

17 | The Economics of Revenue Sharing: Theory and Measurement

The accompanying case, “The Home Video Industry,” discusses the evolution of the home video industry. One important event in the industry’s history was the introduction of revenue sharing contracts in 1998. This technical note has two goals. First, to understand the economic rationale for revenue sharing. Second, to explain some of the pitfalls in the measurement of the effects of revenue sharing contracts.

The economics of revenue sharing

To understand why revenue sharing contracts make economic sense, let us consider a simple example, illustrated by Figure 17.1. A given video rental store must decide how many DVDs of a given title it should keep in the store: 1, 2 or 3. The greater the number of DVDs, the greater the number of checkouts.

The relation between the number of DVDs and the number of checkouts is typically less than proportional: most of the time, one DVD is sufficient to satisfy demand. Occasionally, two customers arrive at the store the same day and looking for the same title. In this event, having two copies instead of one leads to a greater number of checkouts. In the even more unlikely event that three customers come to the store the same day and asking for the same title, then holding three DVDs increases the number of checkouts with respect to holding two DVDs only.

Estimating the relation between the number of holdings and the number of checkouts is a complex operations management problem which we will not solve here. We simply take the result as given — specifically, as given by Figure 17.1: 30 yearly checkouts with a stock of one tape, 42 yearly checkouts with a stock of two tapes, and 44 yearly checkouts with a

stock of three tapes. The fact the third number is so close to the second one reflects the low probability of three customers simultaneously requesting the same title.

We also take as given the average price of a rental, which we assume is 3. Given the price of a rental and the number of checkouts, we can compute gross revenue as a function of DVD holdings. This is given by the second row in Figure 17.1.

□ **Retailer's choice.** The retail store must decide whether to acquire DVDs by purchase or by signing a revenue sharing agreement. Under a purchase contract, the store pays the studio 60 per DVD and keeps all of the rental revenues. This leads to the profit distribution shown in Figure 17.1. For example, if the store selects to have two DVDs, then its profit is given by $126 - 2 \times 60 = 6$. The studio's revenue, in turn, is given by $2 \times 60 = 120$. Total joint profit is simply gross revenue (126) since we are assuming (not a bad approximation) that the physical cost of an extra DVD copy is close to zero.

Consider now the possibility of revenue sharing. Suppose that the contract terms are an upfront fee of 7 and a 50% share of sales revenues. If, again, the store decides to hold two DVD copies, then its profit is now given by $50\% \times 126 - 2 \times 7 = 49$. The studio, in turn, makes $50\% \times 126 + 2 \times 7 = 77$.

We finally come the retailer's optimal choice. This is simply the choice of contract and number of DVDs that maximizes retailer's profit. Direct inspection shows that this is given by two DVDs under revenue sharing, yielding a profit of 49. A different number of DVDs would yield the retailer a lower profit (38 or 45). Shifting to purchase would make the retailer even worse.

As to the studio, having the retailer choose revenue sharing is also a good thing. Notice that, if the retailer were to purchase two DVDs, then the studio would make a profit of 120, greater than what it makes by revenue sharing two DVD. However, were the retailer to choose the purchase option, it would purchase only one DVD: as Figure 17.1 shows, retailer profit is greater with the one DVD purchase option than with two or three.

□ **Incentives.** To summarize the above analysis: by switching from purchase to revenue sharing, both retailer and studio become better off. Magic? Yes and no. The key word is incentives: As Figure 17.1 shows, joint profits (that is, studio plus retailer) are greater the greater the number of copies held in the store. This makes sense: it costs nothing (or next to nothing) to produce additional copies, and there is always the possibility of an extra sale in the event that many customers arrive at the store simultaneously.

The problem is that the store makes its choice of number of DVDs thinking about its profits, not joint profits with the studio. The beauty of revenue sharing is that it (partly) aligns retailer and studio's interests, and consequently moves the choice of number of DVDs closer to the industry optimum.

Close but not quite. As we saw above, under revenue sharing the retailer chooses to hold two DVD copies. However, the jointly optimal number is three. Revenue sharing does not perfectly align retailer and studio incentives. One possibility to further improve the retailer and the studio's situation is for the studio to *force* the retailer to choose three copies if it wants to go with revenue sharing. Why does this work? Because the retailer, while preferring two DVDs to three DVDs under revenue sharing, prefers three DVDs under revenue sharing (a 45 profit) to the best option under purchase (a 30 profit). In fact, Rentrak contracts (exemplified in the accompanying case study) frequently included output forcing

Figure 17.1

Purchase vs revenue sharing.

Number of copies in store	1	2	3
Number of checkouts per year	30	42	44
Gross revenue	90	126	132
Profit levels under purchase			
Store	30	6	-48
Studio	60	120	180
Joint	90	126	132
Profit levels under revenue sharing			
Store	38	49	45
Studio	52	77	87
Joint	90	126	132

Underlying assumptions: Sale price: 60; Revenue sharing upfront fee: 7; Store revenue share: 0.5; Rental price: 3.

clauses of this kind.

□ **Costs and benefits.** So what's the catch, one might ask? If revenue sharing contracts are so good, what took this industry so long to implement them? One important consideration is that the cost of implementing revenue sharing is considerably greater than the cost of implementing a simple sales transaction. In the latter, one transaction takes place at one point in time and no additional information or monitoring is necessary. Revenue sharing, by contrast, requires that studio (or an intermediary between the studio and the retailer) monitor sales.

