

Collection Sales: Good or Bad for Journals?

Mark Armstrong*
Department of Economics
University College London, UK

August 2008

Abstract

This paper discusses the impact of collection sales (i.e., the bundling of several journals for sale by publishers to libraries) on journals. The advent of electronic journal distribution implies that bundling is an efficient sales strategy, and can act to extend the reach of a journal. Current arrangements are discussed and shown to lead to tensions between commercial publishers and non-profit journals. The paper argues that non-profit journals should not necessarily abandon collection sales programmes. Rather, non-profit journals may benefit from withdrawing from commercial publishers which distribute their own for-profit journals, and joining together to be distributed by less commercial publishers who set relatively low prices for their collections.

Keywords: Journal pricing; bundling; price discrimination.

JEL classification: D82, L31, L42, L82.

1 Introduction

Publishers perform several valuable tasks which an academic journal cannot usually do well on its own, such as marketing, negotiations with libraries, maintaining a website, and so on. When it acts for several journals, one function a publisher can perform is *bundling*. In this paper I discuss the pros and cons of bundling journals for sale to libraries (or “collection sales”). The perspective taken is that of an individual journal, rather than overall welfare. For instance, bundling may be used as a vehicle for exclusion of new journals or rival publishers, but this is not discussed here.¹ In particular, in this paper

*I am grateful to Ted Bergstrom, Carli Coetzee, Richard Gedye, Victor Ginsburgh, Martin Green, Jim Hosek, Jean-Charles Rochet, Chris Wallace, and two referees for many helpful comments. All remaining errors and views are my own responsibility. The support of the Economic and Social Research Council (UK) is gratefully acknowledged.

¹See Edlin and Rubinfeld (2004) and Jeon and Menicucci (2006) for discussion of these issues.

the concern about possible high bundle prices from publishers is only that a journal's circulation may not be as large as it desires, rather than that libraries may be exploited.

Library expenditure on scientific journals has significantly increased in recent years.² This has been the result of two main factors: significant merger activity among publishers, and the increased use of bundling as a marketing device as facilitated by the electronic distribution of journals. Regarding the former, Dewatripont, Ginsburgh, Legros, and Walckiers (2007) show that journal prices tend to be higher in those scientific fields where publishers are more concentrated. However, this paper focuses on the second reason for inflation in library expenditure. The introduction of bundling as a business model is a powerful force for price rises, especially when combined with a concentrated market structure in which publishers each distribute many journals. When a publisher distributes many journals, packaging these journals as a bundle acts significantly to reduce the dispersion of libraries' willingness-to-pay for the bundle compared with their willingness-to-pay for individual journals. This reduces the asymmetry of information between publisher and library, which has two effects: first, it means that more libraries will gain access to a given journal (a positive efficiency effect), and second, a commercially-minded publisher will better be able to extract a library's willingness-to-pay for the bundle (a monopoly exploitation effect). It is quite possible that bundling acts to improve overall welfare, but transfers surplus from libraries to publishers (and then possibly onto the journals themselves).³

Journals differ in their ownership and their objectives. In this paper, a journal is termed "for-profit" if it aims to maximize its profit. Usually, such journals are owned by commercial publishers. The main alternative discussed here is that of "non-profit" journals, which are assumed instead to wish to maximize their "reach" subject to a break-even constraint.⁴ (Reach can be measured in various ways, but here we take it to

²Edlin and Rubinfeld (2004) report that during the period 1984–2001 the price of library subscriptions to periodicals in law, medicine and physics rose respectively by 205%, 479% and 615%. (Overall inflation in the United States was 70% over this period.) However, the advent of widespread bundling has meant that library expenditure *per journal* has fallen from a peak in 2000. See page 13 in the Association of Research Libraries' *ARL Statistics 2005-06* (published 2008, Washington, DC.).

³The impact of bundling many products is parallel to so-called first-degree price discrimination. The textbook monopolist of widgets sets price above marginal cost because it does not have accurate information about its buyers' willingness-to-pay for its product. The result is that demand is inefficiently restricted but buyers are left with some surplus. If the monopolist had accurate information about each buyer's willingness-to-pay, it would charge each buyer her reservation price, with the result that demand is expanded to the efficient level but buyers are left with none of the gains from trade.

⁴A third kind of journal, related to the non-profit journals, is the category of society journals.

be the number of libraries which have access to the journal.) Such journals are usually owned independently of commercial publishers, or owned by less commercial publishers (such as some university presses). In practice, the line between for-profit and non-profit journals is not always so easy to draw. For instance, a journal which does not have profit as its ultimate objective may benefit when it receives more revenue, as a better-financed journal may end up publishing better articles (for instance, by being more efficient at processing submissions or by paying authors when they publish a paper in the journal), and this may maximize its reach or impact in the long term.⁵ Nevertheless, in the formal model in section 3 I ignore these subtleties, and suppose that journals are exogenously partitioned into for-profit and non-profit journals.

Subscription prices can vary hugely between these two kinds of journal. Bergstrom (2001) reports that in 2000 the six most-cited economics journals were all non-profit journals, and their (stand-alone) library subscription prices were then an average of \$180. Five of the top 20 most-cited journals were owned by commercial publishers, and their average subscription price was then \$1660 per year.⁶

In the next section, I discuss how current arrangements for collections sales lead to various tensions for non-profit journals, together with some of the options for the future of journal pricing and remuneration. In section 3 a formal model is presented in

Perhaps the main difference between a society journal and other non-profit journals is that the former may wish to ensure there is still a market for personal journal subscriptions, since this historically has been a major reason why people become members of the society. (Other methods are requiring society membership to submit articles or to attend an annual conference.) A wholehearted bundling policy – where essentially all potential members can electronically access the journal for free via their library – would be likely to reduce personal subscriptions and so reduce society membership. Thus, all else equal, a society journal is more likely to be suspicious of bundling policies. Of course, the same point applies also to site licenses for a single journal (where all members of an institution can electronically access the journal when the library has a subscription), but less strongly. With a site license, a journal will lose many of its personal subscribers at subscribing institutions, but with bundling as well it will lose many of its personal subscribers at those institutions which take the collection (as well as those at institutions with full library subscriptions).

⁵This paper does not address why journals have the aims they do. To do this, we would need to model the author side of the market as well. Authors, all else equal, prefer a wide-reaching journal to one with more limited readership. This may be because they like the idea of many people reading their paper, or because a wide-reaching journal is typically a prestigious journal. (More readers will, all else equal, result in more citations and a higher impact factor. And of course causation goes the other way too, and readers are more likely to read prestigious journals.) Thus, a journal’s ultimate aim may be to publish the “best papers”, and one way to attract the best papers is to have a wide readership. See Jeon and Rochet (2007) for more detailed discussion about how journal pricing (both to institutions and to submitting authors) and the quality of articles interact.

