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

Babur De los Santos†
Daniel P. O'Brien‡
Matthijs R. Wildenbeest§

October 2021

Abstract

This paper examines the pricing implications of two types of vertical contracts under bargaining: wholesale contracts, where downstream firms set retail prices after negotiating wholesale prices, and agency contracts, where upstream firms set retail prices after negotiating sales royalties. We show that agency contracts can lead to higher or lower retail prices than wholesale contracts depending on the distribution of bargaining power. We propose a methodology to structurally estimate a model with either contract form under Nash-in-Nash bargaining. We apply our model to the e-book industry, which transitioned from wholesale to agency contracts after the expiration of a ban on agency contracting imposed in the antitrust settlement between U.S. Department of Justice and the major publishers. Using a unique dataset of e-book prices, we show that the transition to agency contracting increased Amazon prices substantially but had little effect on Barnes & Noble prices. We find that the assumption of Nash-in-Nash bargaining explains the data better than an assumption of take-it or leave-it input contracts. Counterfactual simulations indicate that reinstitution of most favored nation clauses, which were banned for five years in the 2012 settlement, would raise the prices of e-books by seven percent but would lower profits of the publishers and Amazon.

Keywords: e-books, agency agreements, vertical restraints, bargaining, most favored nation clause

JEL Classification: C14, D83, L13

*We thank Takanori Adachi, Hanna Halaburda, Julie Holland Mortimer, Ginger Jin, Justin Johnson, Huaxia Rui, Fiona Scott Morton, and Shen Zhang for their useful comments and suggestions. This paper has benefitted from presentations at Boston College, Einaudi Institute for Economics and Finance, European Commission (DG COMP), MIT, Ohio State University, Stanford University, Université Paris Sud, University of Cologne, University of Mannheim, University of Vienna, Wake Forest University, ASSA Annual Meeting, Berlin IO Day at Humboldt University, CEPR Conference on Applied Industrial Organisation, CES North America Conference, EARIE Conference, European University Institute PriTech Workshop, HAL White Antitrust Conference, IIOC Annual Meeting, $(IO)^2$, NBER Summer Institute, and NET Institute Conference. We gratefully acknowledge financial support by the NET Institute, http://www.NETinst.org. De los Santos is a senior economist at Amazon. The views expressed are those of the authors and do not necessarily reflect those of Amazon, Inc. The paper represents work of De los Santos prior to joining Amazon and uses publicly gathered information only.

†John E. Walker Department of Economics, Clemson University, E-mail: babur@clemson.edu.
‡Compass Lexecon and Microfoundations, Inc., E-mail: dobrien@microfoundations.com.
§Kelley School of Business, Indiana University and CEPR, E-mail: mwildenb@indiana.edu.

1 Introduction

In many retail markets, the distribution arrangement involves suppliers charging retailers wholesale prices and retailers setting final prices to consumers (the "wholesale" model). The wholesale model has been extensively-studied in the literature and forms the foundation for much of the economics of vertical contracting, particularly that which informs antitrust policy.[1] Another distribution arrangement that has received much less attention involves agency relationships where suppliers pay retailers sales royalties to distribute products at prices determined by suppliers (the "agency" model).[2] Agency arrangements are used in some conventional markets (e.g., newspapers sold at kiosks, insurance sold by independent agents), but they are especially prevalent in online markets.[3]

Agency arrangements raise interesting questions for both price theory and policy. Key questions include how the choice of pricing institution affects prices and profits. In a recent antitrust case, the Department of Justice (DOJ) alleged that Apple and major book publishers engineered a shift from wholesale to agency pricing in the market for e-books, and that this shift, in combination with retail price most-favored nation ("MFN") clauses, raised the prices of e-books. Empirical evidence (De los Santos and Wildenbeest, 2017) confirms the price increase. A natural question is whether the price increase was caused by the shift to the agency model, the MFN clause, or both.

Recent theoretical literature has begun to address this question, but the literature to date has a significant gap: it abstracts from bargaining, which is an important feature of many intermediate markets, including the e-book market. Johnson (2017) compares the wholesale and agency models under the assumption that input terms are established through take-it or leave-it offers by the entity that is not responsible for setting the downstream price.[4] Under a reasonably weak condition on demand, he finds that the agency model generates lower retail prices than the wholesale model. But suppose that instead of having the non price-setting firm making take-it or leave-it offers, the downstream buyer makes the offers instead. In the wholesale model, the buyer would set the wholesale price equal to upstream marginal cost and thereby eliminate double marginalization. In the agency model, by contrast, the buyer would set the royalty looking ahead to the impact on the upstream firm's pricing decision, and this would generally lead to a degree of double marginalization.[5] Thus, the comparison between wholesale and agency arrangements is sensitive to the distribution of bargaining power. Yet, the theoretical literature on agency

[1] For example, the wholesale model forms the basis for most of the discussion of antitrust treatment of vertical integration and restraints in leading industrial organization textbooks (e.g., Carlton and Perloff, 2004). Note that wholesale pricing is also commonly referred to as linear pricing.

[2] Johnson (2017) distinguishes two other pricing arrangements: the "franchise" model, where the suppliers collect sales royalties from retailers that set the retailer price, and the "consignment" model, where suppliers charge a wholesale price and also control retail prices. Our focus in this paper is on the wholesale and agency models.

[3] For example, third-party sellers on Amazon Marketplace (the "upstream" firms) set the retail price for their products, while Amazon (the "downstream" firm) receives a percentage of the revenue. Other examples include eBay Buy It Now and the Apple App Store.

[4] If the entity with all bargaining power also controls retail prices then it can achieve the vertical integrated outcome. Hence, take-it or leave-it offers in this context are assumed to be made by the entity without control of retail prices.

[5] A form of double marginalization arises in the agency model unless the upstream firm has zero marginal cost.

pricing (Gans, 2012; Gaudin and White, 2014; Abhishek, Jerath, and Zhang, 2015; Foros, Kind, and Shaffer, 2017; Johnson, 2017; Condorelli, Galeotti, and Skreta, 2018; Johnson, 2020) abstracts from bargaining. The literature on the wholesale model, in contrast, has focused extensively on bargaining and considers it to be a fundamental economic factor determining outcomes in many situations.[6]

In this paper we examine the relationship between agency contracts and retail prices when intermediate pricing terms are determined through bargaining, and we propose and estimate a structural model that allows examining both arrangements empirically. We begin in Section 2 by extending the bilateral monopoly models of wholesale and agency pricing in Johnson (2017) to allow for bargaining between the supplier and retailer. We show that agency contracts can lead to higher or lower retail prices depending upon the relative bargaining powers of the upstream and downstream firms. When the upstream firm has high bargaining power, the wholesale price is relatively high in the wholesale model but the royalty paid to the retailer is relatively low in the agency model. In the wholesale model, retailers pass the high input price on to consumers in the form of higher retail price. In the agency model, by contrast, low royalties give the supplier a larger share of the retail price and reduce double marginalization, leading to a lower price than in the wholesale case. The opposite is true when the downstream firm has high bargaining power. In this case, a low wholesale price in the wholesale model reduces double marginalization and leads to a low retail price, while a high royalty paid to the retailer in the agency model causes significant double marginalization and a high retail price. In summary, the retail price tends to be lower in either arrangement when the firm with high bargaining power also determines the retail price, as the price-setting firm has an incentive to establish input terms that mitigate double marginalization. This relationship between bargaining power and retail prices in the wholesale and agency models plays an important role in the identification strategy in our structural model, as we explain in more detail below.

In Section 3 we adapt the theoretical model to make it more amendable to estimation by allowing for multi-product firms and multiple suppliers and retailers, using a logit demand structure. Following recent literature, we use the "Nash-in-Nash" solution to model bargaining.[7] In this framework each pair of firms reaches an asymmetric Nash bargaining solution while taking the terms negotiated by other pairs as given. We extend this literature, which has focused on wholesale

[6]Examples include Horn and Wolinsky (1988) [mergers]; Dobson and Waterson (1997) [countervailing power], O'Brien and Shaffer (2005) [mergers], O'Brien (2014) [price discrimination], Crawford, Lee, Whinston, and Yurukoglu (2018) [vertical integration], and Ho and Lee (2019) [hospital and health insurance pricing].

[7]The "Nash-in-Nash" solution concept was first applied in the wholesale model by Horn and Wolinsky (1988) to study mergers and by Davidson (1988) to study multi-unit bargaining in labor markets (neither set of authors used the term "Nash-in-Nash," which appears to have arisen in the folklore). O'Brien (1989; 2014) provides non-cooperative foundations for this solution concept based on an extension of Rubinstein's (1982) bargaining model to environments with upstream monopoly, downstream oligopoly, and linear input pricing. The extension to the case of multiple upstream firms is straightforward. Collard-Wexler, Gowrisankaran, and Lee (2019) provide a non-cooperative foundation for the Nash-in-Nash solution concept for bargaining that is over fixed transfers that do not affect downstream firms' pricing decisions. Our model is different because we allow wholesale prices and sales royalties to affect downstream pricing decisions.

pricing contracts, to allow for agency contracts between upstream and downstream firms. Moreover, when deriving the bargaining equilibrium for both types of vertical contracts, we let firms take into account retail price reactions to input prices.

We apply our model to the e-book industry. 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 vertical contracts 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 and how these changes affected retail prices. E-books, similar to printed books, were initially sold using the wholesale model. In this period, Amazon pursued a low price strategy for e-books (e.g. $9.99 for newly released e-books). As De los Santos and Wildenbeest (2017) document, publishers were against this pricing arrangement because they believed that it cannibalized profitable hardcover sales, eroded consumer perceptions of the value of a book, and would eventually lead to lower wholesale prices.

With the introduction of the iPad in 2010, major publishers negotiated agency contracts with Apple to offer e-books for sale in Apple's new iBookstore. The terms of the agency contracts with Apple, particularly the MFN clause that required publishers to match lower retail prices at other retailers, prompted five of the six major publishers (the "Big Six") to compel the adoption of agency contracts on Amazon. The industry adoption of agency contracts and the MFN led to higher prices for e-books. In 2012 the Department of Justice sued Apple and five of the Big Six publishers for conspiring to raise e-book prices. All five publishers that were sued settled the lawsuit and agreed to a two-year ban on publisher-set prices, which effectively meant a 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.

The expiration of the two-year ban on agency pricing meant that, by the end of 2014, publishers could again negotiate agency contracts with Amazon and most other retailers that would allow them to control retail prices directly—because Apple had not settled, it was subject to a separate court injunction that banned the use of agency contracts for a longer period. 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 for Hachette titles. These negotiations took over six months, were extensively covered by the media, and they involved public pressure by some of Hachette's bestselling authors. Despite the lengthy bargaining period, by the end of 2015, all of the major publishers had returned to agency contracts with Amazon with publisher-set prices. In Section 4 we also investigate the effect on retail prices following this latest shift towards agency contracts, using price data for e-books sold at Amazon and Barnes & Noble in the period 2014-2015. We exploit the variation in the timing of the implementation of the new agency contracts to estimate the change in retail prices resulting from the switch to the new agency arrangements using a difference-in-differences approach. Our findings indicate that, on average, Amazon prices increased 14 percent and Barnes & Noble prices decreased 2 percent. The estimates also show substantial heterogeneity in price effects across publishers. These findings are difficult to explain using take-it or leave-it contracting models, but they are consistent with a bargaining model in which publishers have different bargaining weights.[8]

In Section 5 we discuss how to structurally estimate the empirical model developed in Section 3 in light of the industry transitions discussed in Section 4, and we present estimates of the bargaining model by jointly estimating demand and supply. The extent to which prices change following a shift to agency contracts is related to the relative bargaining power of the firms involved. To fully exploit this mechanism for identification and estimation, we use data from both before and after the latest switch to agency pricing. Our goal is to obtain estimates of the demand side parameters as well as supply side parameters, which includes a bargaining parameter for each publisher-retailer pair. Although the supply model varies between wholesale and agency contracts, we assume the bargaining parameters do not change when switching. This assumption abstracts away from the initial stage of bargaining between publishers and retailers over the type of contracts (wholesale versus agency) in order to make the model tractable. Then, for a given set of demand and supply parameters, we can use the pricing and bargaining first-order conditions for each model to solve for the margins of the upstream and downstream firms in both periods. We use the margins to back-out upstream marginal costs, which, assuming a log-linear relation between marginal cost and observable cost shifters, allows us to obtain estimates of both demand and supply side unobservables.

We use a covariance restriction approach to deal with price endogeneity in the demand equation. MacKay and Miller (2021) show there is link between the demand and supply side unobservables and the endogenous price coefficient, and that instrument-free identification of the price coefficient can be achieved by using a covariance restriction on the unobserved shocks. MacKay and Miller develop a three-stage estimator for the price coefficient in the simple logit model that is computationally trivial, assuming constant marginal cost and Bertrand competition. Since the bargaining parameters enter our supply side model nonlinearly, we use a GMM estimator that exploits cross-covariance restrictions to create additional moments that are necessary to estimate the additional nonlinear parameters, following the approach put forward by MacKay and Miller (2021) for the case of nonlinear parameters. We assume zero cross-covariance between the unobserved demand and cost shocks and use Monte Carlo experiments to show that this approach effectively deals with price endogeneity in small samples, while also allowing us to recover the nonlinear bargaining parameters of the supply side model.

According to our estimates, the price coefficient for the logit specification implies a median own-

[8] An important difference between the reduced form analysis in this paper and that in De los Santos and Wildenbeest (2017) is that we study transitions between the wholesale and agency model in the absence of an alleged conspiracy, whereas De los Santos and Wildenbeest analyze transitions that resulted from an alleged conspiracy involving Apple and competing publishers. In addition, retail price MFN clauses were not used during the period we study, allowing us to isolate the effect of agency pricing from the effect of the MFN. This is important, as the theoretical results of Johnson (2017) indicate that the MFN would have had a positive effect on prices.

price elasticity of around -1.8 for our main specification. The supply-side estimates suggest that the retailers have more bargaining power than the publishers. However, there are substantial differences in bargaining parameters between different retailer-publisher pairs, and Amazon generally has more bargaining power than Barnes & Noble. The estimates imply an agency royalty of 39 percent on average, which is higher than the thirty percent royalty that was common during the first agency period. It is important to note that these results rely on several simplifying assumptions regarding the bargaining process—to check for robustness with respect to the agency royalty, we have also estimated a specification in which the royalty is fixed to 30 percent and find less bargaining power for Amazon. Moreover, supply side estimates that are obtained using alternative values of the price coefficient show that we get similar bargaining parameter estimates for a range of price coefficients, suggesting our bargaining estimates do not rely critically on the covariance restrictions required by the MacKay and Miller (2021) approach. We compare the fit of the bargaining model to an alternative model with take-it or leave-it offers by the party that does not control retail prices. That is, we estimate a model in which retailers make take-it or leave-it royalty offers to publishers in the agency arrangement, and publishers make take-it or leave-it wholesale price offers to retailers in the wholesale arrangement. We find that the bargaining model gives a better fit to the data than the take-it or leave-it specification.

In Section 5 we also discuss the results of a counterfactual analysis where we use the estimates of the bargaining model to simulate the effect of retail price MFN clauses on retail prices. MFN clauses in this context are price-parity restrictions that guarantee that the same title is sold at the same price everywhere, as in the contracts used during the first agency period in the e-book industry. Price-parity clauses have been used by other online platforms in which agency contracts are used, such as online travel agencies, and even though U.S. courts have mostly upheld MFN clauses (Dennis, 1995), they have been under scrutiny by competition authorities around the world for their potential to reduce price competition.

