \neq k, l = j, \\ m_j^D \frac{\partial^2 s_j}{\partial p_k \partial p_l}, & \text{if } j \neq k, j \neq l, k \neq l. \end{cases}$$

The own-price and cross-price derivatives are given in equation (10). The second derivatives are given by

$$\frac{\partial^2 s_j}{\partial p_k \partial p_l} = \begin{cases} \alpha (1 - 2s_j + \gamma(1 - s_{j|g})) \frac{\partial s_j}{\partial p_j} - \alpha \gamma s_j \frac{\partial s_{j|g}}{\partial p_j} & \text{if } j = k = l, \\ \alpha (1 - 2s_j + \gamma(1 - s_{j|g})) \frac{\partial s_j}{\partial p_l} - \alpha \gamma s_j \frac{\partial s_{j|g}}{\partial p_l} & \text{if } j = k \neq l, \\ -\alpha \left(s_j \frac{\partial s_k}{\partial p_k} + s_k \frac{\partial s_j}{\partial p_k} \right) - \alpha \gamma \left(\frac{\partial s_k}{\partial p_k} s_{j|g} + \frac{\partial s_{j|g}}{\partial p_k} s_k \right) & \text{if } j \neq k = l, \\ -\alpha \left(s_j \frac{\partial s_k}{\partial p_j} + s_k \frac{\partial s_j}{\partial p_j} \right) - \alpha \gamma \left(\frac{\partial s_k}{\partial p_j} s_{j|g} + \frac{\partial s_{j|g}}{\partial p_j} s_k \right) & \text{if } j \neq k, l = j, \\ -\alpha \left(s_j \frac{\partial s_k}{\partial p_l} + s_k \frac{\partial s_j}{\partial p_l} \right) - \alpha \gamma \left(\frac{\partial s_k}{\partial p_l} s_{j|g} + \frac{\partial s_{j|g}}{\partial p_l} s_k \right) & \text{if } j \neq k, j \neq l, k \neq l. \end{cases}$$

The within-group price derivatives are given by

$$\frac{\partial s_{j|g}}{\partial p_k} = \begin{cases} \alpha(\gamma + 1)s_{j|g}(1 - s_{j|g}) & \text{if } j = k, \\ -\alpha(\gamma + 1)s_{k|g}s_{j|g} & \text{if } j \neq k \text{ and in same nest,} \\ 0 & \text{if } j \neq k \text{ and not in same nest.} \end{cases}$$

The (k, l) th element of π_{pw}^D is given by

$$\pi_{pw}^D = T^D(k, l) \frac{\partial^2 \pi_j^D}{\partial p_k \partial w_l}.$$

Straightforward calculations yield the following expression for the derivatives on the right-hand side of this equation:

$$\frac{\partial^2 \pi_j^D}{\partial p_k \partial w_l} = \begin{cases} -\alpha s_j (1 - s_j + \gamma(1 - s_{j|g})) & \text{if } j = k = l, \\ \alpha s_k (s_j + \gamma s_{j|g}) & \text{if } j \neq k, j = l \text{ and in same nest} \\ \alpha s_k s_j & \text{if } j \neq k, j = l \text{ and not in same nest} \\ 0 & \text{otherwise.} \end{cases}$$

Agency model

The total derivative $d\pi_k^U/dr_j$ is given by

$$\frac{d\pi_k^U}{dr_j} = \frac{\partial\pi_k^U}{\partial r_j} + \sum_{k=1}^N \frac{\partial\pi_j^U}{\partial p_k} \frac{\partial p_k^*}{\partial r_j},$$

where

$$\frac{\partial\pi_k^U}{\partial r_j} = \begin{cases} -p_k s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases}$$

The derivative $\partial\pi_j^U/\partial p_k$ is given by

$$\frac{\partial\pi_j^U}{\partial p_k} = \begin{cases} (1 - r_j)s_j + m_j^U \alpha s_j (1 - s_j + \gamma(1 - s_{j|g})) & \text{if } k = j, \\ -m_j^U \alpha s_k (s_j + \gamma s_{j|g}) & \text{if } k \neq j \text{ and in same nest,} \\ -m_j^U \alpha s_k s_j & \text{if } k \neq j \text{ and not in same nest.} \end{cases}$$

Similarly, the total derivative $d\pi_k^D/dr_j$ is given by

$$\frac{\partial\pi_j^D}{\partial r_j} = \frac{\partial\pi_k^D}{\partial r_j} + \sum_{l=1}^N \frac{\partial\pi_j^D}{\partial p_l} \frac{\partial p_l^*}{\partial r_j}, \quad (\text{A3})$$

where

$$\frac{\partial\pi_k^D}{\partial r_j} = \begin{cases} p_k s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases}$$