⁶See Dewatripont, Ginsburgh, Legros, and Walckiers (2007) for more recent data for a wider range of scientific disciplines, which also shows that for-profit journals typically are significantly more expensive than non-profit journals.

an attempt to understand better one plausible future scenario for the industry. Here, bundling is assumed to be widespread and publishers compete for journals along two dimensions: the *revenue* they offer journals and the *reach* of their collection to libraries. For-profit journals care only about the former dimension, while non-profit journals care only about the latter (as long as they receive revenue to cover their fixed costs). Publishers are assumed to make lump-sum payments to attract journals, as opposed to current arrangements which are a messy compromise between stand-alone subscription payments and *ad hoc* revenue allocations from collection sales. In this model, all journals prefer to participate in collection sales programmes, but the two kinds of journals are distributed by publishers with different pricing strategies: for-profit journals are attracted to (or are owned by) publishers who offer them high revenue but relatively low reach, while non-profit journals use publishers who market their collections with relatively low bundle prices and higher reach.

2 Tensions in Current Publishing Arrangements

There are a number of methods for pricing institutional subscriptions and remunerating journals. Historically, all journals were disseminated only in hard copy and bundling was not significantly used.⁷ The result was an extremely simple and transparent process for setting subscription prices and allocating revenue to journals. Each journal sets the price a library must pay to subscribe to the journal, regardless of which other journals the library buys. (There may be some explicit price discrimination, say to developing countries, but journals rarely practised price discrimination among institutions within developed countries.) An independent journal typically has the authority to set its subscription price (with the advice of the publisher), and so the journal has full control over its reach and revenue. An important feature of this “non-bundled” market is that a journal does not really care about which other journals its publisher distributes, or how it markets them. So long as its publisher performs its marketing tasks diligently and offers a competitive rate for its services, a particular journal (for-profit or non-profit) is indifferent about its publisher’s interactions with other journals since that has no impact on a library’s decision to subscribe to that journal.

However, the move to electronic distribution of journals makes a degree of bundling

⁷Of course, though, a journal can itself be regarded as a bundle of individual articles.

an efficient and a profitable strategy. Since it costs virtually nothing for a publisher to supply a library with an extra journal, a publisher will wish to strike a deal with a library to take all the journals to which it does not already have full subscriptions. In more detail, a publisher obviously knows which libraries have full institutional subscriptions to which of its various journals. The publisher may then offer a library access to all the publisher's additional journals (to which the library does not currently subscribe) for an additional top-up fee.⁸ Thus, a journal to which the library already subscribes continues to receive its subscription fee, plus a share (to be discussed) of the top-up fee. Similarly, if a library has not recently subscribed to the journal, the publisher may persuade the library to take its whole collection of journals, in which case the library will then have access to the journal and the journal will receive just its share of the top-up fee.

There are a number of concerns with this arrangement.⁹ First, there is a limit to how long the system can work: why should a library's subscription decisions from many years ago determine how much it should pay for a publisher's collection of journals?

Second, and relatedly, a publisher may have a difficult job policing this arrangement. What if a library cancels its stand-alone subscription to a journal at the next contractual opportunity, and subsequently has access to the journal as part of its top-up bundle? Publishers claim they can persuade libraries not to do this. For instance, if a library cancels one full subscription the publisher can threaten to increase its top-up fee by the corresponding amount the next year. However, there are two problems with this policy:

- If the (commercial) publisher has the ability to increase the top-up fee, we expect that it would do this in any case. For instance, suppose the library originally placed a value v on having access to a particular journal, while the full subscription price for this journal is p . Originally, $v > p$ and the library decided it was worthwhile subscribing. But now for some reason the value v falls below p , and the library decides to cancel its subscription. If the publisher is confident the library is willing to bear a price rise of p in the top-up fee, it seems it could have increased the price

⁸See section B3 in Edlin and Rubinfeld (2004) for more details of Elsevier's various 2002 bundling offers to libraries. More recent information about some of Elsevier's bundling offers can be seen at www.sciencedirect.info/licensing/primary (visited on August 12, 2008), which lists prices for electronic access to its list of journals for (non-research) colleges.

⁹However, one benefit of the arrangement is that a library may wish to continue with print copies of those journals to which it has subscribed in the past, in order to possess the complete print run. Such libraries will be more willing to continue with their regular stand-alone subscriptions, and this partly explains the persistence of current hybrid arrangements.

of the top-up fee before v declined. If not, then the library might cease subscribing to the whole collection when faced with the price rise, which the publisher would not wish to risk.

- Depending on the details of the journal’s contract with the publisher, carrying out this threat may benefit the publisher rather than the cancelled journal itself. If the publisher ceases to pay the journal the full subscription charge after the library has cancelled, the policy results in more revenue coming in the form of the top-up fee rather than journal-specific subscriptions. Therefore, the publisher may gain more revenue (via its own journals in the collection) that it did before, and so may have an incentive to allow libraries to cancel full subscriptions to independent journals and to load more of the total charge onto the top-up fee. However, some publishers may continue to pay the journal its full subscription revenue, at least for some years after cancellation, which mitigates this “perverse” incentive.

It is difficult to find public data to see whether journals in collection sales have better or worse institutional renewal rates than those which choose to be sold outside a publisher’s collection.¹⁰ But in any event a journal can only trust that its publisher is doing its best to maintain its base of historical subscriptions, against each library’s obvious incentive to exploit bundling discounts (and the publisher’s possible incentive to load more onto the top-up fee).

Third, for as long as this arrangement persists, it is virtually impossible for a journal to secure a new full library subscription. (Such a library will surely ask for the journal as part of the top-up package.) Thus, once a subscription has been lost, it will not come back. In addition, in a new market such as China, where there are historically few full subscriptions, a journal will be unable to generate many full subscriptions. Of course, for a really new institution the top-up fee for the bundle of all journals may be much larger than for an older library with many historical institutional subscriptions. In this case, the thorny issue of how the top-fee for the bundle is allocated among journals (see the next point) becomes all the more important.

¹⁰One data-point is *Econometrica*’s institutional subscription figures—see Table 1 in “Report of the Secretary” (2007), *Econometrica*, 75(1), page 292. Here it is reported that, starting from a base of 2438, the journal lost 124 institutional subscribers in the year leading up to the end of 2001, 93 in 2002, 3 in 2003, 189 in 2004 and 80 in 2005. The journal was sold as part of Blackwell’s collection only in the single year 2004, which is when it incurred its largest subscriber loss. However, there is so much noise in this time series that it is hard to conclude much from this with confidence.