The settlements between the DOJ and 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 alleged conspiracy. The role of MFN has been explored theoretically by Johnson (2017), who finds that it tends to raise retail prices. In line with this theoretical finding, our counterfactual simulations indicate that prices would increase an additional seven percent, on average, if retail price MFN clauses were added to the agency contracts. This finding is consistent with recent work by Mantovani, Piga, and Reggiani (2021), who analyze the price effects of laws in several European countries that banned the use of price-parity clauses by online travel agencies and find a significant price reductions in the medium run, especially for hotels affiliated with a chain. Our simulations also show that reinstatement of MFN clauses would lower profits of Amazon and the publishers, which could be a factor that explains why, as far as we know, MFN has not been adopted, despite the ban being lifted. We also compare the counterfactual predictions to the estimated price difference when using a difference-in-differences (DID) approach to compare prices during the first agency period (2010-2012) to prices during the current agency period. Although the DID estimates indicate prices were indeed higher during the first agency period, it is difficult to separate the price effects of MFN clauses from the price effects of the alleged collusion between the publishers in the first agency period. The counterfactual predictions based on the estimates of our structural model are obtained assuming no collusion, and therefore provide more direct estimates of the effect of MFN on agency prices than is possible without using our structural model.

Related Literature

As laid out by Johnson (2017), vertical arrangements between a retailer and supplier can be classified according to who sets the retail price (retailer vs. supplier) and the allocation of rents (revenue sharing vs. wholesale pricing), which leads to four possible business models: the wholesale model, the consignment model, the franchise model, and the agency model. As in the wholesale model, a wholesale (or linear) price is used in the consignment model, but the retail price is set by the upstream firm instead of the retailer. As shown by Johnson (2017), as long as the firm that sets the retail prices is not also setting the wholesale price, the equilibrium retail price will be the same for the two models. The difference between the franchise model and the agency model is that the retail price is set by the retailer in the franchise model and by the upstream firm in the agency model. An important focus of the empirical literature on franchising has been on agency-theoretic explanations for franchising such as moral hazard and risk sharing (see Lafontaine, 1992, and the references therein).

The empirical literature on the agency model is scant—one notable exception is Li and Moul (2015), who study the impact of a switch from wholesale to agency agreements using sales data for mobile phones in a Chinese department store. The focus of Li and Moul (2015) is on how the switch affected service provision—using a structural demand and supply model, they find that demand went up sharply when moving to agency contracts while prices remained relatively flat, which suggests customer service had improved following the switch. This shows that costly retailer effort might be more efficiently coordinated when the upstream firm sets prices (see also Conlon and Mortimer, 2021, for a model of retailer effort provision). Although costly retailer effort provision could play a role in the e-book market as well, a big difference is that because the retailers in our setting are operating online, the upstream firms cannot directly control the retailing environment, which makes it more difficult to coordinate prices and service efforts.

Our paper also fits into a broader empirical literature that studies the role of contracts in vertical markets. Villas-Boas (2007) develops a method to determine which vertical model fits the data best that only requires price and cost data. Part of this literature has focused on the efficiency of revenue-sharing, which, in addition to prices being set by the upstream firm, is an important feature of agency agreements. Mortimer (2008) studies the welfare effects of revenue-sharing contracts in the video rental industry, and finds that both upstream and downstream profits increase when revenue-sharing contracts are adopted. Note that revenue-sharing contracts can usually be written as a two-part tariff contract (see, for instance, Cachon and Lariviere, 2005)—Bonnet and Dubois (2015) and Hristakeva (2020) estimate structural models in which two-part tariff contracts are used to redistribute profits that can be estimated using limited data.

Our paper is also related to a growing empirical literature that use some variant of the Nash-in-Nash bargaining solution to estimate demand and supply models in oligopolistic markets (Draganska, Klapper, and Villas-Boas, 2010; Crawford and Yurukoglu, 2012; Grennan, 2013; Gowrisankaran, Nevo, and Town, 2015; Crawford, Lee, Whinston, and Yurukoglu, 2018; Ho and Lee, 2019; Donna, Pereira, Trindade, and Yoshida, 2021). We extend this literature, which has focused on wholesale pricing contracts, to allow for agency contracts between upstream and downstream firms. Moreover, our identification strategy for the bargaining parameters is based on the observation that the magnitude of retail price changes following a change from wholesale to agency contracts directly relates to how bargaining power is distributed across upstream and downstream firms, and by estimating the model for both agency and wholesale arrangements the bargaining parameters can be identified in a cleaner way than in most of the literature.

Related empirical work on the e-book market includes De los Santos and Wildenbeest (2017), Reimers and Waldfogel (2017), and Li (2021). De los Santos and Wildenbeest (2017) find that the switch from agency to wholesale following the ban on agency pricing in 2012 reduced retail prices by 18 percent at Amazon and 8 percent at Barnes & Noble. Reimers and Waldfogel (2017) find that e-books are priced below *static* profit maximizing levels. Li (2021) estimates a structural model where consumers choose how many books to buy, their format, and platform, and finds that over seventy percent of e-book sales come from cannibalization of print book sales. We refer to Gilbert (2015) for an overview of recent developments in the e-book industry and Baker (2018) for an overview of the lawsuit against Apple and the publishers that led to switch from the wholesale model to the agency model we study in this paper.

2 Vertical Bargaining Model

In this section we extend the bilateral monopoly models of wholesale and agency pricing in Johnson (2017) to allow for bargaining over input terms. Suppose there are two firms, an upstream firm U and a downstream firm D, that produce and sell a product to consumers at retail price p. Consumer demand is given by a continuously differentiable and strictly decreasing function $Q(p)$. Marginal cost is $c^U > 0$ for the upstream firm and $c^D \geq 0$ for the downstream firm. We consider two pricing structures, a wholesale arrangement and an agency arrangement. In the wholesale arrangement, firms first agree to a per-unit wholesale price to be paid by the downstream firm to the upstream firm when units of the product are sold, and then the downstream firm sets the retail price. In the agency model, firms first agree to an ad valorem (percent of price) royalty to be paid by the upstream firm to the downstream firm when units are sold, and then the upstream firm sets the retail price.

2.1 Wholesale Pricing

In the wholesale model, upstream and downstream profits are

$$\pi^U = (w - c^U)Q(p) \text{ and } \pi^D = (p - w - c^D)Q(p).$$

Given the wholesale price w, the downstream firm choose a price p to maximize its profits. The first-order condition is

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

where

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

is a measure of the sensitivity of demand to price. As in Johnson (2017), we assume that $\phi(p)$ and $\phi(p)(2 - \phi'(p))$ have slopes strictly less than 1.[9]

The wholesale price w is determined through asymmetric Nash bargaining (Nash, 1950) between the upstream and downstream firm. Let $p^*(w)$ solve equation (1). Assuming zero disagreement payoff, the Nash product is

$$NP = \left(\pi^U\right)^\lambda \left(\pi^D\right)^{1-\lambda},$$

where the profit functions are evaluated at $(w, p^*(w))$ and $\lambda \in [0, 1]$ is the a bargaining parameter identified with the upstream firm's bargaining weight. This weight is 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 and downstream firms.

The bargaining solution is found by maximizing the Nash product. The first order condition is

$$\lambda \pi^D \pi^{U'} + (1 - \lambda)\, \pi^U \pi^{D'} = 0, \tag{2}$$

where primes ordinarily indicate derivatives with respect to w. However, because $p^*(w)$ is monotonically increasing in w, it is possible to use the first order condition (1) to eliminate w from these profit expressions and express the Nash product as a function of the retail price p (as Johnson (2017) observed for the take-it or leave-it case). It is then possible (and simpler) to characterize the bargaining solution by maximizing the Nash product with respect to the retail price p. To this end, we substitute equation (1) into the profit expressions to express profits in terms of the retail price: $\pi^U = \left(p - \phi(p) - c^U - c^D\right) Q(p)$ and $\pi^D = \phi(p)Q(p)$. Substituting these expressions and

[9] As is well-known (Bagnoli and Bergstrom, 2005; Weyl and Fabinger, 2013), the sign of $\phi'(p)$ determines whether the demand function is log-concave ($\phi'(p) < 0$), log-convex ($\phi'(p) > 0$), or log-linear ($\phi'(p) = 0$). The assumption $\phi'(p) < 1$ ensures that the pass-through rate is positive. The assumption that $\phi(p)(2 - \phi'(p))$ has a slope less than 1 implies a unique solution to the pricing problem in the case where the upstream firm has all the bargaining power.

their derivatives into equation (2) gives

$$\begin{aligned}\lambda\phi(p)Q(p)\left[(1-\phi'(p))Q(p)+(p-\phi(p)-c^U-c^D)Q'(p)\right]\\+(1-\lambda)\,(p-\phi(p)-c^U-c^D)Q(p)\left[\phi'(p)Q(p)+\phi(p)Q'(p)\right]=0.\end{aligned}\tag{3}$$

Dividing both sides of equation (3) by $Q'(p)$ and rearranging gives an expression for the markup as a function of $\phi(p)$, $\phi'(p)$, and λ:

$$p-c^U-c^D=\phi(p)\left(\frac{\lambda+1-\phi'(p)}{\lambda+(1-\lambda)(1-\phi'(p))}\right).\tag{4}$$

2.2 Agency Pricing

In the agency model, upstream and downstream profits are

$$\pi^U=\left((1-r)p-c^U\right)Q(p)\ \text{ and }\ \pi^D=(rp-c^D)Q(p).\tag{5}$$ [10]

Given the royalty r, the upstream firm chooses p to maximize its profits. The first order condition is

$$(1-r)p-c^U=(1-r)\phi(p).\tag{6}$$

We can rewrite the first-order condition for price in equation (6) as

$$r=1-\frac{c^U}{p-\phi(p)}.\tag{7}$$

It is again helpful to substitute the first order condition (7) into the profits in (5) to express profits in terms of the retail price. After some algebra, this gives an upstream profit of

$$\pi^U=\frac{c^U}{p-\phi(p)}\phi(p)Q(p).\tag{8}$$

Downstream profit can be written as the difference between joint profit and upstream profit:

$$\begin{aligned}\pi^D&=\left(p-c^U-c^D\right)Q(p)-\pi^U\\&=\left(p-c^U-c^D\right)Q(p)-\frac{c^U}{p-\phi(p)}\phi(p)Q(p).\end{aligned}\tag{9}$$

The derivative of the upstream profit (8) with respect to p is

$$\pi^{U\prime}=\frac{c^U\left[\phi'(p)Q(p)+\phi(p)Q'(p)\right](p-\phi(p))-c^U\phi(p)Q(p)(1-\phi'(p))}{(p-\phi(p))^2},\tag{10}$$

[10] For brevity we use π^U and π^D to indicate profits in both regimes and will be clear whenever it might cause confusion.

and the derivative of the downstream profit (9) is

$$\pi^{D'} = Q(p) + (p - c^U - c^D)Q'(p) - \pi^{U'}. \tag{11}$$

Substituting the expressions in (8), (9), (10), and (11) into the bargaining first-order condition in (2), where the primes now indicate derivatives with respect to r, gives

$$\lambda \left[\left(p - \phi(p) \frac{c^U}{p - \phi(p)} - c^U - c^D \right) Q(p) \right] \pi^{U'} + (1 - \lambda) \left(\frac{c^U}{p - \phi(p)} \phi(p) Q(p) \right) \\ \times \left[Q(p) + (p - c^U - c^D)Q'(p) - \pi^{U'} \right] = 0. \tag{12}$$

Observe that

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

Dividing both sides of the bargaining first-order condition (12) by $Q'(p)$ and rearranging expresses the markup in the agency model as a function of $\phi(p)$, $\phi'(p)$, and λ:

$$p - c^U - c^D = \phi(p) \left(\frac{(1 - \lambda)(p - \phi(p))^2 + c^U p(1 - \phi'(p))}{(p - \phi(p))\big(p(1 - \lambda\phi'(p)) - \phi(p)(1 - \lambda)\big)} \right). \tag{13}$$

2.3 Comparison of Vertical Contracts

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

Proposition 1 *There exist critical bargaining parameters $\lambda^* \in (0, 1)$ and $\lambda^{**} \in [\lambda^*, 1)$ such that if the upstream firm's bargaining weight exceeds λ^{**}, the equilibrium retail price is higher under wholesale pricing than under agency pricing, and if the upstream firm's bargaining weight is less than λ^*, the opposite is true.*

The proof of Proposition 1 is in Appendix A. To illustrate this proposition, Figure 1 shows optimal retail prices and combined profits when demand has the constant-elasticity form $Q(p) = p^{-1/\kappa}$. In this case, the equilibrium price in the wholesale model is $p^w = (c^U + c^D)(1 + \kappa(\lambda - 1))/(\kappa - 1)^2$, and the equilibrium price in the agency model is $p^a = 2(c^U + c^D(1 - \kappa))/((1 - \kappa) \cdot (1 + \kappa(\lambda - 1)))$.

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. Retail prices are increasing in λ in the wholesale model—the more bargaining power the upstream firm has, the higher the negotiated wholesale price, with higher retail prices as a result. On the other hand, retail prices are decreasing in λ in the agency model as a better bargaining position for the upstream firm leads to lower royalties, which in turn reduces the double marginalization problem and leads to lower prices. Figure 1(a) also illustrates that whether retail prices are higher or lower under agency depends on the exact value of the bargaining parameter. In this example prices are higher under agency than under wholesale for

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

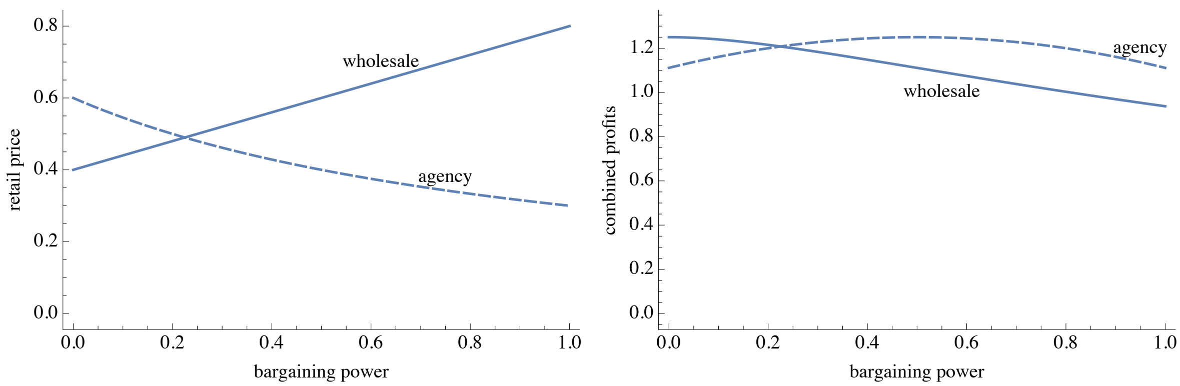

(a) Retail price (b) Combined profits

Notes: Retail price (figure a) and combined upstream and downstream profit (figure b) as a function of the bargaining weight for the wholesale model and agency model. Demand is $Q(p) = p^{-2}$ and $c^U = c^D = 0.1$.

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.[11]

Figure 1(b) shows the combined profits of the upstream and downstream firm as a function of the bargaining power parameter for each of the two models. For this particular example, the joint firm profits are maximized under the agency model when the firms share equal bargaining power. However, under the wholesale model 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 completely eliminates the double marginalization problem and maximizes joint firm profits.

Figure 2(a) compares the upstream firm's profit under the two types of vertical contracts for the same demand parameters, whereas 2(b) makes the same comparison for the downstream firm. In this example, the upstream firm always prefers agency pricing whereas the downstream firm prefers wholesale pricing. The opposite is true when firms use take-it or leave-it offers in which the party that does not control retail prices has all bilateral bargaining power, i.e., $\lambda = 1$ under wholesale pricing and $\lambda = 0$ under agency pricing. It follows that with take-it or leave-it offers, transitioning to agency means higher profits for the downstream firm and lower profits for the upstream firm (see also Proposition 3 of Johnson, 2017).