The derivative $\partial\pi_j^D/\partial p_k$ is given by

$$\frac{\partial\pi_j^D}{\partial p_k} = \begin{cases} r_j s_j + m_j^D \alpha s_j (1 - s_j + \gamma(1 - s_{j|g})) & \text{if } k = j, \\ m_j^D \alpha s_k (s_j + \gamma s_{j|g}) & \text{if } k \neq j \text{ and in same nest,} \\ m_j^D \alpha s_k s_j & \text{if } k \neq j \text{ and not in same nest.} \end{cases}$$

The price derivatives $\partial p_l^*/\partial r_j$ are derived by totally differentiating the retail price-first order conditions in equation (15). The solution is

$$p_{rr}^* = [\pi_{pp}^U]^{-1} [-\pi_{pr}^U].$$

The (k, l) th element of π_{pp}^U is given by

$$\pi_{pp}^U = T^U(k, l) \frac{\partial^2 \pi_j^U}{\partial p_k \partial p_l}.$$

Straightforward calculations yield the following expression for the derivatives on the right-hand side of this equation:

$$\frac{\partial^2 \pi_j^U}{\partial p_k \partial p_l} = \begin{cases} 2(1 - r_j) \frac{\partial s_j}{\partial p_j} + m_j^U \frac{\partial^2 s_j}{\partial p_j \partial p_j}, & \text{if } j = k = l, \\ (1 - r_j) \frac{\partial s_j}{\partial p_l} + m_j^U \frac{\partial^2 s_j}{\partial p_j \partial p_l}, & \text{if } j = k \neq l, \\ m_j^U \frac{\partial^2 s_j}{\partial p_k \partial p_k}, & \text{if } j \neq k = l, \\ (1 - r_j) \frac{\partial s_j}{\partial p_k} + m_j^U \frac{\partial^2 s_j}{\partial p_k \partial p_j}, & \text{if } j \neq k, l = j, \\ m_j^U \frac{\partial^2 s_j}{\partial p_k \partial p_l}, & \text{if } j \neq k, j \neq l, k \neq l. \end{cases}$$

The (k, l) th element of π_{pr}^U is given by

$$\pi_{pr}^U = T^U(k, l) \frac{\partial^2 \pi_j^U}{\partial p_k \partial r_l}.$$

Straightforward calculations yield the following expression for the derivatives on the right-hand side of this equation:

$$\frac{\partial^2 \pi_j^U}{\partial p_k \partial r_l} = \begin{cases} -s_j [1 + \alpha p_j (1 - s_j + \gamma (1 - s_{j|g}))] & \text{if } j = k = l, \\ \alpha p_k (s_j s_k + \gamma s_k s_{j|g}) & \text{if } j \neq k, j = l \text{ and in same nest} \\ \alpha p_k s_j s_k & \text{if } j \neq k, j = l \text{ and not in same nest} \\ 0 & \text{otherwise.} \end{cases}$$

C Derivation m^U (wholesale) and m^D (agency)

Derivation m^U (wholesale)

Equation (14) relates upstream margins to downstream margins, which can be used to solve the upstream margins as a function of the downstream margins. First, rewrite equation (14) as

$$E_j^D (\pi^U - d^U) \frac{1 - \lambda}{\lambda} + \frac{\partial \pi^U}{\partial w_j} = 0, \quad (\text{A4})$$

where $E_j^D = (\pi^D - d^D)^{-1} (\partial\pi^D/\partial w_j)$. Using equation (A1) and taking into account the ownership structure, we can write $\partial\pi^U/\partial w_j$ as

$$\frac{\partial\pi^U}{\partial w_j} = s_j + \sum_{k \in \Omega^U} (m_k^U Z_{jk}^w),$$

which, using matrix notation, can be written as $s + (T^U Z^w)m^U$. Taking into account the ownership structure and using $\pi^U - d^U = (T^U S)m^U$, we can write the bargaining first-order condition in equation (A4) as

$$s + \left(T^U Z^w + E^D (T^U S) \frac{1-\lambda}{\lambda} \right) m^U = 0.$$

Solving for m^U gives

$$m^U = - \left(T^U Z^w + E^D (T^U S) \frac{1-\lambda}{\lambda} \right)^{-1} s.$$

To derive an expression for E^D , first note that we can write equation (A2) as

$$\frac{\partial\pi^D}{\partial w_j} = -s_j + \sum_{k \in \Omega^D} (B_{jk}^D + m_k^D Z_{jk}^w),$$

where $B_{jk}^{Dw} = -s_k (\partial p_k^*/\partial w_j)$. In matrix notation this is $-s + T^D B^{Dw} + (T^D Z^w)m^D$. Moreover, since $\pi^D - d^D = (T^D S)m^D$, we get

$$E^D = ((T^D S) m^D)^{-1} (-s + T^D B^{Dw} + (T^D Z^w) m^D).$$

Derivation m^D (agency)