Fourth, how should the top-up fee be allocated between journals? If a publisher's journals are fairly homogenous, then it does not seem unreasonable to divide the fee equally. (E.g., Elsevier's *Handbooks in Economics* currently all have the same cover price, and revenues from bundled electronic dissemination are allocated equally to all books.) But usually there is significant journal heterogeneity in terms of (i) subscription fees, (ii) impact and prestige, (iii) whether the journal is independent or owned by the publisher, and (iv) the number of historic full subscriptions to the journal. One arrangement is for the top-up fee to be allocated in proportion to journal stand-alone subscription fees. This arrangement is troubling for a number of reasons:

- A publisher then has an incentive to start its own low-quality journals, and to set a very high subscription fee for these. Few libraries will choose to have stand-alone subscriptions to such a journal, but the publisher will syphon off a large fraction of the top-up fee merely as a result of having these journals in its collection.
- Independent journals have an extra incentive to raise their subscription fees, purely in order to obtain a larger slice of the top-up fee, which could lead to inflationary pressure in the market.¹¹
- There is little evidence of positive correlation between a journal's subscription price and the quality of a journal in the collection – if anything, the reverse.¹²

Note that in established regions, where a journal historically has good reach, this injustice is not as great as may at first appear, at least in the short run. For so long as the

¹¹See Ginsburgh and Zang (2007) for analysis of this point.

¹²A look at Wiley's website *InterScience* (visited on August 12, 2008) yields the following data for some economics journals: *Agricultural Economics* has impact factor 0.603 and subscription fee £382; *Contemporary Economic Policy* has impact 0.521 and subscription £161; *Econometrica* has impact 2.972 and subscription £278; *Economic Inquiry* has impact 0.535 and subscription £199; *Economic Journal* has impact 1.548 and subscription £347; *Journal of Economics and Management Strategy* has impact 0.875 and subscription £231; *Journal of Economic Surveys* has impact 0.797 and subscription £427; *Journal of Industrial Economics* has impact 0.86 and subscription £146; *RAND Journal of Economics* has impact 1.44 and subscription £137, and *Review of Economic Studies* has impact 2.539 and subscription £220. (Subscription fees are the "online only" institutional subscription fee in the UK for the 2008 volume.) For further discussion of this point, see Bergstrom (2001) and Dewatripont, Ginsburgh, Legros, and Walckiers (2007).

Note, though, that impact factor is not the best measure of a library's willingness-to-pay for a journal, since it expresses a journal's impact on a per-article basis. A better measure would be impact factor multiplied by the number of articles per year. See the useful website www.eigenfactor.org for rankings of journals in terms of improved measures of per-article impact, total impact, and cost-effectiveness.

publisher is adequately policing the renewals of its institutional subscribers (the second point above), the journal will be receiving its full institutional subscription fee in any case, *plus* a small top up. For instance, consider a simple world in which there are two kinds of journal: the top tier, to which all relevant libraries subscribe on an individual basis, and a lower tier, to which all libraries consume only as part of a bundled package. In such a situation, the top tier journals do not drive the demand for the top-up bundle, and it is not clear why they should receive *any* of the top-up fee. Any revenue which a top-tier journal receives from the top-up fee might be regarded as being “paid twice” for the subscription. Nevertheless, the number of full library subscriptions is likely to continue falling (despite the publishers’ claimed best efforts to stop this). In future, therefore, the top-up fee is sure to be an ever larger fraction of even a top-tier journal’s revenue, and this allocation issue will assume ever greater importance. Moreover, in new regions where libraries are starting from scratch, this issue is important right now. An important point is that an allocation rule for a top-up fee which is broadly acceptable to an established journal in a mature market may not be remotely acceptable for allocating revenues from an entirely new institution, yet some publishers currently apply the same allocation rule in both situations.

Finally, current arrangements leave independent journals with relatively little control over their reach and revenues. In contrast to the previous regime of stand-alone journal sales, where a journal knew exactly how much it would receive from each subscribing institution, a journal is now left to the mercy of a publisher in its private negotiations with libraries. If a publisher claims to have struck a good deal to supply a consortia of, say, Chinese universities, the journals in its collection have to take that at face value (even if they accept the publisher’s rule for how it divides collection revenue).

Future arrangements: For all of these reasons, it is hard to imagine how current arrangements can persist for long, and the industry will move to other methods for remunerating journals. Some journals have reacted to these current tensions by reverting to the historical system of stand-alone pricing and distribution, and ceasing to be included in their publisher’s collection sales programme. Such a response has some advantages: it eliminates the danger of the journal cross-subsidizing the publisher’s own second-rate journals; it gives the journal full control over the price libraries pay, and it safeguards

the journal's full subscriptions when the publisher is unable or unwilling to adequately police full subscription renewals for the journals in its collection programme. However, it also eliminates all efficiency gains from the use of bundling. (These efficiency gains, which can be dramatic, are described in more detail in section 3.) For this reason, such a policy cannot be recommended for most journals.^{13,14}

Possible incremental adjustments to current arrangements include the following two options. First, the top-up fee could be divided on the basis of, say, impact factor or usage instead of subscription charges. This would largely remove the danger of cross-subsidy from prestigious non-profit journals to the publisher's own less good journals. However, in mature markets at least, there is real danger of the prestigious journals being "paid twice", which might mean that the cross-subsidy goes in the opposite direction.

Second, the allocation of the top-up fee could take into account which journals already hold a full subscription at the library. For instance, it might be that a journal with an existing subscription receives no part of the top-up fee, so that there is no danger of being paid twice. Those journals without an existing subscription might then divide the top-up fee in proportion to some measure of impact or usage relative to those other journals without subscriptions. Such an arrangement might mean that a prestigious, non-profit journal obtains a much larger payment (e.g., £50, as opposed to something like £5 or less under current arrangements) when it is included in a library's top-up collection. This arrangement would in essence implement different rules in mature and in emerging markets. In mature markets, established prestigious journals already receive the revenue from their historical full subscriptions, and for them the details of how the relatively small top-up fee is allocated are not so important. But in new markets, prestigious journals should receive an allocation of the collection fee which reflects their importance within the collection.

¹³At the time of writing the *Review of Economic Studies* has approximately 1500 full library subscriptions while about 2000 further libraries have access to the journal through collection sales. If the journal withdrew from the collection sales programme, its reach would fall significantly (even if a portion of the second group of libraries take out a full subscription in response to being unable to access the journal via collection sales).

¹⁴If the journal can price discriminate to institutions in such a way that a library with a low willingness-to-pay for the journal is offered a low charge to subscribe, this journal may be able to maximize its reach without using bundling. However, it is not clear how feasible such a policy would be. Another scenario in which bundling is not useful for a journal is when the journal is so prestigious that essentially *all* relevant institutions are willing to pay a (uniform) subscription charge which allows the journal to cover its costs. In such cases, the journal can attain maximum reach without being bundled with other journals.

More radical arrangements are also possible. Many people have suggested that an “open access” model is the best way for the market to evolve, and that journals should not charge readers for access at all but should cover their costs by levying fees on published authors (and/or by the journal’s staff offering their labour for free, or with other more explicit charitable donations). This policy has several attractions: if electronic access to journals is free then the whole vexed issue of bundling goes away; and clearly open access journals have the widest conceivable reach. Moreover, nowadays very many published articles are available in similar “pre-publication” versions on authors’ own webpages, which limits how much publishers can feasibly charge libraries for access to journals. Doubtless, some open-access journals will thrive in the long run.