Note that in our bargaining model it is not always the case that the downstream firm prefers wholesale agreements. Figures 3(a) and 3(b) show profits as a function of the bargaining weight when the marginal cost for the upstream firm is 0.6 instead of 0.1—for intermediate values of the bargaining parameter, both firms now prefer agency pricing.

[11]This result is consistent with the conditions of Lemma 2 of Johnson (2017) for lower retail prices under the agency model compared to the wholesale model.

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

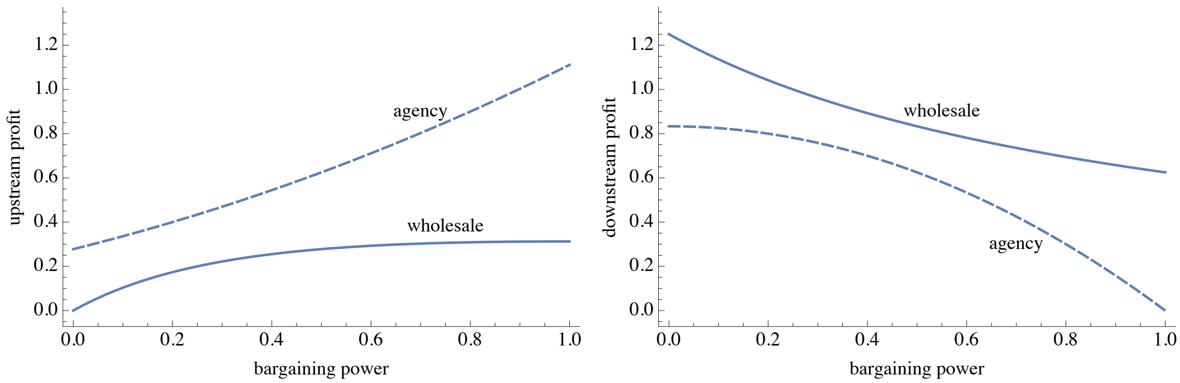

(a) Upstream profits (b) Downstream profits

Notes: Upstream profits (figure a) and downstream profits (figure b) as a function of the bargaining weight for the wholesale model and agency model. Demand is $Q(p) = p^{-2}$ and $c^U = c^D = 0.1$.

3 Empirical Model

To make the model amendable for estimation, we extend the model to allow for multiple upstream and downstream firms, as well as multi-product firms. In addition, we model consumer demand using a logit discrete choice framework. In this section, we derive the equilibrium conditions of the model.

3.1 Demand

We consider an industry with multiple upstream suppliers where each produces one or more goods and sells a set of these goods non-exclusively through multiple downstream retailers. Upstream producers can sell the same good through different retailers and retailers can sell goods of different suppliers. We define a product as a good-retailer pair. That is, a product is a specific good sold by a specific retailer. This means that product j sold by one retailer may be the same good as product k sold by another retailer. The idea is that different good-retailer pairs (different products) represent different points in product space. The utility consumer i derives from product j is given by

$$u_{ij} = \alpha \log(p_j) + x_j'\beta + \xi_j + \varepsilon_{ij}, \tag{14}$$

where p_j is the price of product j, x_j and ξ_j are observed and unobserved characteristics of product j, α and β are demand parameters, and ε_{ij} is a consumer-product specific utility shock. We allow for an outside option with utility $u_{i0} = \varepsilon_{ij}$. Assuming ε_{ij} follows a Type I Extreme Value distribution, and letting $\delta_j = \alpha \log(p_j) + x_j'\beta + \xi_j$, the market share of product j is

$$s_j(p) = \frac{\exp(\delta_j)}{1 + \sum_{k=1}^{N} \exp(\delta_k)}.$$

Figure 3: Upstream and downstream profits as a function of bargaining power ($c^D = 0.6$)

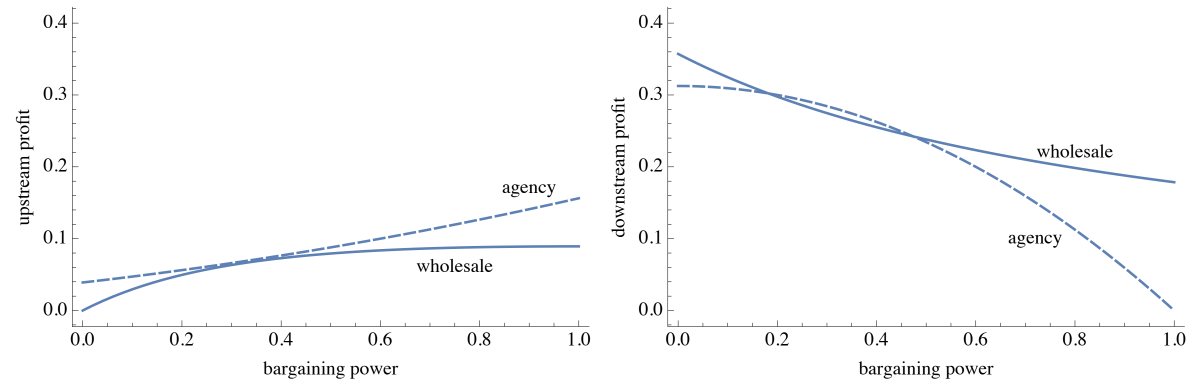

(a) Upstream profits (b) Downstream profits

Notes: Upstream profits (figure a) and downstream profits (figure b) as a function of the bargaining weight for the wholesale model and agency model. Demand is $Q(p) = p^{-2}$ and $c^U = 0.1$ and $c^D = 0.6$.

We consider two pricing structures: (1) the "wholesale model," where firms first agree to wholesale contracts for each book and the retailers set prices; and (2) the "agency model," where firms first agree to royalty contracts for each book and then the publishers set retail prices.

3.2 Wholesale Model

We model wholesale and retail pricing as a two stage game. In stage one, the supplier and retailer of each product j agree to a wholesale contract in which the retailer pays the supplier a wholesale price w_j for product j.[12] 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 downstream variable profit from selling product j is given by

$$\pi_j^D(p) = \left(p_j - w_j - c_j^D\right) s_j(p), \tag{15}$$

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

[12] The actual wholesale contracts that were used for e-books in the period 2012-2014 are typically called agency contracts because the retailer keeps a fraction r_j of the recommended price ρ_j for every product sold and the supplier receives the remainder. However, during this period, the retailer was free to set a discount, which means that these contracts are equivalent to wholesale agreements. To see this, observe that variable profit of the retailer from selling product j is

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

where $\rho_j - p_j$ reflects the discount the retailer may set. Note that the term $(1 - r_j)\rho_j$ is effectively a per-product wholesale price w_j paid to the supplier, which is the notation we use in this section. To distinguish between the two types of agency agreements (with and with without discounting), in the remainder of the paper we will use the term agency agreements only for agency agreements that do not allow the retailers to give discounts, whereas we will use the term wholesale agreements when discounting is allowed.

of product j. The upstream variable profit from selling j is

$$\pi_j^U(p) = \left(w_j - c_j^U\right) s_j(p), \tag{16}$$

where c_j^U is the upstream supplier's marginal cost of product j. The variable joint profit of the supplier 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 that sells products in the set Ω^D are given by

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

We assume a pure-strategy Nash equilibrium in retail prices. The first-order condition for product j is given by

$$s_j + \sum_{k \in \Omega^D} m_k^D \frac{\partial s_k}{\partial p_j} = 0, \tag{17}$$

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

$$\frac{\partial s_k}{\partial p_j} = \begin{cases} \alpha s_k \left(1 - s_k\right)/p_k & \text{if } k = j; \\ -\alpha s_j s_k / p_j & \text{if } k \neq j. \end{cases} \tag{18}$$

Upstream Market

We assume that wholesale prices are the outcome of a bilateral bargaining process between suppliers and retailers, and separate wholesale prices are chosen for each product. Overall profits of an upstream firm that sells products in the set Ω^U 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 and c_j^U is the upstream firm's marginal cost for product j.

We assume that wholesale prices are determined through simultaneous Nash bargaining ("Nash-in-Nash" bargaining) between the upstream and downstream firm associated with each product. The Nash product for downstream firm d and upstream firm u is

$$NP_{du}(w_{du}; w_{-du}) = \left(\pi^U - d_{du}^U\right)^{\lambda} \left(\pi^D - d_{du}^D\right)^{1-\lambda}, \tag{19}$$

where w_{du} is the vector of wholesale prices of the products associated with the upstream-downstream pair du, w_{-du} is the vector of wholesale prices for products associated with other upstream-downstream pairs, d^U_{du} and d^D_{du} are disagreement payoffs (discussed below), and $\lambda \in [0,1]$ is the bargaining weight of upstream firm u. Although we do not index λ to keep the notation simple, in our empirical application we allow λ to vary across supplier-retailer pairs. The Nash-in-Nash bargaining solution is the vector of wholesale prices for all products such that w_{du} maximizes NP_{du} for all upstream-downstream pairs du, given the results of the negotiations between other upstream-downstream pairs.

Following Horn and Wolinsky (1988) and Crawford and Yurukoglu (2012), we assume rival firms do not observe a bargaining breakdown, which means that rival firms do not adjust input or retail prices if negotiations between a specific upstream-downstream pair fail.[13] However, we do allow market shares to adjust in case of disagreement. Specifically, we assume disagreement payoffs for each du combination are given by

$$d^U_{du} = \sum_{k \in \Omega^U \setminus \{k \in du\}} m^U_k s^{-du}_k \qquad \text{and} \qquad d^D_{du} = \sum_{k \in \Omega^D \setminus \{k \in du\}} m^D_k s^{-du}_k .$$

In these expressions, s^{-du}_k is defined as the (counterfactual) market share for product k when products of du are not offered, i.e.,

$$s^{-du}_k = \frac{\exp(\delta_k)}{1 + \sum_{l \in \mathcal{J}_g \setminus \{l \in du\}} \exp(\delta_l)}. \tag{20}$$

So the disagreement payoff for the pair du consists of the profits for d from products not supplied by u and profits for u for products sold by other retailers that are not available at retailer d, considering that the demand for products $-du$ 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 (19) with respect to w_{du} equal to zero for *all* products that belong to the set of products offered by each du combination. Let j be such a product. The first-order condition for product j is

$$\lambda \left(\pi^U - d^U_{du}\right)^{\lambda - 1} \left(\pi^D - d^D_{du}\right)^{1-\lambda} \frac{\partial \pi^U}{\partial w_j} + (1-\lambda) \left(\pi^U - d^U_{du}\right)^{\lambda} \left(\pi^D - d^D_{du}\right)^{-\lambda} \frac{\partial \pi^D}{\partial w_j} = 0. \tag{21}$$

Since the Nash bargaining model defines an equilibrium payment for the set of products sold (and not just for an individual product j), in this first-order condition π^U and π^D are the *total* profits

[13] Iozzi and Valletti (2014) show in a setting with one upstream firm and two downstream firms that the disagreement payoff will depend on whether breakdowns are observable or not by the other downstream firm, which in turn may have implications for how input prices are determined. Crawford and Yurukoglu (2012) point out that an alternative model in which other firms renegotiate based on disagreeing pairs dropping out is computationally much more challenging, and therefore estimate the simpler model in which breakdowns are unobservable by other firms.

of the upstream and downstream firm. Equation (21) can be simplified to

$$\lambda \left(\pi^D - d^D_{du}\right) \frac{\partial \pi^U}{\partial w_j} + (1-\lambda) \left(\pi^U - d^U_{du}\right) \frac{\partial \pi^D}{\partial w_j} = 0. \tag{22}$$

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

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

where $d\pi^U_k/dw_j$ and $d\pi^D_k/dw_j$ are the *total* derivatives of π^U_k and π^D_k 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.[14] Condition (22) together with condition (17) yield the equilibrium wholesale input prices and retail prices.

3.3 Agency Model

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

$$\pi^D_j(p) = \left(r_j p_j - c^D_j\right) s_j(p).$$

The upstream variable profit from selling product j is

$$\pi^U_j(p) \quad = \quad \left((1-r_j)p_j - c^U_j\right) s_j(p).$$

The variable joint profit of the supplier and retailer associated with product j is $\pi^J_j = (p_j - c^D_j - c^U_j) s_j(p)$.

Upstream Market

In the agency model, the upstream supplier determines the retail price p_j. Overall profits of the supplier that sells products in the set Ω^U are given by

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

[14] An 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. In addition to treating the retail prices as fixed (which does not mean that retail price are independent of equilibrium retail prices), this literature also assumes that the derivative of the disagreement payoff with respect to input prices is zero. While we depart from this literature by assuming input prices are determined taking into account that retail prices may change in response (i.e., we allow for a non-zero derivative of retail prices with respect to input prices), we do keep the assumption that there are no disagreement payoff derivatives with respect to input prices. This implies that even if a firm is involved in multiple contract negotiations, it will treat each separately. As pointed out by Sheu and Taragin (2021), this assumption is important for maintaining tractability.

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

$$(1 - r_j)s_j + \sum_{k \in \Omega^U} m_k^U \frac{\partial s_k}{\partial p_j} = 0, \tag{23}$$

where $m_k^U = (1 - r_j)p_j - c_j^U$ is the upstream margin on product k and the derivative of the market share of product k with respect to p_j is given by equation (18).

Downstream Market

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

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

for each upstream-downstream pair du, where r_{du} and r_{-du} are vectors of royalties for the pairs du and $-du$, respectively. The bargaining first-order condition is found by setting the derivative of NP_{du} with respect to r_{du} equal to zero for all products that belong to the set of products offered by each du combination. The bargaining first-order condition for one such product—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

$$\lambda \left(\pi^D - d_{du}^D\right) \frac{\partial \pi^U}{\partial r_j} + (1 - \lambda) \left(\pi^U - d_{du}^U\right) \frac{\partial \pi^D}{\partial r_j} = 0. \tag{24}$$

Similar to the wholesale model, π^U and π^D are not just the profits for product j, but the total profits of the firms. 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 (24) together with condition (23) yield the equilibrium agency royalties and retail prices.

Upstream-Downstream–Specific Royalties. The analysis above assumes that a separate royalty is set for every product, whereas we assume one royalty for each upstream-downstream pair when structurally estimating the model. Modifying the analysis to allow for a setting in which each upstream firm and retailer choose a single royalty for all of the upstream firm's products carried by the retailer is relatively straightforward. In bargaining over the profit-maximizing royalty to set, retailer d and upstream firm u recognize that a change in royalty r_j changes the royalties for the other products from upstream firm u that are carried by retailer d by the same amount. The first-order condition that reflects this behavior sets the sum of the derivatives in equation (24) across the products sold by the pair du equal to zero and evaluates this sum at a common royalty, r^{du}. That is,

$$\sum_{j \in du} \frac{\partial NP_{du}}{\partial r_j}\bigg|_{r_j = r^{du}\ \forall\ j \in du,\ \forall du} = 0, \quad \text{for all } du. \tag{25}$$

The components of the left hand side of equation (25) are the same as the components of equation (24).

4 Vertical Contracts in the E-Book Industry

In this section we focus on vertical contracts in the e-book industry. We first provide a description of important changes in the contracts between upstream book publishers and downstream book retailers. We then use a large dataset on retail prices in the period 2014-2015 to show how retail prices changed at Amazon and Barnes & Noble as a result of the switch from wholesale to agency contracts between publishers and bookstores. This transition to agency occurred after a period of intense bilateral bargaining between retailers and publishers.