Measuring the impact of revenue sharing

Having established the economics reasoning behind revenue sharing, the next question is how to measure its effects. One possibility is to compare studio and retailer profitability under purchase agreements and under revenue sharing. The theory described in the previous section predicts that both are higher under revenue sharing.

Figure 17.3 shows profitability per title when the contract form was purchase and when it was revenue sharing. The data are divided by title type: A, B and C, in decreasing order of box-office revenues (A: greater than \$40m; B: \$15–40; C: \$0.25–15m). As expected, profits are greater for A titles than for B titles, and for B titles than for C titles: rental-market popularity is proportional to box-office popularity. Moreover, for a given movie type, both retailer profit and distributor profit are greater under purchase than under revenue sharing.

This result runs contrary to the theory prediction. The problem is that a direct comparison of profits per title is potentially problematic. To understand why, consider Figure 17.2, which replicates the analysis of Figure 17.1 for the case of a high-demand title. Specifically, all values of rental sales are doubled with respect to Figure 17.1. When we now compare

Figure 17.2

Purchase vs revenue sharing: the case of a high-demand title.

Number of copies in store	1	2	3
Number of checkouts per year	60	84	88
Gross revenue	180	252	264
Profit levels under purchase			
Store	30	6	-48
Studio	60	120	180
Joint	90	126	132
Profit levels under revenue sharing			
Store	120	132	84
Studio	60	120	180
Joint	180	252	264

Underlying assumptions: Sale price: 60; Revenue sharing upfront fee: 7; Store revenue share: 0.5; Rental price: 3.

the retailer's profit levels under each combination of contract and number of copies, we conclude that the optimal solution is to outright purchase two copies of the DVD. The reason is that, if demand levels are higher, than the fixed cost of purchasing a DVD is relatively lower, and the purchase option becomes more attractive (since under purchase the retailer keeps 100% of rental revenues).

The problem in performing the direct, unconditional comparison in Figure 17.3 is that it does not control for variations in title popularity other than those predicted by previous box-office revenue. But there are many variations in demand in addition to those predicted by previous box-office revenue. Although these are not observed by us, they are observed by the rental stores who must choose between purchase and revenue sharing. In other words, Figure 17.3 may simply be capturing the difference between titles like that in Figure 17.1 and titles like that in Figure 17.2. The fact that profits are higher in the case of Exhibit 17.2's title (where purchase is the contract option chosen) than in the case of Exhibit 17.1's title (where revenue sharing is the contract option chosen) does not mean that switching to the purchase option increases profits — it simply means that the title corresponding to Figure 17.2 is a more popular title.

In technical terms, the problem with the unconditional approach implicit in Figure 17.3 is that it ignores *unobserved heterogeneity* across titles and so does not account for *endogeneity* in the choice of contracts, that is, the fact that contract choice depends critically on variables that are observed by the retail store but not included in the model.

□ **Controlling for unobserved heterogeneity.** In order to properly measure the effects of revenue sharing versus purchase, the statistician must follow one of several alternative paths. One is to effectively eliminate unobservable heterogeneity: to formulate a statistical model that, based on all movie characteristics, store characteristics, store neighborhood demographics and other observable variables, effectively predicts the demand for a given

Figure 17.3

Retailer and distributor profits: Comparison of purchased and revenue shared titles within store. Source: see Endnote 16.

Contract form	Purchase	Revenue sharing
A titles		
Retailer profit	539	440
Distributor profit	1,021	715
B titles		
Retailer profit	339	280
Distributor profit	618	384
C titles		
Retailer profit	131	125
Distributor profit	239	170

title. And then compare the profits of two titles with a comparable demand level, one rented under a purchase contract and one rented under revenue sharing. This approach will succeed if we can find enough observable determinants such that the residual choice between purchase and revenue sharing is nearly random.

A second possibility is to find an *instrumental variable*, that is, a variable that is correlated with the retailer's contract choice but not with movie demand. Consider an extreme example: suppose that all retailers in North Dakota must choose purchase, whereas all retailers in South Dakota must choose revenue sharing, perhaps due to state-level regulations. This would be an ideal instrument, to the extent that demand and store characteristics are unlikely to vary systematically between North and South Dakota.

Still a third possibility is to formulate a model of retailer behavior so that, by observing each retailer's choice, we may infer what the particular characteristics of the retailer and the title must be (basically market demand). Essentially, this amounts to uncovering ("reverse engineering"), for each retailer, the values leading to a table like Figures 17.1 or 17.2. The details of such a task go beyond the present technical note. Economist Julie Holland Mortimer of Harvard University did so with a sample of thousands of titles in the 1999–2001 period.¹ Based on her estimated model, she predicts that, if a retailer were forced to move back to purchasing tapes instead of revenue sharing, then the average number of copies of an A title would drop from 23 to 17. Moreover, retailer profit would decrease by 11% and distributor profit by 10%. Similar results are obtained for B and C titles. These results are consistent with the theoretical analysis in the previous section: the misalignment of incentives implied by purchase contracts (a phenomenon also known as *double marginalization*).

Endnotes

1. Julie Holland Mortimer, "Vertical Contracts in the Video Rental Industry," Harvard University, 2006.