But there are a number of reasons why it seems unlikely that open access will become universal in the medium term, at least without major external pressure from some source. First, the size of the author fee needed to allow major journals to break even could be large, and these fees could deter some people from publishing their research.¹⁵ Second, when journals receive their revenue only from the author side, on a per-acceptance basis, there is a danger that a journal might lower its quality standard and accept more papers. It is possible that equilibrium quality might be lower in a market with many open access journals.¹⁶ Finally, it is important to recognize that journals have to compete for authors in a way they do not have to compete for libraries. Since authors can only publish a given article in a single journal, journals hold a monopoly over giving libraries access to this article and can therefore charge high prices to libraries.¹⁷ If authors care sufficiently about the reach of their paper, it is possible they may be willing to pay substantial fees in order for their article be freely available to readers, but it is not clear there is much author demand for this service.¹⁸ When journals compete harder for authors than for libraries, the market equilibrium may be one with no author charges and costs being covered (or more than covered) by the library side of the market.

The alternative considered in the remainder of this paper is the widespread use of

¹⁵A journal with operating costs of, say, £200,000 per year which publishes 50 articles per year would need to levy a publication fee of £4,000 per article. (These figures are similar to those for the *Review of Economic Studies* at the time of writing, for instance.)

¹⁶See McCabe and Snyder (2005) and Jeon and Rochet (2007) for further discussion of this point.

¹⁷See Armstrong (2006) for more details on this “competitive bottleneck” story, and McCabe and Snyder (2006) for the application of related ideas to the journal market.

¹⁸For instance, many publishers offer authors the option for their article to be open access in return for a substantial fee, but apparently few authors take advantage of this.

collection sales but without using historical data on library subscriptions. As is clear from the discussion so far, once bundling is involved, it is hard to establish (i) a journal's contribution to the overall demand for the collection and (ii) how to determine the price for the collection as a whole. Instead of using data on past stand-alone subscriptions, publishers could compete to offer an independent journal an annual fixed payment for the right to market the journal as the publisher sees fit. Such an arrangement bypasses the revenue allocation issue altogether, and it gives journals no explicit control over the total charge for the bundle. Implicit in this method is a promise from the publisher to attain a specified reach for its bundle. While for-profit journals do not care about reach, a non-profit journal will need to consider both its remuneration and reach when choosing its publisher. Two kinds of publisher exist in this market: commercial publishers which aim to serve the for-profit journals and less commercial publishers (perhaps some university presses, or some non-profit entity set up for the purpose) aiming to distribute the non-profit journals. Commercial publishers would set a profit-maximizing bundle price to libraries, and then pass this back to their journals according to their value in the bundle. Less commercial publishers would charge a lower bundle price which just allows their journals to break even but which expands their reach. I present a stylized model of such a market in the next section.

3 A Model of Collection Sales

In this section I provide a model of collection sales.¹⁹ The model presumes that bundling is the norm, that historic records of which libraries previously held full subscriptions are irrelevant, that publishers pay lump-sum fees to journals for the right to market their journals, and that publishers offer libraries a bundling tariff for the right to access all or part of their collections of journals.

Suppose the various journals are indexed by $i = 1, 2, \dots, N$, and that a particular library is willing to pay v_i to obtain access to journal i . Journals are assumed not to substitute for each other, and a library's willingness-to-pay for one journal is unaffected

¹⁹Related analysis is found in Bergstrom and Bergstrom (2004). Their interest is in the situation where a single journal is sold to many readers (via a site license), unlike my focus which is on selling many journals to a single "reader" (the library). However, in economic terms the analysis of the two situations is almost identical. In addition, their focus is more on the impact of bundling arrangements on the welfare of readers, while here only the interests of journals are discussed.

by which other journals it buys. Relatedly, suppose that libraries are not budget constrained, and can afford to buy any journal for a price up to the valuation v_i .²⁰ In sum, a library’s total willingness-to-pay for the journal collection $\mathcal{I} \subset \{1, \dots, N\}$ is $\sum_{i \in \mathcal{I}} v_i$. We assume that publishers cannot observe a particular library’s list of valuations (v_1, \dots, v_N) but they know the distribution of (v_1, \dots, v_N) in the population of libraries. Also, suppose the marginal cost of supplying a journal to an additional library is zero, and the costs in the publishing layer of the market are normalized to zero.

In this model two kinds of publisher will emerge in order to cater to the different aims of the two kinds of journals: commercial publishers will cater to (or own) for-profit journals, and concentrate their efforts on extracting as much revenue as possible from libraries; “non-profit” publishers will attract the non-profit journals by setting a low subscription price for their collections (but enough to cover the journals’ costs) and thereby enhance their reach.²¹

The analysis is divided into two parts. First, I focus on the case of small collections, where a publisher has only a few journals to distribute. This analysis apparently requires the choice of specific distributions for valuations to make much progress. Second, I examine cases where a publisher has many journals to distribute, which facilitates a more general analysis.

3.1 Small Collections

For the sake of simplicity consider the case where each v_i is independently and uniformly distributed on the interval $[0, 1]$. We discuss the cases of stand-alone journal pricing and bundling in turn.

Stand-alone pricing: Consider the case of a for-profit journal. Without bundling, the optimal price for such a journal maximizes $p(1 - p)$, i.e., $p = \frac{1}{2}$ is the subscription charge for a stand-alone journal. The profit to be shared between the journal and the publisher is $\frac{1}{4}$, and the journal sells to 50% of libraries. If the publisher market is competitive, all

²⁰Clearly this is a significant simplification. If a library faces an exogenous and *ad hoc* budget constraint, journals will have to compete with each other to be purchased, even if they are not viewed as substitutes by end readers. See Nevo, Rubinfeld, and McCabe (2005) and Jeon and Menicucci (2006) for analysis of this point.

²¹Note that we do not necessarily consider the two kinds of publisher to have different intrinsic objectives. All publishers could be profit-maximizers, and in a competitive publishing market some publishers succeed by offering non-profit journals the contracts those journals prefer.

that profit goes to the journal and the outcome is just as if the journal sold directly to libraries (e.g., via its own website), so long as the journal was as talented as a publisher at marketing activities. For a non-profit journal, the non-bundled price which just covers its fixed cost will obviously depend on that fixed cost: for example, if the fixed cost is $k = \frac{1}{10}$, the lowest break-even price is $p \approx 0.11$, and so this non-profit journal reaches 89% of libraries.

Bundling two journals: Turning next to bundling arrangements, suppose first that a publisher distributes just two journals. The publisher charges libraries a price p_1 for subscribing to either single journal on its own, while p_2 is its price for subscribing to the second journal if the library also subscribes to one journal (so that the price for subscribing to the collection of both journals is $p_1 + p_2$). Then, under the realistic assumption that there is a discount for consuming the bundle (i.e., $p_2 \leq p_1$), the pattern of library demand for the publisher's journals is as depicted on Figure 1.