4.1 Background

Initially e-books were sold using wholesale contracts. Publishers would set a list price for the e-book and would sell the book to a retailer for roughly half the list price. The retailer then would set a retail price at which to sell the product to the consumers. This vertical contract changed in 2010 with the introduction of the iPad when Apple, together with five of the (then) Big Six publishers, developed the agency model to sell e-books at the iPad's new iBookstore. Publisher's welcomed Apple's entrance to the e-book industry to provide a counterweight to Amazon's dominance and saw it as an opportunity to increase retail prices. Publishers believed that low e-book prices, especially Amazon's pricing of $9.99 for new releases, cannibalized hardcover sales and eroded consumers' perception of a book's value. As an MFN clause required the publishers to match retail prices at the iBookstore to the lowest price retailer, publishers compelled Amazon to adopt the agency model. Furthermore, the agency contracts included a mapping between list prices of newly released hardcover titles and the agency retail prices for the corresponding e-books, where this mapping was virtually identical across publishers.[15] 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, and 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.[16] Moreover, the retail price most-favored nation clauses that were seen as

[15]The two basic price tiers were $12.99 for hardcover prices between $25 and $27.50 and $14.99 for hardcover prices between $27.51 and $30.

[16]Note that termination dates of the bans were intentionally staggered to minimize the likelihood of collusive action

fundamental for making the switch to the agency model, were banned for a period of five years.[17] The U.S. district court argued that the two-year ban on agency and the five-year ban on retail price MFNs was necessary to provide a reset of the bilateral bargaining relationship between retailers and publishers. Apple did not settle, but eventually lost the case after further appeals. As part of the federal court's injunction, which went into effect in October 2013, Apple could net enter agency agreements with the publishers that were part of the lawsuit, with expiration dates ranging from 24 months for agreements with Hachette to 48 months for agreements with Macmillan.[18] De los Santos and Wildenbeest (2017) show that the transition from agency to wholesale contracts following the ban resulted in an 18 percent decrease in retail prices at Amazon and 8 percent at Barnes & Noble.

Table 1: New contract announcement and switch dates for Amazon

	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	Feb 01, 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 agency agreement announcement dates as well as approximate switch dates were widely reported by several media outlets (including a series of articles by Jeffrey Trachtenberg in the Wall Street Journal). Actual switch dates are verified using screenshots from Amazon, from which it can be inferred whether the price was set by the publisher or by Amazon. The dates of the start of the agency ban, which correspond to the switch to the wholesale model under the terms of the settlements, are taken from De los Santos and Wildenbeest (2017).

The first column in Table 1 displays the effective date of the start of the ban on agency contracts observed in the period after the settlement with the DOJ for each of the now Big Five publishers (De los Santos and Wildenbeest, 2017).[19] The second column of Table 1 displays the dates when it was

at the time new contracts had to be negotiated.

[17]Early 2017 Amazon agreed to stop enforcing e-book MFN clauses in Europe as part of a settlement with the European Commission. Although there was a similar lawsuit in 2012 in Europe as in the United States, with similar settlements (a two year ban on agency and five year ban on pricing MFNs), this does not necessarily imply that Amazon was using *retail price* MFN clauses in their agreements with the publishers prior to 2017 in the United States. Even though Amazon was not part of the 2012 lawsuits, the publishers that were part of the lawsuit were banned from using pricing MFNs for a period of five years, which effectively meant that also Amazon could also not use price MFNs in their agreements with these publishers during this period. However, since publishers that were not part of the lawsuit are not covered by the settlements, Amazon could have potentially used MFNs in contracts with other publishers. Moreover, the 2017 settlement agreement with the European Commission was about the use of MFNs in general, which includes other restrictive ebook contract clauses, such as the requirements to disclose to Amazon the terms of contracts with other retailers.

[18]According to the final judgment and order entering permanent injunction (see https://www.justice.gov/atr/case-document/file/486651/download), "Apple shall not enter into or maintain any agreement with a Publisher Defendant that restricts, limits, or impedes Apple's ability to set, alter, or reduce the Retail Price of any E-book or to offer price discounts or any other form of promotions to encourage consumers to purchase one or more E-books."

[19]Because of a merger between Penguin and Random House in July 2013, the Big Six was renamed the Big Five. Although Random House was not a conspirator defendant in the DOJ lawsuits, Random House adopted the terms of the settlement after the merger, which is prior to the sample period for this paper.

reported in the media that Amazon and each of the publishers had reached a bilateral agreement. The third column displays the dates when new agency agreements took effect and the switch to agency can be identified in the data. The table shows that even though each publisher announced an agreement with Amazon prior to the end of the two-year ban, the actual implementation dates of the new agreements varied between January and September of 2015, which means that most agreements did not go into effect immediately.

Note that while the media has reported extensively on Amazon's dealings with each of the Big Five publishers, we were unable to find reports on new agency agreements between Barnes & Noble and the publishers. Moreover, unlike Amazon, Barnes & Noble does not mention on a book's product page whether or not the price was set by the publisher. Nevertheless, the two-year ban on agency went into effect the same time for e-books sold at both Amazon and Barnes & Noble, which meant that contracts had to be renewed around the same time for both retailers. Moreover, new selling terms for Harper Collins e-books went into effect at the same date for *all* retailers.[20] We therefore make the assumption that for each publisher, the switch at Barnes & Noble happened at the same date as at Amazon. Also note that we exclude Apple from the analysis below, since it was banned from using agency agreements for a much longer period than the other retailers and therefore did not switch back to agency agreements during our sample period.

Figure 4: Amazon inventory and e-book pricing decisions

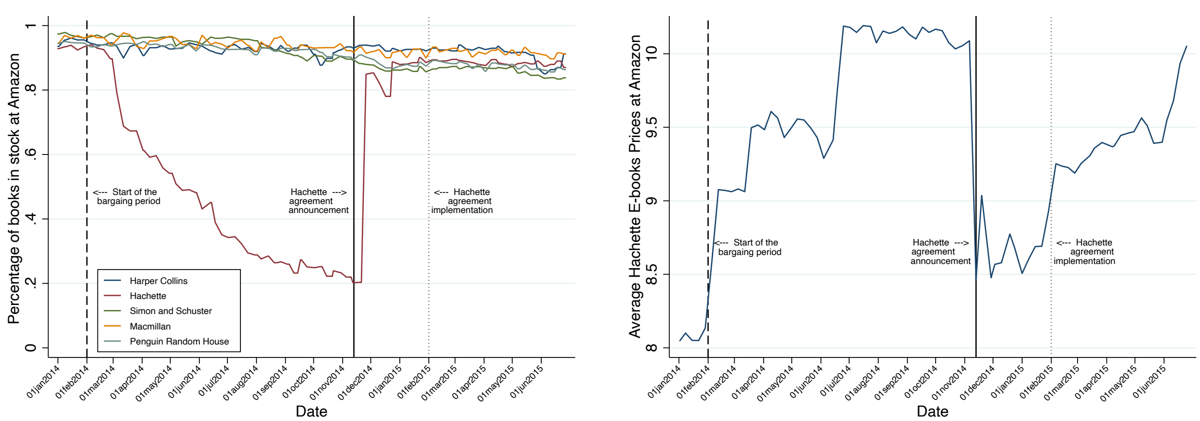

(a) Book Inventory by publisher (b) Hachette e-book prices at Amazon

Notes: Figure (a) gives the percentage of printed books in stock at Amazon over time for each of the Big Five publishers. Figure (b) gives the average Hachette e-book price at Amazon over time.

[20] According to Publishers Lunch (https://lunch.publishersmarketplace.com/2015/04/harper-readies-return-to-full-agency), "Multiple retailers report that Harper has informed them their selling terms will change as of Tuesday, April 14. (The change is actually effective midnight Pacific time, rather than Eastern. Amazon would be among those companies that naturally end their business day on Pacific time.) Harper is requiring retailers to implement all price changes within 24 hours. Going forward Harper will require that their e-books be sold at the publisher's listed consumer price, without any discounts." The article also notes that "Harper's notice to retailers is an 'interim' measure, in advance of more permanent new agency contracts," which means that even though other retailers could no longer offer discounts, they could still negotiate new terms regarding agency royalties.

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. The negotiations between Amazon and Hachette became well known publicly as they included various pressure tactics. Amazon reduced the inventory, delayed delivery, increased e-book prices, and removed the pre-order button of Hachette titles. Hachette started a public campaign to pressure Amazon, which included the involvement of support of some of their bestselling writers. The dispute between Amazon and Hachette started in February 2014 when Amazon did not allow customers to pre-order and reduced inventories of newly released Hachette books. Figure 4(a) shows that the percentage of printed Hachette books that were in stock at Amazon declined sharply from levels around 90 percent, which was similar to books from other publishers, to around 20 percent in November 2014.[21] After the agreement was announced, the percentage of Hachette books in stock immediately returned to 80 percent, which was similar to inventory levels at other large publishers. The figure also shows 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 last publisher to reach an agreement with Amazon. Figure 4(b) shows that e-book prices of Hachette titles increased sharply at the same time of the inventory reduction from average price levels of around $8 to $9 and continued increasing during the bargaining period up to levels around $10. After the announcement of the agreement, Amazon dropped prices sharply to levels around $8.50. Prices increased again after the implementation of the agency agreement which gave control of retail prices to Hachette.

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 November 5, 2014 (seven weeks before the first Big Five publisher switched) and October 21, 2015 (seven weeks after the last Big Five publisher switched) and consists of daily prices (obtained using a web scraper) for a large number of e-book titles. All titles are new and former New York Times bestseller books. Books that appear in the New York Times bestseller lists are added to the sample from the moment of the appearance on the list.[22] 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 publisher, number of pages, and ratings and number of reviews at Amazon. For the analysis, we focus on the Big Five publishers and aggregate the data to weekly observations. Table 2 summarizes the variables.

For the analysis in this section, we use a similar difference-in-differences (DID) approach as

[21] The data used to create Figures 4(a) and 4(b) contains the same e-book titles used in the analysis in the next subsection, but also includes observations prior to November 5, 2014 as well as stock information for the printed version of the e-books in the main sample.

[22] We modified the collection method for technical reasons in July 21, 2015. Because of this, the number of books we could track was reduced and was restricted to mostly popular books, as defined by the sales rank.

Table 2: Summary statistics

	Harper Collins	Hachette	Simon & Schuster	Macmillan	Penguin Random House
Price e-book					
Amazon	9.98	9.46	11.30	9.68	9.28
	(3.52)	(2.57)	(2.93)	(2.81)	(2.68)
Barnes & Noble	11.17	10.07	11.68	9.98	11.16
	(3.78)	(2.64)	(2.68)	(2.72)	(2.71)
Book characteristics					
Ratings	4.31	4.23	4.30	4.21	4.25
	(0.37)	(0.40)	(0.39)	(0.40)	(0.39)
Number of reviews	1022.19	1264.89	944.45	882.05	1205.53
	(2300.17)	(2152.31)	(1528.53)	(1478.83)	(2792.97)
Number of months since release	26.74	29.64	30.29	26.85	31.90
	(23.34)	(23.46)	(28.45)	(19.43)	(41.12)
Number of pages	374.62	394.39	383.56	396.54	398.20
	(133.19)	(134.90)	(142.50)	(152.95)	(219.30)
Number of titles	381	316	419	252	1,326
Number of observations	20,016	20,055	26,334	15,632	83,551

Notes: The table gives the mean of each variable, with standard deviations in parentheses.

in De los Santos and Wildenbeest (2017). But where 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 occurred after the two-year ban on agency had expired in the period 2014-2015. An important difference is that during the first period several of the key players in the industry were found to be colluding. Another important difference is that MFN clauses were not used during the second agency period and therefore do not play a role explaining the higher agency prices, as argued by Johnson (2017). Note that in Section 5.5 we make an explicit comparison between the first agency period and the second agency period when discussing the results of a counterfactual exercise to study the effects of MFN on agency prices.

As was shown in Table 1, new contracts were announced between Amazon and the major publishers at different points in time, resulting in the staggering of the actual switching dates at Amazon. We exploit this cross-publisher variation in the timing of the switch in a difference-in-differences setup. Specifically, the baseline specification we estimate is

$$\log(price_{jt}) = \gamma \cdot agency_{jt} + \beta \cdot X_j + \nu_p + \nu_w + \varepsilon_{jt}, \tag{26}$$

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 ν_p and ν_w are publisher and week fixed effects. The difference-in-differences estimator is captured by γ and gives the price effect of the switch to the agency model.

Table 3: Results difference-in-differences analysis

	Amazon				Barnes & Noble			
	Publisher fixed effects		Book fixed effects		Publisher fixed effects		Book fixed effects	
Difference-in-differences estimator								
Agency	0.132		0.127		-0.015		-0.022	
	(0.008)		(0.007)		(0.006)		(0.005)	
Agency × Harper Collins		0.139		0.118		-0.048		-0.066
		(0.018)		(0.015)		(0.016)		(0.011)
Agency × Hachette		0.072		0.073		0.014		0.007
		(0.013)		(0.010)		(0.012)		(0.009)
Agency × Simon & Schuster		0.126		0.138		-0.020		-0.014
		(0.013)		(0.012)		(0.008)		(0.006)
Agency × Macmillan		0.097		0.078		-0.054		-0.077
		(0.019)		(0.019)		(0.017)		(0.017)
Agency × Penguin Random House		0.246		0.258		0.053		0.058
		(0.016)		(0.014)		(0.015)		(0.009)
Other controls								
Rating	0.020	0.020	0.027	0.031	0.012	0.012	-0.018	-0.023
	(0.013)	(0.013)	(0.016)	(0.016)	(0.012)	(0.012)	(0.015)	(0.015)
Months since release	-0.002	-0.002	0.001	0.001	-0.001	-0.001	0.001	0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Pages	0.071	0.070			0.091	0.090		
	(0.050)	(0.050)			(0.049)	(0.050)		
Constant	2.122	2.128	2.039	2.029	2.321	2.322	2.423	2.448
	(0.059)	(0.059)	(0.069)	(0.070)	(0.055)	(0.055)	(0.064)	(0.064)
R-squared	0.152	0.155	0.829	0.832	0.078	0.080	0.855	0.857
Number of observations	75,842	75,842	75,841	75,841	76,649	76,649	76,624	76,624

Notes: Dependent variable is log(*price*). All specifications include week fixed effects. Standard errors (clustered by book) in parentheses. The number of pages is divided by thousand.

Table 4: Placebo tests on the effect of switch to agency on print book prices

	Amazon		Barnes & Noble	
Difference-in-differences estimator				
Agency	-0.001		-0.003	
	(0.004)		(0.003)	
Agency × Harper Collins		-0.005		-0.011
		(0.009)		(0.007)
Agency × Hachette		-0.021		-0.010
		(0.009)		(0.005)
Agency × Simon & Schuster		0.009		0.004
		(0.007)		(0.005)
Agency × Macmillan		0.015		0.013
		(0.008)		(0.008)
Agency × Penguin Random House		0.006		
		(0.020)		
Other controls				
Rating	-0.042	-0.040	0.015	0.016
	(0.020)	(0.016)	(0.013)	(0.013)
Months since release	0.001	0.001	-0.002	-0.002
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	2.784	2.773	2.740	2.733
	(0.085)	(0.086)	(0.055)	(0.055)
R-squared	0.896	0.896	0.839	0.839
Number of observations	75,818	75,818	67,267	67,267

Notes: Dependent variable is log(*price*) of the printed version of the e-book. All specifications include book and week fixed effects. Standard errors (clustered by book) in parentheses.

Table 3 gives the main results for the difference-in-differences analysis. We estimate equation (26) 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 14 percent as a result of the switch from wholesale to agency; for Barnes & Noble prices went down by approximately 2 percent. Table 3 also 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: the percentage increase in e-book prices following the switch ranges from 7 percent for Hachette to 28 percent for Penguin Random House. The findings for books sold at Barnes & Noble are also mixed. Prices for Macmillan and Harper Collins titles decreased following the switch to agency, while prices for books published by Penguin Random House increased by approximately 5 percent. The evolution of prices for the publishers are also shown in Figure 5, which show the average price changes at Amazon by publisher around the agreement announcement (dotted vertical line) and the date of the switch to agency (solid vertical line). The figures show jumps in prices for Amazon around the time of the switch, but much less so for Barnes & Noble.