Equation (16) relates upstream margins to downstream margins, which can be used to solve the downstream margins as a function of the upstream margins. First, rewrite equation (16) as

$$E_j^U (\pi^D - d^D) \frac{\lambda}{1-\lambda} + \frac{\partial\pi^D}{\partial r_j} = 0, \tag{A5}$$

where $E_j^U = (\pi^U - d^U)^{-1} (\partial\pi^U/\partial r_j)$. Using equation (A3) and taking into account the ownership structure, we can write $\partial\pi^D/\partial r_j$ as

$$\frac{\partial\pi^D}{\partial r_j} = A_j + \sum_{k \in \Omega^D} (B_{jk}^D + m_k^D Z_{jk}^r),$$

where $A_j = p_j s_j$ and $B_{jk}^D = r_k s_k (\partial p_k^* / \partial r_j)$. In matrix notation this is

$$A + T^D B^D + (T^D Z^r) m^D.$$

Taking into account the ownership structure and using $\pi^D - d^D = (T^D S) m^D$, we can write the bargaining first-order condition in equation (A5) as

$$A + T^D B^D + \left(T^D Z^r + E^U (T^D S) \frac{\lambda}{1 - \lambda} \right) m^D = 0.$$

Solving for m^D gives

$$m^D = - \left(T^D Z^r + E^U (T^D S) \frac{\lambda}{1 - \lambda} \right)^{-1} (A + T^D B^D).$$

Alternatively, solving for $L = \lambda / (1 - \lambda)$ gives

$$L = -(A + T^D B^D + (T^D Z^r) m^D)^{-1} (E^U (T^D S)).$$

D Additional Specifications DID Analysis

Table A1 gives the results for several alternative specifications. To see to what extent serial correlation is an issue we aggregate the daily prices into weekly and monthly price observations. As shown in the first two rows of the table, the estimates are very similar to those of the main specification for both Amazon and Barnes & Noble. The dataset also contains prices of non-Big Five publishers. These smaller publishers typically kept using wholesale contracts throughout the sample. Including other publishers as an additional control group slightly reduces the magnitude of the price change at Amazon and makes the effect insignificant at Barnes & Noble. Including print book prices has no effect on the estimates in comparison to the main estimates.

To analyze if the effects can be attributed to the switch in selling method and not to other shocks that may have occurred around the switching dates, we run a placebo test in which we replicate the analysis using prices for the printed version of the e-book. Since there was no change in selling method for printed titles, one would not expect to see a significant change in prices for those books. Table A2 shows the results for the main specification in which we do not make a distinction by publisher, as well as the results split out by publishers. The table shows that for print books sold at Amazon there is no overall effect. The effects by publishers is in most cases

Table A1: Robustness analysis

Specification	DID Estimator agency \times Big Five		<i>R</i> -squared	Number of observations
Panel A: Amazon				
Aggregated by week	0.145***	(0.009)	0.437	92,889
Aggregated by month	0.139***	(0.009)	0.466	21,545
Including other publishers	0.128***	(0.009)	0.374	886,047
Including print-book prices	0.151***	(0.010)	0.423	473,793
Panel B: Barnes & Noble				
Aggregated by week	-0.014***	(0.007)	0.275	94,654
Aggregated by month	-0.020***	(0.007)	0.280	21,940
Including other publishers	0.011	(0.007)	0.301	943,308
Including print-book prices	-0.019***	(0.008)	0.255	462,781

Notes: The table presents difference-in-differences coefficients estimates for different sample and control specifications. Dependent variable is $\ln(\text{price})$. The specifications include week fixed effects (month fixed effects when aggregated by month) and controls as the main specification. Standard errors (clustered by book) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

small or not significantly different from zero. However, Hachette is an exception: prices for print books went down by 11 percent around the time of the switch. This can be explained by some of the bargaining tactics that were used in the period before the switch, including selling books at list prices. Notice that we also find an effect for Hachette books sold at Barnes & Noble; however, the effect goes the other way. For the other publishers we do not find an effect at Barnes & Noble.

Table A2: Placebo tests: effect of switch to agency on print book prices

	Amazon		Barnes & Noble	
	Overall	By publisher	Overall	By publisher
Agency × Big Five	-0.006 (0.005)		0.012*** (0.004)	
Agency × Harper Collins		0.000 (0.010)		-0.002 (0.008)
Agency × Hachette		-0.108*** (0.012)		0.051*** (0.008)
Agency × Simon & Schuster		0.030*** (0.009)		0.005 (0.008)
Agency × Macmillan		0.025** (0.010)		-0.002 (0.008)
Agency × Penguin Random House		0.026 (0.021)		
<i>R</i> -squared	0.647	0.648	0.657	0.657
Number of observations	484,368	484,368	464,131	464,131

Notes: The table presents difference-in-differences coefficient estimates by publisher and for Big Five publishers using $\ln(\text{price})$ of print books as dependent variable. The specification includes switching interaction coefficients for each publisher, week fixed effects and controls as the main specification. Standard errors (clustered by book) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.