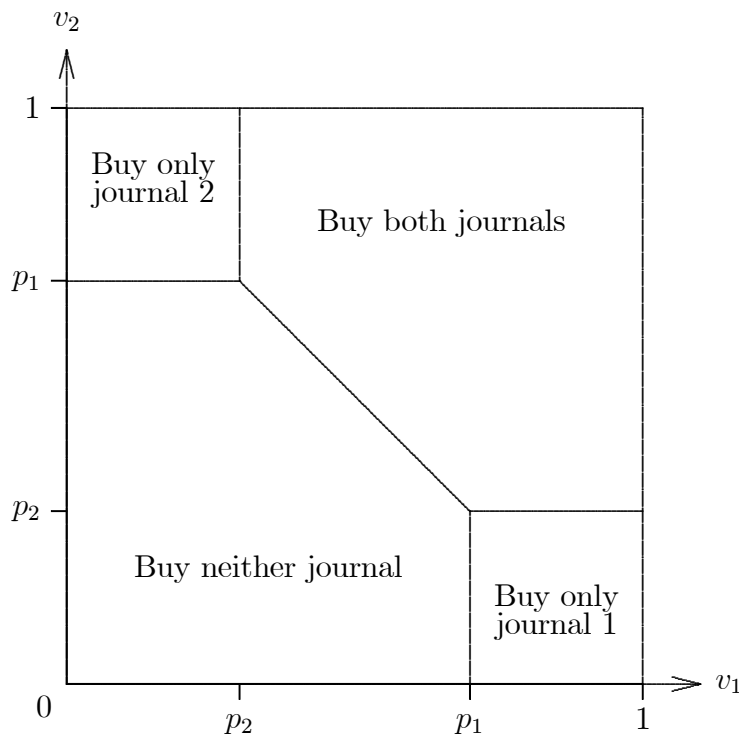


Figure 1: Pattern of Library Demand with Bundling (Two Journals)

When these journals are for-profit, the publisher will aim to maximize the profit from marketing these journals. By calculating the areas of the various regions on Figure 1,

one can show that the most profitable way to sell these journals to libraries is mixed bundling of the form:

$$\text{price for any single journal} = \frac{2}{3} ; \text{ price for both journals} = \frac{4}{3} - \frac{\sqrt{2}}{3} \approx 0.86 . \quad (1)$$

Thus the price for a single journal rises relative to the situation with stand-alone pricing (where the price was $\frac{1}{2}$), while the charge for the bundle falls. Bundling generates total profit of approximately 0.55, so that the incremental profit of the “second” journal is 0.3, instead of 0.25 with stand-alone pricing. One can show that each journal now sells to about 60% of libraries (compared to 50% with stand-alone pricing), and so bundling expands the reach of these journals (even though these for-profit journals do not care about their reach). By examining Figure 1, one can see that the fraction of libraries who buy nothing from the publisher with the tariff (1) is precisely one-third.²²

Suppose instead that the journals are non-profit, so that the publisher chooses its pricing policy to maximize the journals’ reach, subject to allowing the journals to cover their fixed cost, say k for each journal. In this example at least, it turns out that *pure* bundling is optimal, and the journals are not made available individually (i.e., $p_2 = 0$). Again, the price for the bundle depends on the size of k . For instance, if $k = \frac{1}{10}$, the bundle price for the two journals is approximately 0.2 and the collection’s reach is about 98%. Thus, bundling again acts to boost the reach of a journal. (With stand-alone pricing and $k = \frac{1}{10}$, a non-profit journal reached 89% of libraries.) Of course, it is generally true that bundling can expand the reach of a non-profit journal, since it adds to the range of pricing instruments which could be employed: the publisher could sell its journals without bundling, but except in knife-edge cases it can do strictly better with a degree of bundling. (The only case in which bundling cannot do better is when a library’s willingness-to-pay for one journal is *perfectly* correlated with that for the other.)

Bundling three journals: One can perform the same exercise when a publisher distributes three journals, by brute-force calculation of the volumes of the various regions in the three-dimensional version of Figure 1, which is depicted as Figure 2. Here, p_1 is the price for subscribing to the first journal, p_2 is the price for the second, and p_3 is the price for

²²In this example, the aggregate surplus of libraries increases slightly when bundling is employed, although those libraries which have little interest in one of the journals are worse off since the stand-alone price rises.

the third. (Thus, the total charge for getting two journals is $p_1 + p_2$, while $p_1 + p_2 + p_3$ is the charge for the entire collection.) The diagram illustrates the case where $p_1 \geq p_2 \geq p_3$. The region in which valuations (v_1, v_2, v_3) make it optimal to buy the set of S journals is labeled as S on the diagram. (Thus, those libraries who buy the entire collection lie in the region marked $\{1, 2, 3\}$.) Not all eight possible demand regions can be shown on a single diagram.




Figure 2: Pattern of Library Demand with Bundling (Three Journals)

A profit-maximizing publisher with three journals will set the following bundling tariff:

$$\text{price for any single journal} = \frac{3}{4}; \text{ price for all three journals} \approx 1.218. \quad (2)$$

Note that there is no separate price for subscribing to two journals (i.e., $p_3 = 0$), and libraries are simply offered the choice between subscribing to a stand-alone journal or getting the publisher's entire list. In this example, the reach of each journal is now about 70%, while the incremental profit of the third journal over what the publisher could obtain with two journals (which was about 0.55) is about 0.32. In this example,

the fraction of libraries who buy nothing from the publisher with the tariff (2), i.e., the volume of region $\{\emptyset\}$ on the figure, is precisely 25%.²³

3.2 Large Collections

For this author at least, calculating the precisely optimal mixed bundling tariff for the Uniform example if the publisher has four or more journals seems close to intractable.²⁴ Moreover, even with a collection consisting of just a few journals it is clear that the form of the optimal tariff will depend heavily on the precise distributions of valuations. The reader may rightly be worried about the strong assumptions: (i) that all journals are symmetric; (ii) that all libraries have their valuations taken from the same distribution, let alone (iii) that valuations are Uniformly distributed.

It is technically much easier, and in many respects more general, to suppose a publisher has “many” journals, and thereby invoke the powerful *law of large numbers*. Such an assumption is also not too unrealistic.²⁵ This discussion is divided into two parts:

²³One can show for general n journals that with a Uniform distribution for valuations that the profit-maximizing bundling tariff will induce the fraction $1/(n + 1)$ of libraries to buy nothing (as verified in our analysis for 1, 2 and 3 journals). These examples with the uniform distribution and 1, 2 and 3 journals makes it tempting to extrapolate, and to hypothesize that with n journals the price for a stand-alone journal is $n/(n + 1)$. Moreover, it would be interesting and worthwhile to analyze in more detail when the optimal strategy for a profit-maximizing publisher is to set a price for a stand-alone journal and a price for the whole bundle, and not to offer any intermediate bundles.