To analyze if the effects can be attributed to the switch from wholesale to agency model and not to other shocks that may have occurred around the switching dates, we run a placebo test in

Figure 5: E-book prices at Amazon and Barnes & Noble by publisher

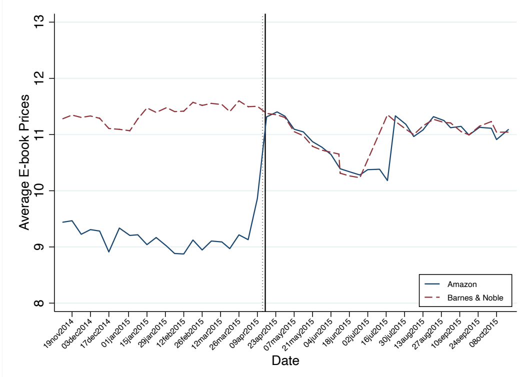

(a) Harper Collins

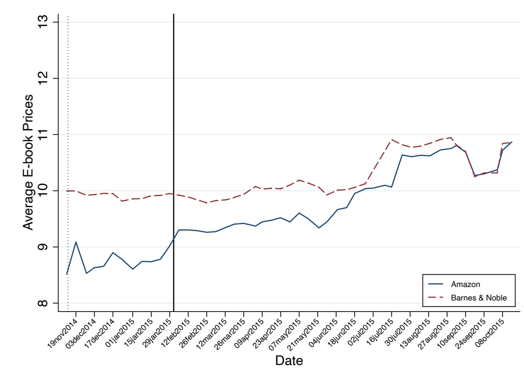

(b) Hachette

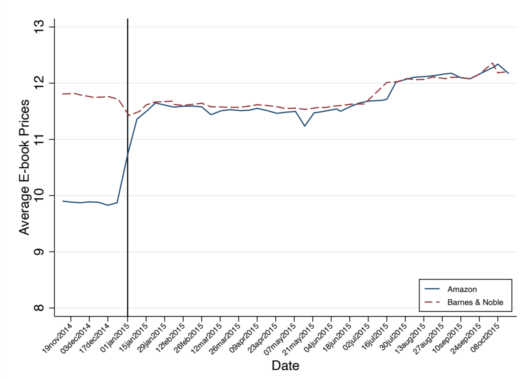

(c) Simon & Schuster

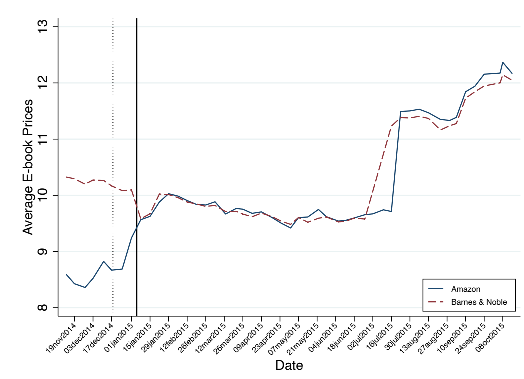

(d) Macmillan

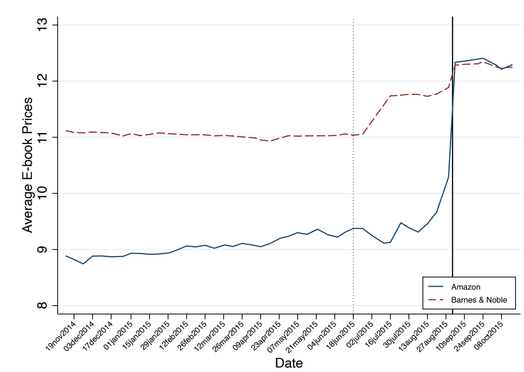

(e) Penguin Random House

Notes: Average e-book prices at Amazon and Barnes & Noble by publisher.

which we replicate the analysis using prices for the printed version of the e-book. Since printed titles remained under the wholesale model, one would not expect to see a significant change in prices for those books. Table 4 shows the results for a specification that has book fixed effects in which we do not make a distinction by publisher, as well as the results split out by publishers.

The table shows that if we look at the overall effect of the switch to agency, there is no effect at both Amazon and Barnes & Noble. If we split out the effect by publisher, we find some effect for Hachette titles, although the effect is very small and not very precisely estimated.

The results from the difference-in-differences analysis point to two important observations. First, the effect of the switch to agency agreements was different for Amazon than for Barnes & Noble. Second, 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

5.1 Data

We use a subset of the data used in the previous section to estimate the structural model—for the estimation of the wholesale model we use the first seven weeks of the sample, while we use the last seven weeks of the sample for the estimation of the agency model. Table 5 provides summary statistics of the main variables we use by publisher. Panel A of the table gives summary statistics for the seven-week period in which wholesale agreements were used by all of the publishers (November/December 2014) whereas panel B gives summary statistics for the seven-week period in which agency agreements were used by all the Big Five publishers (September/October 2015). As before, we use weekly observations. For both periods we use the 100 most popular titles in our sample, using the restriction that for a title to be included, we need to have observations for both Amazon and Barnes & Noble throughout the entire sample.[23] In total we have 14 weeks of data, which corresponds to a total of 2,800 weekly observations (1,400 for the wholesale period and 1,400 for the agency period). The largest Big Five publisher, Penguin Random House, represents most of the observations in our sample. Macmillan is the smallest with 6 titles. As we lack quantity data for each book title, we use the approach suggested by Chevalier and Goolsbee (2003) to infer sales from the observed sales rank data (see Appendix D for details). A comparison of the average prices under the two selling regimes indicates that even though average prices were $2.59 higher at Barnes & Noble than at Amazon during the wholesale period, average prices under agency are very similar across the retailers, despite the five-year ban on the use of retail price MFN clauses during this period.

5.2 Estimation Approach

Our estimation approach is to jointly estimate demand and supply and to use covariance restrictions to deal with endogeneity concerns (MacKay and Miller, 2021). The unobserved characteristic

[23] Unfortunately, we do not have data sales rank data for Apple, so we cannot include Apple when estimating the model.

Table 5: Summary statistics

	Panel A. Wholesale period					Panel B. Agency period				
	Harper Collins	Hachette	Simon & Schuster	Mac- millan	Penguin Random House	Harper Collins	Hachette	Simon & Schuster	Mac- millan	Penguin Random House
Price e-book										
Amazon	9.98	7.67	6.17	7.47	7.59	9.97	9.46	10.27	11.03	10.81
	(3.86)	(2.33)	(2.18)	(3.42)	(2.42)	(1.29)	(2.06)	(1.19)	(3.16)	(1.95)
Barnes	11.38	9.47	10.66	11.44	10.13	10.03	9.46	10.66	11.03	10.82
& Noble	(3.39)	(2.75)	(0.73)	(2.55)	(1.97)	(1.22)	(2.05)	(0.73)	(3.16)	(1.95)
Weekly sales										
Amazon	4,225	6,627	736	1,074	4,057	1,165	484	679	366	2,724
	(5,891)	(24,038)	(585)	(881)	(14,717)	(3,778)	(318)	(2,954)	(205)	(17,298)
Barnes	1,557	1,040	128	144	520	269	316	107	149	390
& Noble	(2,864)	(3,869)	(168)	(93)	(1,373)	(268)	(346)	(120)	(153)	(1,136)
Book characteristics										
Ratings	4.30	4.20	4.40	4.40	4.29	4.24	4.22	4.43	4.41	4.29
	(0.25)	(0.38)	(0.16)	(0.20)	(0.32)	(0.26)	(0.34)	(0.15)	(0.20)	(0.33)
Months	12	21	32	25	48	22	31	42	35	58
since release	(10)	(20)	(18)	(16)	(84)	(10)	(20)	(18)	(16)	(84)
Pages	401	443	497	499	506	401	443	497	499	506
	(76)	(109)	(207)	(273)	(229)	(76)	(109)	(207)	(273)	(229)
Titles	7	18	8	6	61	7	18	8	6	61
Observations	98	252	112	84	854	98	252	112	84	854

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

ξ_j in the utility function (14) captures unobserved quality, which is likely to be correlated with a book's price. Our main specification below includes product and store fixed effects, so the product- and store-specific variation in unobserved quality that does not vary over time is captured by these fixed effects. However, these 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 and retail prices. The traditional approach to resolve price endogeneity is to use supply-side instruments and estimate the demand side of the model using instrumental variables. However, 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. Instead, we follow the approach suggested by MacKay and Miller (2021), which is to restrict the covariance between unobserved demand and cost shocks to be zero. Intuitively, the supply side model imposes restrictions on how prices respond to demand shocks, and by restricting how the demand and supply unobservables move together, it is possible to obtain causal estimates of the endogenous price parameter without needing supply-side instruments. Although an assumption of zero correlation between the unobserved demand and cost shocks may seem ad hoc, we include fixed effects in both the demand and cost function to make the zero-covariance restriction more credible. For instance, a potential concern here is that books with greater unobserved quality may also have higher marginal costs—the product fixed effects will absorb the quality component from the demand and costs shocks, making the zero-correlation assumption more plausible.

The parameters to be estimated are the parameters of the demand and marginal cost functions, as well as the bargaining parameters. 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 type of vertical contract). In addition to assuming the bargaining parameters do not vary with the type of vertical contract, we assume that the demand parameters as well as the marginal cost parameters remain the same across the two periods, which means that identification of the bargaining parameters is largely driven by the change in prices we observe as a result of switching from the wholesale model to the agency model.

Start with the demand side of the model. Since we are assuming the model has a logit structure, we can express demand for product j in period t as

$$\log(s_{jt}) - \log(s_{0t}) - \alpha \log(p_{jt}) = X_{jt}\beta + \xi_{jt}, \tag{27}$$

where p_{jt} is likely to be correlated with ξ_{jt}. Given α, the demand side residuals can be obtained as the residuals of a regression of the left-hand side of this expression on X_{jt}.

Next, consider the supply side of the model. Since the retailers' marginal costs are likely to be negligible, we set the retailers' marginal cost to zero.[24] The combined upstream and downstream markup of product j equals the difference between price and marginal cost, which means upstream marginal cost c_{jt} equals

$$c_{jt} = p_{jt} - m_{jt}^D - m_{jt}^U,$$

where, depending on whether agency or wholesale agreements were used in period t, we use the first-order conditions for the wholesale or agency period derived in Section 3 to solve for m_{jt}^D and m_{jt}^U as a function of the parameters.

Consider first the wholesale model. The vector of downstream margins m^D can be written as

$$m^D = -\left(T^D \cdot \Delta\right)^{-1} s, \tag{28}$$

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 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 (18).[25] The upstream margin when wholesale contracts are used can be found by solving the equilibrium condition in equation (22) for

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

[25] When using matrix notation, we use $\cdot$ to indicate an entrywise (Hadamard) product.

m^U—we show in Appendix C that, using matrix notation, this can be written as

$$m^U = -\left(T^U \cdot Z^w + \left[(T^D \cdot S)m^D \left(\Delta^D_w\right)^{-1}\right]^{-1} \left(T^U \cdot S\right) \frac{1-\lambda}{\lambda}\right)^{-1} s. \tag{29}$$

In this equation 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, Z^w is a matrix that captures how market shares change through changes in equilibrium prices (derived in Appendix C), $\Delta^D_w = \partial\pi^D/\partial w$, and S is a matrix with market shares on the diagonal and the differences in markets shares when product j is not offered as off-diagonal elements, i.e.,

$$S = \begin{bmatrix} s_1 & -\Delta s_2^{-1} & \dots & -\Delta s_N^{-1} \\ -\Delta s_1^{-2} & s_2 & \dots & -\Delta s_N^{-2} \\ \vdots & \vdots & \ddots & \vdots \\ -\Delta s_1^{-N} & -\Delta s_2^{-N} & \dots & s_N \end{bmatrix}. \tag{30}$$

In equation (30), Δs_k^{-j} is defined as the additional market share for product k when product j (and all other products that are part of the downstream-upstream combination du) is not offered, i.e., $\Delta s_k^{-j} = s_k^{-du} - s_k$, with s_k^{-du} defined as in equation (20).[26]

Consider next the agency model. Since the upstream firm sets retail prices, we can solve the pricing first-order condition in equation (23) for the vector of upstream margins m^U, which gives using matrix notation

$$m^U = -\left(T^U \cdot \Delta\right)^{-1} (1-r)s. \tag{31}$$

The downstream margin under agency is a directly related to the royalty r and, assuming zero marginal cost for the retailers, is given by

$$m^D = rp. \tag{32}$$

Equations (31) and (32) show that to obtain the upstream and downstream margins in the agency model, we need a vector of agency royalties r in addition to the price parameter α. The agency royalties are not observed, but, since they are determined through bargaining, can be obtained by solving the bargaining first-order condition in equation (24) for r, given bargaining weights λ and price parameter α.

To obtain the supply side residuals η we assume upstream marginal cost for product j in market t is given by

$$\log(c_{jt}) = Z_{jt}\gamma + \eta_{jt}. \tag{33}$$

[26]The notation used for S follows Draganska, Klapper, and Villas-Boas (2010). Note that if product j and k are part of the same downstream-upstream combination du, then if j is not offered, k is not offered as well, resulting in $\Delta s_k^{-j} = -s_k$. Also not that using matrix S, the difference between profits and disagreement profits can be written as $\pi^U - d^U = (T^U \cdot S)m^U$ for the upstream firm and as $\pi^D - d^D = (T^D \cdot S)m^D$ for the downstream firm.

Equations (27) and (33) give us the demand side residuals $\xi_{jt}(\alpha, \beta, \lambda; X_{jt})$ and supply side residuals $\eta_{jt}(\alpha, \gamma, \lambda; X_{jt})$ given the data and the parameters of the model. Following MacKay and Miller (2019), we use the following orthogonality condition for identification

$$E_t[\xi_{jt} \cdot \eta_{jt}] = 0 \quad \forall j, k,$$

where the expectation is over markets. The corresponding empirical moments that are used to estimate the parameters of the model σ (which includes $\alpha, \beta, \gamma, \lambda$) by GMM are then

$$\hat{\sigma} = \arg\min_{\sigma} \left[\frac{1}{J^2} \sum_{j,k} \left(\frac{1}{T} \sum_{t} \xi_{jt}(\sigma; X_{jt}) \cdot \eta_{jt}(\sigma; X_{jt}) \right)^2 \right]. \tag{34}$$

Computational Details. Due to the inherent complexity of the bargaining first-order condition in the agency model, solving this first-order condition for r has to be done numerically. Because of this, it is computationally much faster to start from a vector of royalties (one for each retailer-publisher pair) and to use the bargaining first-order condition to obtain the vector of bargaining weights λ (again, one for each retailer-publisher pair) as a function of the agency royalties and the other parameters of the model for which the bargaining first-order condition holds. Specifically, solving equation (24) for $(1 - \lambda)/\lambda$ gives

$$\frac{1-\lambda}{\lambda} = -\frac{(\pi^D - d^D)\left(\partial \pi^D / \partial r\right)^{-1}}{(\pi^U - d^U)\left(\partial \pi^U / \partial r\right)^{-1}}. \tag{35}$$

Taking into account the ownership structure and using matrix notation, we can write $\pi^D - d^D = (T^D \cdot S)m^D$ and $\pi^U - d^U = (T^U \cdot S)m^U$. We can then write equation (35) more compactly using matrix notation, i.e.,

$$\frac{1-\lambda}{\lambda} = \left[(T^U \cdot S)m^U \left(\Delta_r^U\right)^{-1} \right]^{-1} (T^D \cdot S)m^D \left(\Delta_r^D\right)^{-1}, \tag{36}$$

where $\Delta_r^U = \partial \pi^U / \partial r$ and $\Delta_r^D = \partial \pi^D / \partial r$. By reorganizing the bargaining first-order condition for the agency model this way, we end up with a closed-form expression for the bargaining parameters as a function of the price parameter α as well as the agency royalties r (see equation (A10) in Appendix C). Then, when estimating the model, instead of starting from an initial vector of bargaining parameters λ and solving for r, one can starting from an initial vector of agency royalties r and solve for λ, which significantly speeds up estimation.[27] Since we are assuming the bargaining weights do not depend on the selling method, we can then use equation (36) to substitute for $(1 - \lambda)/\lambda$ in equation (29).