²⁴An example which is much more tractable involves each library having a *binary* distribution for its willingness-to-pay for a journal. For instance, consider the parametrization where the low valuation is 1 and the high valuation is 2, and where each library has an independent probability of $\frac{7}{10}$ of having a high valuation for any journal. With just one journal (or, more generally, with stand-alone pricing for journals), a profit-maximizing publisher will sell only to high-valuation libraries and set the stand-alone price 2. With several journals, the publisher’s decision reduces to deciding how many full-price journals (i.e., with price 2) the library must get before it can obtain all remaining journals at the discounted price 1 per journal. One can show the following is the optimal bundling tariff when the publisher has n journals to distribute:

number of journals in collection:	1	2	3	4	5
number of “full-price” subscriptions needed for discount:	1	2	2	3	3
reach of each journal:	70%	70%	85%	80%	90%
incremental profit of journal:	1.40	1.40	1.50	1.47	1.51

Note that with two journals, no bundling is used. For more journals, bundling is used to some extent, and the reach of a journal and the incremental profit of a journal are higher than the corresponding stand-alone values. (However, reach and incremental profit are not monotonic in the number of journals, and “integer constraints” for the number of necessary full-price journals appear to play a major role in this.) For further analysis of the binary case, but for welfare-maximizing rather than profit-maximizing tariffs, see Fang and Norman (2008).

²⁵It is important to remember that it is a publisher’s entire collection of journals which is usually relevant for an institution’s decision to subscribe, not, say, just the number or quality of its economics

the case where all libraries have a similar valuation for collections of many journals, and the more plausible case where libraries differ systematically in their journal demand.

Similar journals: Suppose that each library has a valuation for journal i , v_i , taken from some distribution for v_i with mean \bar{v}_i . Suppose that a library's valuation for journal i is independently distributed from its valuation for journal j .²⁶ (This assumption will be relaxed in the next section.) When a collection \mathcal{I} consists of many journals, a library's actual willingness-to-pay for \mathcal{I} is tightly distributed around its expected willingness-to-pay, so that most libraries satisfy

$$\sum_{i \in \mathcal{I}} v_i \approx \sum_{i \in \mathcal{I}} \bar{v}_i.$$

For simplicity, suppose each library's willingness-to-pay is exactly $\sum_{i \in \mathcal{I}} \bar{v}_i$ for the bundle of journals \mathcal{I} . Thus, all libraries have the same willingness-to-pay for a large collection, and a library's private information is eliminated through the use of collection sales.

Consider first a publisher which distributes a collection \mathcal{I} of for-profit journals. If the collection is large, the publisher should therefore set the pure bundle price:

$$\text{price for collection } \mathcal{I} \text{ of journals} \approx \sum_{i \in \mathcal{I}} \bar{v}_i. \quad (3)$$

(Note that no journals are made available individually, and the publisher retails its entire list on a take-it-or-leave-it basis.) All libraries will just be willing to pay this, with the result that each journal is available in all libraries (in contrast to the non-bundling situation, where a for-profit journal was taken by just half the libraries in the Uniform distribution example). In a competitive publisher market each journal will receive its contribution to the collection revenue, i.e., journal i will receive the incremental profit \bar{v}_i . This incremental profit is equal to aggregate willingness-to-pay of all libraries, i.e., journal i will be paid the entire area under its demand curve to be included in the collection, so that it is just as if the journal (or publisher) could engage in first-degree price discrimination between libraries. Commercial publishers here extract the entire surplus of libraries. (This contrasts with the previous two-journal example where bundling made the libraries better off.)

journals. At the time of writing, for instance, Wiley-Blackwell distributes more than 1400 academic journals and Elsevier distributes more than 2300.

²⁶See Armstrong (1999) and Bakos and Brynjolfsson (1999) for details of this analysis.

Table 1 summarizes the profit-maximizing outcome for the Uniform example, where $\bar{v}_i = \frac{1}{2}$, for various numbers of journals:

Table 1: Profit-maximizing collection sales in Uniform example

number of journals in collection:	1	2	3	...	∞
reach of each journal:	50%	60%	70%	...	100%
incremental profit of a journal to the collection:	0.25	0.30	0.325	...	0.50
average price of journal if entire collection purchased:	0.5	0.43	0.41	...	0.5

This example serves to demonstrate the very considerable potential benefits of collection sales in terms of enhancing a for-profit journal’s (i) reach (which *doubles* in this example relative to stand-alone pricing) and (ii) revenue (which also doubles).

Consider next the non-profit journals. As long as the publisher sets its collection price below that in (3), each library is willing to take the collection. In particular, a non-profit journal benefits from enhanced reach when bundling is employed. (In this example with $k = \frac{1}{10}$, when bundling is used a non-profit journal’s reach expands from 89% to 100% coverage, although almost all of this circulation gain is obtained when just two journals are bundled.) Thus, regardless of whether it is for-profit or non-profit, a journal benefits from bundling. An artificial feature of this model is that library demand for a collection of n journals is essentially inelastic up to the choke price (3). As such, non-profit (and for-profit) journals do not face a meaningful trade-off between the revenue they receive and their reach. Therefore, a non-profit journal is not harmed if it is distributed by a commercial publisher, since it obtains a 100% reach in any event. A more sensible model where there is a real trade-off between revenue and reach is presented in the next section.

The following table presents results for the reach-maximizing bundling tariff in the Uniform example for the case where $k = \frac{1}{5}$ for a variety of collection sizes.²⁷ (This simulation assumes that pure bundling is the optimal strategy, as was found to be the case in the two-journal example earlier. If it turns out that mixed bundling does strictly better, the results reported below are a lower bound on what non-profit journals could

²⁷These simulations, and those reported in Table 3, need to use the distribution for the sum of n Uniform variables, the cumulative distribution function for which is

$$F_n(x) = \sum_{k=0}^n (-1)^k \frac{([x-k]_+)^n}{k!(n-k)!}.$$

Here, $[x-k]_+ = x-k$ if $x \geq k$ and zero otherwise.

achieve with bundling. As mentioned, in the limit with many journals, any average journal price between 0.2 and 0.5 would still achieve 100% reach, and the table reports the lowest price consistent with journals breaking even.)

Table 2: Reach-maximizing collection sales in Uniform example ($k = 0.2$)

number of journals in collection:	1	2	3	4	5	...	∞
reach of each journal:	72%	90%	96%	98%	99%	...	100%
average price of journal in collection:	0.28	0.22.	0.21	0.20	0.20	...	0.2

Thus, again it appears that the benefits of bundling are quite quickly exhausted, and there is little benefit for a non-profit journal being part of a very large collection.

Table 1 shows that there is a tendency for incremental profit to increase as the size of the collection grows.²⁸ Thus, a publisher with a greater number of other journals will in many cases be prepared to pay more to attract an additional journal than a publisher with fewer existing journals. In this model, the publishing market exhibits a natural tendency towards a concentrated market structure, with large publishers better able to attract further journals than smaller publishers.²⁹ A similar point applies to competition for non-profit journals: a publisher with a greater number of other journals will be able to offer greater reach to a journal than a publisher with fewer existing journals (as illustrated in Table 2 above).

Table 3: Reach-maximizing collection sales in Uniform example ($k = 0.3$)

number of journals in collection:	1	2	3	4	5	...	∞
reach of each journal:	n/a	n/a	n/a	80%	90%	...	100%
average price of journal in collection:	n/a	n/a	n/a	0.38	0.33	...	0.3

It is possible that a “niche” journal (either for-profit or non-profit) with low overall demand from libraries would not be able to cover its fixed cost in a world without bundling, and so would not be supplied. But the extra revenue generated with bundling may allow cost-recovery. This is illustrated in Table 3, which is the same example as for

²⁸The argument shows that a publisher with many journals is willing to pay more for an additional journal than a publisher with few journals. But we cannot be sure that the incremental profit of a further journal is monotonic in the size of the collection. Indeed, the table in footnote 24 illustrates that this is not always the case.

²⁹This point is made in section 3 of Bakos and Brynjolfsson (2000).

Table 2 but with journal costs increased to $k = 0.3$. In this example, a journal cannot break even unless it is bundled with at least three other journals. Thus, it is possible that the use of bundling leads to the introduction of new, smaller, journals. (This effect is quite distinct from the possibility, discussed in section 2, that with bundling publishers may introduce spurious low-value journals to boost their share of the top-up fee.) This observation is an instance of the general point that price discrimination can act to open markets which would not otherwise be served.

Heterogeneous libraries: In the previous discussion, a publisher which bundles together many journals in its collection does not face a trade-off between its reach and the price of its collection. Since this seems unrealistic, consider the following extended framework.³⁰

Suppose that libraries differ systematically in their willingness-to-pay for collections of journals, captured by the scalar parameter θ . (This parameter might reflect the size of the institution, its degree of research-orientation, or its wealth.) A library's willingness-to-pay for journal i is $v_i = \theta w_i$, and if the library subscribes to journal collection \mathcal{I} , its total willingness-to-pay for the bundle is $\theta \sum_{i \in \mathcal{I}} w_i$, so that θ shifts a library's willingness-to-pay multiplicatively. Even if w_i and w_j are independently distributed, v_i and v_j are now positively correlated via the common factor θ . Suppose that θ is distributed throughout the population of libraries with the distribution function $G(\theta)$, and that the $N+1$ parameters θ, w_1, \dots, w_N are independently distributed. Suppose that each library's preference parameter w_i for journal i is drawn from some distribution with mean \bar{w}_i .

As before, assume that a publisher distributes many journals, so that a library's actual willingness-to-pay for collection \mathcal{I} , given θ , is approximately equal to the expected willingness-to-pay:

$$\theta \sum_{i \in \mathcal{I}} w_i \approx \theta \sum_{i \in \mathcal{I}} \bar{w}_i .$$

Because the parameter θ remains private information to a library, publishers can no longer accurately predict a library's willingness-to-pay for even a large collection. In particular, a commercial publisher cannot fully extract library surplus.

What is the optimal strategy for a commercial publisher who distributes a large number of for-profit journals? If a publisher has the set of \mathcal{I} journals and chooses the collection price P , the marginal library type willing to pay P for the collection has θ

³⁰See section 3 of Armstrong (1999) for more details.

such that $\theta \sum_{i \in \mathcal{I}} \bar{w}_i = P$. Therefore, if it chooses to make θ the marginal library, the revenue generated from the collection is

$$(1 - G(\theta)) \times P = (1 - G(\theta)) \times \left(\theta \sum_{i \in \mathcal{I}} \bar{w}_i \right) .$$

Therefore, regardless of its list of journals, the publisher will chose θ^* to maximize $(1 - G(\theta))\theta$. Each of its journals will then (in a competitive publisher market) receive its incremental contribution to the bundle revenue, so that journal i will be paid $(1 - G(\theta^*))\theta^*\bar{w}_i$. Notice that a profit-maximizing journal does not care about the quality of the other journals in the same collection (as represented by the other journals' \bar{w}_j), since that does not affect its own remuneration.

Next, consider competition for non-profit journals. For these journals, a publisher will aim to maximize the reach of its collection (i.e., minimize the marginal library type θ), while covering its journals' fixed costs. It is clear that journals of (approximately) the same quality (as represented by \bar{w}), will group together with a publisher, and higher quality journals will have a greater reach.³¹ To see this, suppose that a publisher has 100 high-quality journals with $\bar{w} = 3$ and 100 lesser-quality journals with $\bar{w} = 2$, so that the value of this collection to a type- θ library is $500 \times \theta$. If each journal's fixed cost is k , then the publisher must ensure that

$$500(1 - G(\theta))\theta \geq 200k ,$$

and it will choose the smallest θ which satisfies this break even constraint. But if another publisher “poached” the 100 high-quality journals and offered libraries this high-quality bundle (worth $300 \times \theta$ to the type- θ library), its break even constraint is relaxed to

$$300(1 - G(\theta))\theta \geq 100k .$$

Thus this rival publisher can offer the high-quality journals an expanded reach, while still covering their costs, relative to the initial case. Thus, higher quality journals have lower per-journal prices and greater reach. Nevertheless, bundling leads to greater reach for a non-profit journal of any given quality than does stand-alone pricing.

This discussion presumes there are “many” non-profit journals of broadly similar qualities, which might seem implausible. However, there are two reasons to be reassured

³¹More generally, if cost differs across journals, journals with the same quality/cost ratio μ/k will be distributed together.

that this analysis is relevant. First, it is not necessary that within each subject area (such as economics) there be many journals of each quality level, only that across *all* subject areas there be enough similar journals. Second, for non-profit journals it appears that the benefits of bundling are quickly exhausted—being in a relatively small collection achieves most of the benefits of being in a large collection (see Table 2, for instance)—and so in practice it will probably not be necessary to have many similar journals in a collection. For instance, even bundling together, say, the “top five” non-profit journals in a single subject area would likely capture most of the available efficiency gains from bundling.