[27] It is important to stress that starting from a vector of λ's and using the bargaining first-order condition to numerically solve for r (given the other parameters) is identical to proceeding the other way around, but takes much longer since the bargaining first-order condition has to be solved numerically.

5.3 Small Sample Performance

To investigate identification as well as small sample performance of our estimation procedure, we run a Monte Carlo exercise. The setup of the experiment is as follows. We simulate two retailers who each sell two products, where the two products are provided by two different upstream firms, so there are four differentiated products in total. We let the firms compete for a number of periods, where in half the periods retail and input prices are set according to the wholesale model and in half the periods retail and input prices are set using the agency model. We assume disagreement payoffs are zero when bargaining. Furthermore, we let consumer i's indirect utility for product j in period t be given by $u_{ijt} = 10 - 4\log(p_{jt}) + \xi_{jt} + \varepsilon_{ijt}$, whereas product j's marginal cost during period t is given by $\log(c_{jt}) = 1 + \eta_{jt}$. We generate data assuming both demand and supply shocks are uniform, with $\xi \sim U(0, 0.5)$ and $\eta \sim U(0, 0.5)$. Table 6 reports the mean and standard deviation of the estimated parameters, where each column represents a different number of markets.

Table 6: Monte Carlo simulations

		10 markets		40 markets		100 markets	
Variable	True	Coeff.	Std. Dev.	Coeff.	Std. Dev.	Coeff.	Std. Dev.
PANEL A. DEMAND MODEL ONLY (OLS)							
log(price)	-4.000	-3.539	(0.333)	-3.549	(0.159)	-3.553	(0.099)
PANEL B. DEMAND AND SUPPLY MODEL							
Demand parameters							
log(price)	-4.000	-4.300	(0.468)	-4.045	(0.230)	-3.987	(0.148)
Bargaining parameters Retailer 1 (R1)							
Upstream firm 1 (U1)	0.200	0.199	(0.074)	0.190	(0.065)	0.188	(0.063)
Upstream firm 2 (U2)	0.300	0.306	(0.107)	0.289	(0.095)	0.287	(0.090)
Bargaining parameters Retailer 2 (R2)							
Upstream firm 1 (U1)	0.400	0.410	(0.110)	0.394	(0.099)	0.388	(0.095)
Upstream firm 2 (U2)	0.250	0.253	(0.095)	0.238	(0.085)	0.235	(0.078)
Number of observations		40		160		400	

Notes: Results for each column are based on 10,000 simulations. Utility is given by $u_{ijt} = 10 - 4\log(p_{jt}) + \xi_{jt} + \varepsilon_{ijt}$ and marginal cost by $\log(c_{jt}) = 1 + \eta_{jt}$. Data is generated assuming $\xi \sim U(0, 0.5)$ and $\eta \sim U(0, 0.5)$.

Panel A of Table 6 shows the estimates when only the demand side is estimated, using OLS—failing to deal with price endogeneity results in the price parameter being underestimated, even when 400 observations are used, as in the last column of the table. Panel B shows that by jointly estimating supply and demand using the cross-covariance restrictions approach, the bias in the price parameter disappears. The average price parameter is relatively close to the true parameter value even when the number of markets it small, and the standard deviation of the estimates decreases in the number of observations used. Furthermore, the nonlinear bargaining parameters are very close to the true parameter values even when data from only ten markets is used, and the estimates become more precise when the number of markets increases. As argued before, the identification of the bargaining parameters follows from observed changes in prices when switching selling methods, and the Monte Carlo results show that using data from both pricing regimes allows one to successfully pin down the bargaining parameters while simultaneously dealing with price endogeneity, even in relatively small samples.

Figure 6: Simulated prices for each upstream-retailer pair

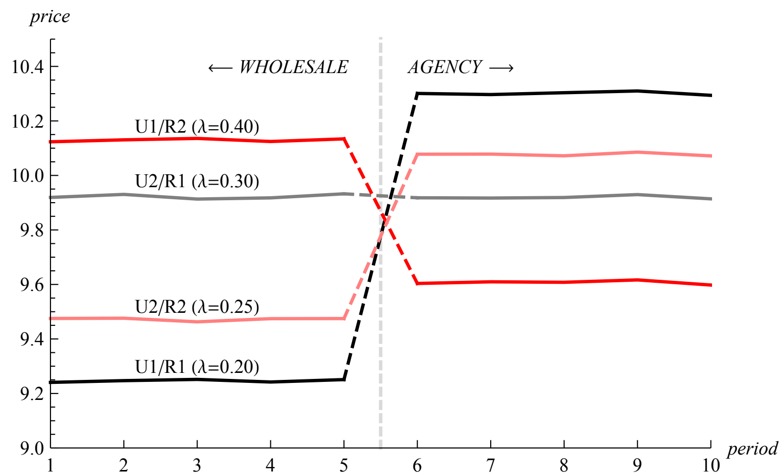

Notes: Simulated prices for each upstream-retailer pair for the 10 period case (averaged across 10,000 simulations). Utility is given by $u_{ijt} = 10 - 4\log(p_{jt}) + \xi_{jt} + \varepsilon_{ijt}$ and marginal cost by $\log(c_{jt}) = 1 + \eta_{jt}$. Data is generated assuming $\xi \sim U(0, 0.5)$ and $\eta \sim U(0, 0.5)$. Demand and bargaining parameters are given by Panel B of the first column of Table 6.

The identification argument can also be seen from Figure 6, which shows simulated prices for each upstream-retailer pair for the 10 period case (averaged across the 10,000 simulations). The figure shows that for upstream-retailer pairs with a relatively low bargaining parameter (U1/R1 and U2/R2), prices increase when switching to agency, whereas prices remain constant (U2/R1) or go down (U1/R2) for the upstream-retailer pairs that face higher values of their bargaining parameter λ.

5.4 Parameter Estimates

Table 7 gives the parameter estimates for various specifications of the bargaining model. Specification (A) of Table 7 gives estimates of the demand parameters as well as the bargaining parameters when using publisher fixed effects in the demand equation. Other demand shifters we include are the log of the price, the five-star rating of the title on Amazon, months since the title was released, the number of pages, store fixed effects, and week fixed effects. As marginal cost shifters we include book fixed effects. In specification (B) we use book fixed effects in the demand specification instead of publisher fixed effects, which means we can no longer include the number of pages, but leave everything else the same. In both specifications the price coefficient is estimated precisely, with demand estimated to be slightly less elastic in specification (B) than (A). Rating has the expected sign in both specifications, but is only estimated precisely in the specification that has book fixed effects. Months since release and the number of pages do not appear to significantly affect the

Table 7: Parameter estimates of the bargaining model

Variable	(A) Publisher Fixed Effects Coeff.	(A) Publisher Fixed Effects Std. Err.	(B) Book Fixed Effects Coeff.	(B) Book Fixed Effects Std. Err.
Demand parameters				
log(price)	-2.066	(0.188)	-1.805	(0.139)
rating	0.481	(0.353)	1.603	(0.696)
months since release	-0.004	(0.002)	-0.384	(0.361)
number of pages	0.638	(0.543)		
Bargaining parameters Amazon				
Harper Collins	0.352	(0.112)	0.347	(0.099)
Hachette	0.239	(0.051)	0.256	(0.043)
Simon & Schuster	0.037	(0.144)	0.103	(0.081)
Macmillan	0.097	(0.142)	0.149	(0.091)
Penguin Random House	0.153	(0.052)	0.188	(0.032)
Average	0.176		0.208	
Bargaining parameters Barnes & Noble				
Harper Collins	0.477	(0.100)	0.441	(0.080)
Hachette	0.400	(0.035)	0.378	(0.027)
Simon & Schuster	0.430	(0.125)	0.389	(0.077)
Macmillan	0.460	(0.130)	0.418	(0.079)
Penguin Random House	0.368	(0.022)	0.351	(0.020)
Average	0.427		0.395	
Objective function	0.020		0.007	
Own price elasticity	-2.066		-1.804	
Number of observations	2,800		2,800	

Notes: Bootstrapped standard errors shown in parentheses (clustered by book title). Demand specification includes store and week fixed effects. Marginal cost specification includes book fixed effects.

marginal utility of an e-book.

All of the estimated bargaining power parameters are less than a half, which, given our assumption that the bargaining weights are not affected by the selling method, suggests the retailers have more bargaining power than the publishers. The bargaining parameters for Amazon are generally lower than those for Barnes & Noble, which is consistent with the results of the difference-in-difference analysis in Section 4, which indicated a bigger switch-related price effect for Amazon than for Barnes & Noble, and hence suggests relatively more bargaining power for Amazon than for Barnes & Noble. However, there is substantial variation in the estimated bargaining parameters across publishers, which to some extent can be explained by the magnitude of the price effects we found in Section 4. Especially the estimated bargaining parameters for Barnes & Noble show a clear relation with the results of the difference-in-differences analysis: for both Harper Collins and Macmillan we found a negative price effect of the switch to agency, which indicates these publishers having relatively more bargaining power, exactly as their above average bargaining weight estimates show. The relation with the reduced-form findings is less clear for Amazon, although we do find bargaining power for Simon & Schuster and Penguin Random House—the publishers for which we found the biggest price effects—to be below-average.

Table 8: Royalty parameter estimates

	Amazon		Barnes & Noble		Average
SPECIFICATION A. PUBLISHER FIXED EFFECTS					
Harper Collins	0.322	(0.064)	0.254	(0.061)	0.288
Hachette	0.373	(0.036)	0.291	(0.042)	0.332
Simon & Schuster	0.472	(0.065)	0.277	(0.044)	0.375
Macmillan	0.441	(0.072)	0.262	(0.055)	0.352
Penguin Random House	0.420	(0.028)	0.307	(0.039)	0.363
Average	0.406		0.278		0.342
SPECIFICATION B. BOOK FIXED EFFECTS					
Harper Collins	0.371	(0.072)	0.311	(0.058)	0.341
Hachette	0.418	(0.032)	0.346	(0.035)	0.382
Simon & Schuster	0.504	(0.044)	0.339	(0.033)	0.422
Macmillan	0.476	(0.056)	0.323	(0.041)	0.400
Penguin Random House	0.461	(0.026)	0.361	(0.029)	0.411
Average	0.446		0.336		0.391

Notes: Bootstrapped standard errors shown in parentheses (clustered by book title).

Table 8 gives the estimated royalty parameters during the agency period for the two specifications in Table 7. The average royalty share across retailers and publishers is 0.391 for the model with book fixed effects and 0.342 for the model with publisher fixed effects, which is higher than 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 is in line with Amazon's more favorable bargaining power estimates. Publishers with below average bargaining power tend to have lower agency royalties, although a publisher's bargaining *position* matters here too. For instance, despite having lower than average bargaining weight, Penguin Random House has above average royalities, which can be explained by Penguin's better bargaining position due to its size relative to other publishers (which affects its disagreement payoffs).[28]

Table 9 reports the implied margins for the bargaining model estimates reported in column (B) of Table 7. The table reports figures for both the wholesale and agency model and shows that the publishers' margins went up for all of the publishers when switching to agency contracts. Amazon's average downstream margin also went up when switching to agency contracts, but Barnes & Noble's margins went down as a result of the switch. Note that higher margins (and possibly profits) for Amazon are not inconsistent with the theoretical model of Section 2—Figure 3 illustrates this point by showing that while upstream profits are higher under agency, whether downstream profits are higher under agency depends on the bargaining power parameters.

[28] According 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.

Table 9: Prices and margins

	Panel A. Wholesale period					Panel B. Agency period				
	Retail price	Down-stream margin	Whole-sale price	Up-stream margin	Mar-ginal cost	Retail price	Down-stream margin	Whole-sale price	Up-stream margin	Mar-ginal cost
Retailers										
Amazon	7.65	4.32	3.33	0.70	2.63	10.48	4.74	5.74	3.20	2.54
Barnes & Noble	10.24	5.70	4.55	1.45	3.10	10.52	3.69	6.83	3.81	3.02
Publishers										
Harper Collins	10.68	5.97	4.71	1.57	3.14	10.00	3.41	6.59	3.66	2.94
Hachette	8.57	4.80	3.77	1.09	2.67	9.46	3.61	5.85	3.24	2.61
Simon & Schuster	8.55	4.79	3.76	0.95	2.81	10.47	4.40	6.07	3.36	2.70
Macmillan	9.45	5.29	4.16	1.13	3.03	11.03	4.41	6.62	3.67	2.95
Penguin Random House	8.86	4.96	3.90	1.02	2.88	10.82	4.44	6.37	3.56	2.81

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

Robustness

Table 10 gives estimates for alternative models and estimation approaches, all estimated using book fixed effects in the demand side of the model. In column (A) of the table we allow for a more flexible demand system that allows for different price coefficients for the two retailers. Our main specification has a single price parameter, which means that if Amazon and Barnes & Noble cater to consumers who differ in their price sensitivity, the difference in price sensitivity will be picked up by the bargaining parameters. As shown by the price parameter estimates in column (A) of Table 10, Barnes & Noble customers do appear to be less price sensitive on average than Amazon customers, although the difference is small. The bargaining parameters are not much affected by allowing for this additional layer of heterogeneity, although Barnes & Noble is now estimated to have slightly more bargaining power on average than in our main specification.

In column (B) of Table 10 we estimate the model 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. The estimated demand parameters for this specification indicate demand is slightly more elastic, whereas the other demand parameters do not differ much from the ones estimated in column (B) of Table 7. A comparison of the objective function values for the two specifications suggests 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 types of vertical contracts, while with take-it or leave-it contracts the publishers have all the bargaining power in the wholesale model ($\lambda = 1$) and retailers have all the bargaining power in the agency model ($\lambda = 0$). To formally test which model gives the best fit to the data, we use the non-nested test proposed by Rivers and Vuong (2002). The test statistic is -200.51 which means we can strongly reject the take-it or leave-it model against the bargaining model.[29]

[29] The Rivers and Vuong (2002) statistic is calculated as $T_N = (\sqrt{N}/\hat{\sigma}_N)(\hat{Q}_1 - \hat{Q}_2)$, where N is the number of observations, $\hat{Q}_1$ and $\hat{Q}_2$ are the objective function values of the bargaining model and the take-it or leave-it model,

Table 10: Parameter estimates alternative supply side model and estimation approaches

	(A) Retailer-specific price parameter		(B) Take-it or leave-it contracts		(C) Royalties fixed to 0.30		(D) Demand and supply separately	
Variable	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Demand parameters								
log(price)			-2.130	(0.182)	-2.427	(0.082)	-1.341	(0.178)
Amazon	-1.726	(0.120)						
B&N	-1.588	(0.096)						
rating	1.606	(0.693)	1.607	(0.709)	1.612	(0.806)	1.888	(0.810)
months since release	-0.400	(0.373)	-0.327	(0.326)	-0.276	(0.439)	-0.251	(0.089)
Bargaining parameters Amazon								
Harper Collins	0.333	(0.027)			0.287	(0.038)	0.277	(0.035)
Hachette	0.249	(0.093)			0.280	(0.034)	0.225	(0.073)
Simon & Schuster	0.108	(0.040)			0.283	(0.053)	0.133	(0.075)
Macmillan	0.151	(0.081)			0.279	(0.080)	0.160	(0.025)
Penguin Random House	0.187	(0.085)			0.288	(0.029)	0.193	(0.045)
Average	0.206				0.283		0.197	
Bargaining parameters Barnes & Noble								
Harper Collins	0.417	(0.028)			0.273	(0.030)	0.356	(0.014)
Hachette	0.366	(0.065)			0.274	(0.024)	0.321	(0.083)
Simon & Schuster	0.376	(0.022)			0.273	(0.047)	0.336	(0.074)
Macmillan	0.399	(0.078)			0.273	(0.046)	0.350	(0.015)
Penguin Random House	0.347	(0.075)			0.274	(0.031)	0.311	(0.009)
Average	0.381				0.274		0.335	
Objective function	0.007		0.016		0.014		0.051	
Own price elasticity	-1.657		-2.130		-2.427		-1.341	
Number of observations	2,800		2,800		2,800		2,800	

Notes: Bargaining parameters are fixed to 1 for wholesale model and 0 for agency model in specification (A). Bootstrapped standard errors shown in parentheses (clustered by book title), except for demand parameters in column (D), which are TSLS standard errors. Demand specifications include book, store, and week fixed effects. Marginal cost specification includes book fixed effects.

As shown in Table 8, the average royalty estimate in our main specification is 39 percent, which is higher than the 30 percent that was used during the first agency period. To analyze how sensitive our bargaining parameter estimates are to the royalty estimates, Column (C) of Table 10 gives estimates of the demand and supply parameters of the model when we estimate the model using the same covariance restrictions as we use for the main specification, but now fix the royalties for both Amazon and Barnes & Noble to 30 percent instead of estimating them. Comparing the estimates to those in specification (B) of Table 7, two things stand out. First of all, the estimated price coefficient is more negative—everything else given, with more price sensitive consumers it becomes optimal to set lower royalties. Secondly, the average estimated bargaining parameter increases for Amazon, but decreases for Barnes & Noble, which corresponds to a decrease in bargaining power for Amazon but an increase for Barnes & Noble. In comparison to the royalty estimates from the main

respectively, and $\hat{\sigma}_N$ is the estimated standard deviation of the difference between $\hat{Q}_1$ and $\hat{Q}_2$ (estimated using 100 bootstrap replications). The Rivers and Vuong test statistic has to be evaluated against a standard normal distribution, which, given the test-statistic of -200.51, results in a p value of 0.000, so we can strongly reject the take-it or leave-it model against the bargaining model.

specification, Amazon's royalty is substantially lower, which implies Amazon has less bargaining power than in our main specification. In general, the lower the royalties, the more negative the estimated price parameter and the higher the estimated bargaining weights. For instance, if both retailers obtain a 20 percent royalty, the estimates price coefficient drops to -3.330 and the average upstream bargaining parameter faced by Amazon increases from an average of 0.208 to 0.344.

In column (D) of Table 10, we estimate demand and supply separately. To estimate the demand side we use two-stage least squares, using lagged prices as an instrument for prices. 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 well as online marketplaces (Jin, Lu, Zhou, and Fang, 2021). As can be seen in column (D) of Table 10, the estimated price coefficient decreases in magnitude in comparison to our main specification. Also shown in the table are the estimated supply side parameters, which are estimated in a separate step by minimizing the sum of squared residuals of the marginal cost equation, taking the demand parameters as given. The estimated bargaining parameters are again close to those for the main model, which suggests estimation of the supply side parameters of the model is robust to using a more traditional IV-based approach to estimate the model.

Table 11: Supply side estimates for different price parameters

	(A)		(B)		(C)	
Variable	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Price parameters						
log(price)						
Amazon	-1.800		-2.200		-2.200	
B&N	-1.800		-1.800		-2.200	
Bargaining parameters Amazon						
Harper Collins	0.340	(0.097)	0.353	(0.128)	0.353	(0.125)
Hachette	0.249	(0.041)	0.234	(0.052)	0.232	(0.048)
Simon & Schuster	0.086	(0.082)	0.011	(0.053)	0.011	(0.078)
Macmillan	0.137	(0.086)	0.082	(0.079)	0.081	(0.090)
Penguin Random House	0.180	(0.028)	0.131	(0.029)	0.134	(0.033)
Average	0.199		0.162		0.162	
Bargaining parameters Barnes & Noble						
Harper Collins	0.447	(0.076)	0.446	(0.074)	0.490	(0.102)
Hachette	0.384	(0.024)	0.383	(0.023)	0.405	(0.032)
Simon & Schuster	0.408	(0.098)	0.408	(0.089)	0.438	(0.111)
Macmillan	0.433	(0.091)	0.432	(0.083)	0.471	(0.105)
Penguin Random House	0.359	(0.019)	0.361	(0.019)	0.370	(0.022)
Average	0.406		0.406		0.435	
Objective function	0.042		0.036		0.041	
Own price elasticity	-1.800		-2.000		-2.200	
Number of observations	2,800		2,800		2,800	

Notes: Bootstrapped standard errors shown in parentheses (clustered by book title). Marginal cost specification includes book fixed effects.

To get a better sense of how the uncertainty related to the price parameters affects the supply side estimates, in Table 11 we take a different approach and calibrate the price coefficients instead of estimating them. We then estimate the supply side model only by minimizing the sum of squared (upstream) marginal cost residuals. In column (A) of this table we fix the price parameter to -1.8 for both Amazon and Barnes & Noble, which corresponds to the estimate we obtained from our main specification. Although the estimation approach is different (i.e., minimizing the sum of squared marginal cost residuals instead of targeting zero cross-covariance between demand and supply side residuals), we do get very similar bargaining parameter estimates. Next, we keep the price parameter for Barnes & Noble fixed at -1.8, but change the coefficient for Amazon to -2.2, which corresponds to Amazon's internal estimate of the average price elasticity for e-books.[30] As shown in column (B) of the table, more elastic demand at Amazon leads to lower bargaining weight estimates on average, which implies more bargaining power for Amazon. When changing both price coefficients to -2.2, as in column (C) of the table, both Amazon and Barnes & Noble are estimated to have more bargaining power relative to the publishers. Nevertheless, the differences with the estimates for the main specification are relatively small, which suggests the supply side estimates are relatively robust to variation in the price parameters.

5.5 Counterfactual Analysis of the Most Favored Nations Clause

The settlements between DOJ and the Big Five publishers in 2012 explicitly banned the use of retail price MFN clauses for a period of five years. In this section we study what happens to agency prices when retail price 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 the largest publishers are again using the agency model and MFN clauses were not instrumental for making the switch to this second period of agency pricing, this does not mean that MFN clauses are 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, if it raises the royalty for one publisher, the retail price will remain the same relative to other retailers. This encourages retailers to push for higher royalties, which results in higher retail prices (see also Johnson, 2017).

To simulate what happens to retail prices when MFN agreements are used, we use the estimates from specification (B) 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 retail pricing first-order condition but also the bargaining first-order condition, assuming that contracts between publishers and retailers are renegotiated. The implications of using MFN on prices, royalties, margins, and profits are shown in Table 12. The table shows that royalties increase across the board, with Barnes & Noble seeing the biggest changes

[30] According to a message from the Amazon Books Team to the reader published at readersunited.com in 2014, "for every copy an e-book would sell at \$14.99, it would sell 1.74 copies if priced at \$9.99," which implies an own-price elasticity of $\%\Delta Q/\%\Delta P = 74/(100 \times (-5/14.99)) = -2.2$.

Table 12: Prices, royalties, margins, and profits MFN

	No MFN				MFN			
	Price	Royalty	Margin	Profits	Price	Royalty	Margin	Profits
Retailers								
Amazon	10.51	0.451	4.75	586,661	11.29	0.476	5.15	561,508
Barnes & Noble	10.54	0.351	3.70	84,267	11.29	0.455	6.05	122,063
Publishers								
Harper Collins	9.85	0.341	3.60	22,517	10.90	0.423	3.37	18,420
Hachette	9.60	0.382	3.29	38,788	10.52	0.479	3.04	32,727
Simon & Schuster	10.55	0.422	3.39	7,734	10.96	0.544	3.15	7,075
Macmillan	11.32	0.400	3.77	12,102	12.87	0.508	3.96	9,836
Penguin Random House	10.79	0.411	3.55	395,577	11.45	0.445	2.66	374,213

Notes: Estimates are based on the bargaining model estimates reported in specification (B) of Table 7. Figures in the Margin columns reflect downstream margins for the retailers and upstream margins for the publishers.

among the retailers and Hachette, Simon & Schuster, and Macmillan among the publishers. Even though retail prices go up as well, by an average of seven percent across titles, this price increase is not enough to prevent margins from going down for all but one of the publishers. Moreover, there is a lot of variation in price changes across the publishers, ranging from a sales-weighted average price increase of 3 percent for Penguin Random House to 13 percent for Macmillan. Table 12 also shows that publishers' profits decrease, which is driven by the higher share of revenue going to the retailers combined with higher retail prices. Overall retail profits are up when MFN clauses are used, although only Barnes & Noble benefits from this increase. Moreover, combined industry profits are down, so even though Barnes & Noble would benefits substantially from reinstatement of MFN clauses, this might not be enough for MFN agreements to make a comeback.

Table 13: Most-favored nation clause difference-in-differences analysis

	Amazon		Barnes & Noble	
	Publisher fixed effects	Book fixed effects	Publisher fixed effects	Book fixed effects
Difference-in-differences estimator				
Agency	0.143	0.137	-0.023	-0.039
	(0.009)	(0.008)	(0.007)	(0.006)
Agency × MFN	0.078	0.073	0.121	0.114
	(0.013)	(0.012)	(0.015)	(0.012)
Other controls				
Rating on Amazon	0.033	-0.023	0.045	-0.026
	(0.010)	(0.010)	(0.010)	(0.008)
Months since release	-0.002	0.002	-0.002	0.001
	(0.000)	(0.000)	(0.000)	(0.000)
Number of pages	0.048		0.068	
	(0.044)		(0.043)	
Constant	2.056	2.219	2.168	2.431
	(0.046)	(0.040)	(0.047)	(0.032)
R-squared	0.161	0.675	0.074	0.653
Number of observations	255,194	255,239	250,434	250,443

Notes: Dependent variable is log(*price*). All specifications include week fixed effects. Standard errors (clustered by book) in parentheses. The number of pages is divided by thousand.

Although for the main analysis we only include observations for the second period in which the industry shifted to agency contracts (2014-2015), for some of the books in our sample we also have price data for the first period in which agency contracts were used (2010-2012). Table 13 provides results using a similar difference-in-differences approach as in Section 4.2, but with the sample expanded to include the first agency period. Since MFN clauses were used in first agency period but not in the second, the expanded sample allows us to directly compare the effect of MFN on agency prices. As shown by the estimate on the Agency × MFN coefficient, prices where between 7 and 8 percent higher during the first agency period for Amazon, while 12 to 13 percent higher for Barnes & Noble. Note that these differences are somewhat higher than predicted by the counterfactual exercise. A potential explanation for the difference is the alleged collusion during the first agency period, which is captured by the difference-in-difference comparison between the two periods, but is not picked up by the counterfactual simulations. The higher price increase in the first agency period is consistent with this, and also shows the usefulness of our structural analysis, since it allows us to separate the effect of MFN from the effect of collusion.

6 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 from 2014 to 2015. Using a difference-in-differences analysis, we have shown that prices went up by 14 percent following the switch at Amazon, but went down 2 percent at Barnes & Noble. We have theoretically shown that if an upstream and downstream firm are bargaining over input prices, retail prices will be higher or lower under agency depending on the relative bargaining power of the firms.

Our structural model extends this theoretical model to allow for competition among publishers and retailers, multi-product firms, and 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 retailers have most of the bargaining power, although there are large differences in estimated bargaining weights 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 seven percent. Nevertheless, the counterfactual analysis also shows that adoption of MFN would lower profits of the publishers as well as Amazon, which might be a contributing factor for why MFN contracts so far have not made a return.

What complicates our analysis of the market for e-books is that the objective function of the publishers and retailers has not always been clear. Our structural model assumes static profit maximization within a single market, which means complementaries with other products as well as dynamic considerations are not directly taken into account. For instance, Amazon not only sells e-books, but also the hardware that goes with it, which may affect pricing or royalty decisions (see Gaudin and White, 2014). Dynamic competition could be important too, especially since consumers may face switching costs when moving between the retailers' different e-book platforms (see Johnson, 2020). Furthermore, we know from the 2012 lawsuit that the publishers were mostly concerned about how low e-book prices affected the market for hardcover books and supported the move to agency despite earning less from agency contracts than they were from wholesale contracts. Another limitation of our study is that litigation of the initial agency contracts has likely affected subsequent developments in this industry, so our results may not necessarily generalize to other contexts in which contracts are not litigated. Despite these limitations, our study shows the importance of the interaction of bargaining power and contract types (agency or wholesale) in determining retail prices.

References

Abhishek, Vibhanshu, Kinshuk Jerath, and Z. John Zhang: "Agency Selling or Reselling? Channel Structures in Electronic Retailing," *Management Science* 62, 2258–2280, 2015.

Bagnoli, Mark and Ted Bergstrom: "Log-concave Probability and Its Applications," *Economic Theory* 26, 445–469, 2005.

Baker, Jonathan: "Cartel Ringmaster or Competition Creator? The Ebooks Case Against Apple," in John E. Kwoka Jr. and Lawrence J. White (Editors), *The Antitrust Revolution*, 7th edition, 2018.

Berry, Steven, James Levinsohn, and Ariel Pakes: "Automobile Prices in Market Equilibrium," *Econometrica* 63, 841–890, 1995.

Bonnet, Céline and Pierre Dubois: "Identifying Two Part Tariff Contracts with Buyer Power: Empirical Estimation on Food Retailing," Mimeo, 2015.

Cachon, Gérard P. and Martin A. Lariviere: "Supply Chain Coordination with Revenue-Sharing Contracts: Strengths and Limitations." *Management Science* 51, 30–44, 2005.

Carlton, Dennis W. and Jeffrey M. Perloff: *Modern Industrial Organization*, Pearson, 2004.

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," *Journal of Political Economy* 127, 163–195, 2019.

Condorelli, Daniele, Andrea Galeotti, and Vasiliki Skreta: "Selling through Referrals," *Journal of Economics & Management Strategy* 27, 669–685, 2018.

Conlon, Christopher T. and Julie Holland Mortimer: "Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence," forthcoming in the *Journal of Political Economy.*

Crawford, Gregory S. and Ali Yurukoglu: "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review* 102, 643–685, 2012.

Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu: "The Welfare Effects of Vertical Integration in Multichannel Television Markets," *Econometrica* 86, 891–954, 2018.

Davidson, Carl: "Multiunit Bargaining in Oligopolistic Industries," *Journal of Labor Economics* 6, 397–422, 1988.

De los Santos, Babur and Matthijs R. Wildenbeest: "E-book Pricing and Vertical Restraints," *Quantitative Marketing and Economics* 15, 85–122, 2017.

Dennis, Anthony J.: "Most Favored Nation Contract Clauses under the Antitrust Laws," *University of Dayton Law Review* 20, 821–854, 1995.

Dobson, Paul W. and Michael Waterson: "Countervailing Power and Consumer Prices," *Economic Journal* 107, 418–430, 1997.

Donna, Javier D., Pedro Pereira, Andre Trindade, and Renan C. Yoshida: "Direct-to-Consumer Sales by Manufacturers and Bargaining," Mimeo, 2021.

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.

Gans, Joshua S.: "Mobile Application Pricing," *Information Economics and Policy* 24, 52–59, 2012.

Gaudin, Germain and Alexander White: "On the Antitrust Economics of the Electronic Books Industry," Mimeo, 2014.

Gilbert, Richard J.: "E-books: A Tale of Digital Disruption," *Journal of Economic Perspectives* 29, 165–184, 2015.

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.

Ho, Kate and Robin S. Lee: "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets," *American Economic Review* 109, 473–522, 2019.

Horn, Henrick and Asher Wolinsky: "Bilateral Monopolies and Incentives for Merger," *RAND Journal of Economics* 19, 408–419, 1988.

Hristakeva, Sylvia: "Vertical Contracts with Endogenous Product Selection: An Empirical Analysis of Vendor-Allowance Contracts," Mimeo, 2020.

Iozzi, Alberto and Tommaso Valletti: "Vertical Bargaining and Countervailing Power," *American Economic Journal: Microeconomics* 6, 106–135, 2014.

Jin, Ginger Zhe, Zhentong Lu, Xiaolu Zhou, and Lu Fang: "Flagship Entry in Online Marketplaces," Mimeo, 2021.

Johnson, Justin: "The Agency Model and MFN Clauses," *Review of Economic Studies* 84, 1151–1185, 2017.

Johnson, Justin: "The Agency and Wholesale Models in Electronic Content Markets," *International Journal of Industrial Organization* 69, Article 102581, 2020.

Francine Lafontaine: "Agency Theory and Franchising: Some Empirical Results," *RAND Journal of Economics* 23, 263–283, 1992.

Li, Hui: "Are E-Books a Different Channel? Multichannel Management of Digital Products," *Quantitative Marketing and Economics* 19, 179–225, 2021.

Lia, Jia and Charles C. Moul: "Who Should Handle Retail? Vertical Contracts, Customer Service, and Social Welfare in a Chinese Mobile Phone Market," *International Journal of Industrial Organization* 39, 29–43, 2015.

MacKay, Alexander and Nathan H. Miller: "Estimating Models of Supply and Demand: Instruments and Covariance Restrictions," Mimeo, 2021.

Mantovani, Andrea, Claudio Piga, and Carlo Reggiani: "Online Platform Price Parity Clauses: Evidence from the EU Booking.com Case," *European Economic Review* 131, Article 103625, 2021.

Mortimer, Julie Holland: "Vertical Contracts in the Video Rental Industry," *Review of Economic Studies* 75, 165-199, 2008.

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.

O'Brien, Daniel P., "The Uniform Settlements Policy in International Telecommunications: A Noncooperative Bargaining Model of Intermediate Product Third Degree Price Discrimination," doctoral dissertation, Northwestern University, 1989.

O'Brien, Daniel P., "The Welfare Effects of Third-Degree Price Discrimination in Intermediate Good Markets: The Case of Bargaining," *RAND Journal of Economics* 45, 92–115, 2014.

O'Brien, Daniel P. and Greg Shaffer, "Bargaining, Bundling, and Clout: The Portfolio Effects of Horizontal Mergers," *RAND Journal of Economics* 36, 573–595, 2005

Reimers, Imke and Joel Waldfogel: "Throwing the Books at Them: Amazon's Puzzling Long-Run Pricing Strategy," *Southern Economic Journal* 83, 869–885, 2017.

Rivers, Douglas and Quong Vuong: "Model Selection Tests for Nonlinear Dynamic Models," *Econometrics Journal* 5, 1–39, 2002.

Rubinstein, Ariel: "Perfect Equilibrium in a Bargaining Model," *Econometrica* 50, 97–109, 1982.

Sheu, Gloria and Charles Taragin: "Simulating Mergers in a Vertical Supply Chain with Bargaining," forthcoming in *RAND Journal of Economics*, 2021.

Shiller, Benjamin R.: "Digital Distribution and the Prohibition of Resale Markets for Information Goods," *Quantitative Marketing and Economics* 11, 403–435, 2013.

Villas-Boas, Sofia Berto: "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data," *Review of Economic Studies* 74, 625–652, 2007.

Villas-Boas, J. Miguel and Russell S. Winer: "Endogeneity in Brand Choice Models," *Management Science* 45, 1324–1338, 1999.

Weyl, E. Glen and Michal Fabinger, "Pass-through as an Economic Tool: Principles of Incidence Under Imperfect Competition," *Journal of Political Economy* 121, 528–583, 2013.

A Proof of Proposition 1

Proof. We show that the wholesale model leads to a strictly higher price than the agency model when the upstream firm has all of the bargaining power and a strictly lower price when the opposite is true. The proposition then follows from the continuity of the equilibrium prices with respect to λ.

Suppose first that the upstream firm has all the bargaining power, i.e., $\lambda = 1$. Using equation (4), the first order condition in the wholesale model is

$$p^w - c^U - c^D = \phi(p^w)\left(2 - \phi'(p^w)\right) \tag{A1}$$

where the subscript 'w' denotes the equilibrium price in the wholesale model. Note that this condition corresponds to the first order condition in the take-it or leave-it case analyzed by Johnson (2017). Using equation (13), the first-order condition for the agency model is

$$p^a - c^U - c^D = \phi(p^a)\left(\frac{c^U}{p^a - \phi(p^a)}\right) \tag{A2}$$

where the subscript 'a' denotes the agency price. We show by contradiction that the agency price must be lower than the wholesale price.

Suppose not, i.e., suppose $p^a \geq p^w$. Because the slope of the right hand side of equation (A1) is less than 1 by assumption, there is a unique solution to equation (A1). Further, because $p^a \geq p^w$, it must be true that $p^a - c^U - c^D \geq \phi(p^a)\left(2 - \phi'(p^a)\right)$. Combining this with equation (A2) yields

$$\phi(p^a)\left(\frac{c^U}{p^a - \phi(p^a)}\right) \geq \phi(p^a)\left(2 - \phi'(p^a)\right),$$

Using $1 - r = c^U/(p^a - \phi'(p^a))$, this gives $1 - r \geq 2 - \phi'(p^a)$, which is a contradiction because $1 - r \leq 1$ and $2 - \phi'(p^a) > 1$. This establishes that $p^a < p^w$ when the upstream firm has all the bargaining power.

Next, consider the case in which he downstream firm has all the bargaining power, i.e., $\lambda = 0$. The first-order conditions for this case are

$$p^w - c^U - c^D = \phi(p^w) \tag{A3}$$

and

$$p^a - c^U - c^D = \phi(p^a)\left(1 + c^U \frac{p^a(1 - \phi'(p^a))}{(p^a - \phi(p^a))^2}\right). \tag{A4}$$

Note that the condition for p^a corresponds to agency take-it or leave-it case analyze by Johnson (2017). Proceeding again by contradiction, suppose $p^a \leq p^w$. Because the right hand side of equation (A3) has slope less than 1, there is a unique solution for p^w, and the supposition that

$p^a \leq p^w$ implies $p^a - c^U - c^D \leq \phi(p^a)$. Combining this with equation (A4) yields

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

But since $\phi'(p^a) < 1$, all terms in parentheses on the left hand side of this inequality are positive, which means that the term between brackets exceeds one, which yields a contradiction. Thus, $p^a > p^w$ when the buyer has all the bargaining power. ■

B Price Derivatives

Wholesale Model

From equation (18) (swapping k and j) we have

$$\frac{\partial s_j}{\partial p_k} = \begin{cases} \alpha s_j (1 - s_j)/p_j & \text{if } k = j; \\ -\alpha s_k s_j / p_k & \text{if } k \neq j. \end{cases}$$

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

$$\frac{d\pi_k^U}{dw_j} = \frac{\partial \pi_k^U}{\partial w_j} + \sum_{k=1}^{N} \frac{\partial \pi_j^U}{\partial p_k} \frac{\partial p_k^*}{\partial w_j}, \quad \text{where} \quad \frac{\partial \pi_k^U}{\partial w_j} = \begin{cases} s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases} \tag{A5}$$

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

$$\frac{\partial \pi_j^U}{\partial p_k} = \begin{cases} m_j^U \alpha s_j (1 - s_j)/p_j & \text{if } k = j, \\ -m_j^U \alpha s_k s_j / p_k & \text{if } k \neq j. \end{cases}$$

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

$$\frac{\partial \pi_j^D}{\partial w_j} = \frac{\partial \pi_k^D}{\partial w_j} + \sum_{l=1}^{N} \frac{\partial \pi_j^D}{\partial p_l} \frac{\partial p_l^*}{\partial w_j}, \quad \text{where} \quad \frac{\partial \pi_k^D}{\partial w_j} = \begin{cases} -s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases} \tag{A6}$$

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

$$\frac{\partial \pi_j^D}{\partial p_k} = \begin{cases} s_j + m_j^D \alpha s_j (1 - s_j)/p_j & \text{if } k = j, \\ -m_j^D \alpha s_k s_j / p_k & \text{if } k \neq j. \end{cases}$$

The price derivatives $\partial p_l^* / \partial w_j$ are derived by totally differentiating the retail price-first order conditions in equation (17). The solution is

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

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

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

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

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

The own-price and cross-price derivatives are given in equation (18). The second derivatives are given by

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

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

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

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

$$\frac{\partial^2 \pi^D_j}{\partial p_k \partial w_l} = \begin{cases} -\alpha s_j\left(1 - s_j\right)/p_j & \text{if } j = k = l, \\ \alpha s_k s_j / p_k & \text{if } j \neq k,\ j = l \\ 0 & \text{otherwise.} \end{cases}$$

Agency Model

The total derivative $d\pi^U_k / dr_j$ is given by

$$\frac{d\pi^U_k}{dr_j} = \frac{\partial \pi^U_k}{\partial r_j} + \sum_{k=1}^{N} \frac{\partial \pi^U_j}{\partial p_k}\frac{\partial p^*_k}{\partial r_j}, \quad \text{where} \quad \frac{\partial \pi^U_k}{\partial r_j} = \begin{cases} -p_k s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases}$$

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

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

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

$$\frac{\partial \pi_j^D}{\partial r_j} = \frac{\partial \pi_k^D}{\partial r_j} + \sum_{l=1}^{N} \frac{\partial \pi_j^D}{\partial p_l} \frac{\partial p_l^*}{\partial r_j}, \quad \text{where} \quad \frac{\partial \pi_k^D}{\partial r_j} = \begin{cases} p_k s_k & \text{if } k = j, \\ 0 & \text{if } k \neq j. \end{cases} \tag{A7}$$

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

$$\frac{\partial \pi_j^D}{\partial p_k} = \begin{cases} r_j s_j + m_j^D \alpha s_j (1-s_j)/p_j & \text{if } k = j, \\ m_j^D \alpha s_k s_j / p_k & \text{if } k \neq j. \end{cases}$$

The price derivatives $\partial p_l^* / \partial r_j$ are derived by totally differentiating the retail price-first order conditions in equation (23). 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:

$$\frac{\partial^2 \pi_j^U}{\partial p_k \partial r_l} = \begin{cases} -s_j[1+\alpha(1-s_j)] & \text{if } j=k=l, \\ \alpha s_j s_k & \text{if } j \neq k,\ j=l \\ 0 & \text{otherwise.} \end{cases}$$

C Derivation $\mathbf{m^U}$ (Wholesale) and $(\mathbf{1}-\lambda)/\lambda$ (Agency)

Derivation $\mathbf{m^U}$ (Wholesale)

Equation (22) 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 (22) as

$$E_j^D \left(\pi^U - d^U\right) \frac{1-\lambda}{\lambda} + \frac{\partial \pi^U}{\partial w_j} = 0, \tag{A8}$$

where $E_j^D = \left(\pi^D - d^D\right)^{-1} \left(\partial \pi^D / \partial w_j\right)$. Using equation (A5) and taking into account the ownership structure, we can write $\partial \pi^U / \partial w_j$ as

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

which, using matrix notation, can be written as $s + (T^U Z^w) m^U$. Taking into account the ownership structure and using $\pi^U - d^U = (T^U \cdot S) m^U$, we can write the bargaining first-order condition in equation (A8) as

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

Solving for m^U gives

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

To derive an expression for E^D, first note that we can write equation (A6) as

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

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

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

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

$$Z^w_{jk} = \alpha s_k \left(1 - s_k\right)/p_k \frac{\partial p^*_k}{\partial w_j} - \sum_{l \neq k} \alpha s_l s_k / p_l \frac{\partial p^*_l}{\partial w_j},$$

Derivation $(1 - \lambda)/\lambda$ (Agency)

The bargaining first-order condition (24) relates the bargaining parameters to the upstream and downstream margins, which can be used to obtain the bargaining parameters as a function of the margins. First, rewrite equation (24) as

$$E^U_j \left(\pi^D - d^D\right) \frac{\lambda}{1-\lambda} + \frac{\partial \pi^D}{\partial r_j} = 0, \tag{A9}$$

where $E^U_j = \left(\pi^U - d^U\right)^{-1} (\partial \pi^U / \partial r_j)$. Using equation (A7) and taking into account the ownership structure, we can write $\partial \pi^D / \partial r_j$ as

$$\frac{\partial \pi^D}{\partial r_j} = A_j + \sum_{k \in \Omega^D} \left(B^D_{jk} + m^D_k Z^r_{jk}\right),$$

where $A_j = p_j s_j$ and $B^D_{jk} = r_k s_k (\partial p^*_k / \partial r_j)$. In matrix notation this is

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

Taking into account the ownership structure and using $\pi^D - d^D = (T^D \cdot S) m^D$, we can write the bargaining first-order condition in equation (A9) as

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

Solving for $L = (1-\lambda)/\lambda$ gives

$$L = -(A + T^D \cdot B^D + (T^D \cdot Z^r) m^D)^{-1} (E^U (T^D \cdot S) m^D). \tag{A10}$$

where E^U is a vector whose jth element is given by $E^U_j = \left(\pi^U - d^U\right)^{-1} (\partial \pi^U / \partial r_j)$, 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

$$Z^r_{jk} = \alpha s_k \left(1 - s_k\right)/p_k \frac{\partial p^*_k}{\partial r_j} - \sum_{l \neq k} \alpha s_l s_k / p_l \frac{\partial p^*_l}{\partial r_j}.$$

D Sales Rank Data

Following Chevalier and Goolsbee (2003), we assume that book sales quantities follow a Pareto distribution, i.e., the probability that an observation τ exceeds a level of Sales is

$$\Pr(\tau > \text{Sales}) = \left(\frac{k}{\text{Sales}}\right)^{\theta}, \tag{A11}$$

where k and θ are the scale and shape parameters of the Pareto distribution. Since the fraction of books that have more sales than a particular title is (Rank − 1)/(Total number of books), we can write equation (A11) 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

$$\log(\text{Sales}) = \gamma_0 + \gamma_1 \log(\text{Rank} - 1), \tag{A12}$$

where $\gamma_0 = \log(k) + \frac{1}{\theta}(\text{Total number of books})$ 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 daily Kindle sales data.[31] OLS estimates of equation (A12) give us a coefficient θ of 1.19, which is very close to the estimate of 1.2 that Chevalier and Goolsbee (2003) use throughout their paper.[32] We use the fitted sales to transform Kindle sales rank data into quantity data. To obtain Barnes & Noble sales we use a title's observed Barnes and & Noble sales ranks in the same estimated equation, but shift the intercept to reflect that aggregate e-book sales at Barnes & Noble are approximately one-quarter of those at Amazon.[33] To obtain weekly sales we multiply the daily sales data for each of the retailers by seven.

[31]See https://kindlepreneur.com/amazon-kdp-sales-rank-calculator.

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

[33]Although precise figures are not available, according to their sales data the digital book publishing platform Vook (rebranded as Pronoun in 2015, acquired by Macmillan in 2016, and discontinued in 2018) estimates Amazon's market share to be 60 percent, while Barnes & Noble's market share is 15 percent. See tinyurl.com/wynynw6.