To illustrate some of this discussion, suppose each w_i is independently and Uniformly distributed on $[0, 1]$, so all journals are the same quality, and suppose that θ is uniformly distributed on $[0, 2]$. (This configuration ensures that each journal has expected value $\frac{1}{2}$, which makes for comparable discussion with the examples in section 3.1.) Suppose a journal’s fixed cost is $k = \frac{1}{10}$. When bundling is used the commercial publishers, catering to for-profit journals, will choose θ^* to maximize $\theta(1 - \frac{1}{2}\theta)$ and so set $\theta^* = 1$ and sell to 50% of libraries. Each for-profit journal will obtain revenue $\frac{1}{4}$ (well in excess of their cost). By contrast, in a stand-alone pricing regime, a for-profit journal in this environment would set the subscription price $p \approx 0.57$, which is attractive to just 36% of libraries. Therefore, each such journal will obtain revenue of around 0.2.³²

A publisher catering to non-profit journals will choose the marginal subscribing library to have the smallest parameter θ which satisfies the break even constraint

$$\frac{1}{2}\theta(1 - \frac{1}{2}\theta) = \frac{1}{10} ,$$

which entails selling the collection to about 89% of libraries. With stand-alone prices, a non-profit journal will set a subscription price of $p \approx 0.13$, which is attractive to around 75% of libraries. To summarise this example from the viewpoint of a non-profit journal:

- Without bundling, the journal can cover its costs by selling to 75% of libraries.
- With bundling, provided the journal uses a “non-commercial” publisher, it can cover its costs and increase its reach to 89% of libraries

³²A library’s willingness-to-pay for a given journal is θw , which is the independent product of a $U[0, 1]$ and a $U[0, 2]$ variable. The probability that θw is greater than p , where $0 \leq p \leq 2$, is therefore $\frac{1}{2}(2 - p + p \ln \frac{p}{2})$. The profit-maximizing stand-alone price maximizes $p[\frac{1}{2}(2 - p + p \ln \frac{p}{2})]$, which entails $p \approx 0.57$.

- However, if the journal is included in the collection of a commercial publisher which normally deals with for-profit journals, it will be available in just 50% of libraries. In particular, if a non-profit journal feels itself to be locked into dealing with a particular commercial publisher³³, it may be better off withdrawing from the publisher’s collection sales programme, and opt instead to be distributed on a stand-alone basis at its own chosen subscription price (in which case it can reach 75% of the libraries). In general, however, the comparison of reach between (i) non-bundling sales and (ii) bundling with a commercial publisher is ambiguous and depends on the details of the environment.³⁴

In sum, the best outcome for a non-profit journal is to be part of a collection sales programme, but with a non-commercial publisher which typically distributes other non-profit journals of similar quality and which sets a price for the collection which just allows its journals to meet their costs.

4 Conclusions

In the past, journals were sold to libraries on a stand-alone basis, and an independent journal was free to set its own subscription price. There are two main advantages to this policy (which could still be adopted by a journal now). First, the journal has full control over its revenues and reach. Second, there are no difficult issues to do with how to allocate the publisher’s revenues—the journal’s contribution to overall revenue is transparent. In such a world, a journal is (all else equal) indifferent to how its publisher otherwise deals with libraries. Similarly, a journal does not care (all else equal) about the quality of other journals distributed by its publisher, since a library makes its purchase decision regardless of which other journals it buys.³⁵

The electronic revolution means that it is essentially costless for a publisher to give one more library access to one of its journals. As a result, is it socially efficient for all

³³For instance, the embarrassment of terminating an existing contract may be a source of inertia. (Those involved in managing non-profit journals are seldom hardened business people.)

³⁴For instance, suppose the fixed cost k was much higher, so that only a price close to the profit-maximizing non-bundling price ($p \approx 0.56$) could allow the journal to break-even when the journal was sold on a stand-alone basis. Then the journal’s reach is improved (from 36% of libraries to 50%) if it is distributed as a bundled package by a commercial publisher.

³⁵Similarly, a journal which participates in a publisher’s collection sales programme does not care about the number or quality of the publisher’s journals which have chosen to remain outside that publisher’s collection.

libraries to have electronic access to all journals. Bundling acts to improve the trade-off between revenue and reach for a journal, so that for a given level of target revenue a journal can expand its reach if it is bundled with other journals. Bundling has the potential to enhance profit (and reach) of for-profit journals and to expand the reach of non-profit journals.

Current arrangements are an awkward mishmash of stand-alone and bundling elements. The price paid by a library for a collection depends on what journals it historically took, and the extra journals are offered to the library in return for a top-up fee in addition to the stand-alone subscription prices. The division of this top-up fee between the participating journals is, almost inevitably, *ad hoc* and perceived as unfair by many journals. Of particular concern is the potential for a commercial publisher to cross-subsidise its own in-house journals from the revenue generated by the independent non-profit (and often more prestigious) journals. In addition, the reliance on historic data about library subscriptions will become ever less tenable. This arrangement is unstable and will likely change to something more sensible. In the meantime, one option for a non-profit journal is to leave its publisher's collection, and to return to the historic business model of stand-alone sales. However, this decision will deny the journal the circulation benefits involved in collection sales.

This paper focussed on one (I hope, plausible) model for the evolution of the publishing market. This model involved: the widespread bundling of journals; publishers bidding to attract journals using lump-sum fees and promised library circulation; non-profit journals are offered relatively little revenue in return for a wide circulation, while for-profit journals are offered the reverse. The efficiency gains from the use of bundling will flow into higher profits for for-profit journals (often together with higher reach relative to a stand-alone business model), and into higher reach for non-profit journals. The model predicts that non-profit journals of roughly similar quality will group together in the same publisher.

In the historic regime of stand-alone journal sales there was little tension when a non-profit journal was published by a highly commercial publisher. Now, though, there is a tension, and non-profit journals might benefit from disentangling themselves from the more commercial publishers. Both for-profit and non-profit journals, however, should surely make full use of the powerful instrument of bundling. In particular, relative

to any stand-alone sales strategy, a non-profit journal will be better off if it joins the collection sales programme of a noncommercial publisher. The example of *JSTOR*—which distributes the content of a collection of largely non-profit journals with a lag of several years—might be one possible guide for how to disseminate also the current output of non-profit journals.³⁶

References

- ARMSTRONG, M. (1999): “Price Discrimination by a Many-Product Firm,” *Review of Economic Studies*, 66(1), 151–168.
- (2006): “Competition in Two-Sided Markets,” *Rand Journal of Economics*, 37(3), 668–691.
- BAKOS, Y., AND E. BRYNJOLFSSON (1999): “Bundling Information Goods: Pricing, Profits, and Efficiency,” *Management Science*, 45(12), 1613–1630.
- (2000): “Bundling and Competition on the Internet,” *Marketing Science*, 19(1), 63–82.
- BERGSTROM, C., AND T. BERGSTROM (2004): “The Costs and Benefits of Library Site Licenses to Academic Journals,” *Proceedings of the National Academy of Sciences*, 101(3), 897–907.
- BERGSTROM, T. (2001): “Free Labor for Costly Journals,” *Journal of Economic Perspectives*, 15(4), 183–198.
- DEWATRIPONT, M., V. GINSBURGH, P. LEGROS, AND A. WALCKIERS (2007): “Pricing of Scientific Journals and Market Power,” *Journal of the European Economic Association*, 5(2), 400–410.

³⁶ *JSTOR* (www.jstor.org) distributes a large number of journals from several disciplines, with a preponderance of non-profit and society journals. Articles are distributed with a lag so as not to unduly cannibalize a journal’s library subscriptions. *JSTOR* is available to libraries on a bundled basis (with scope for libraries to choose only particular subject areas). Since *JSTOR* has always distributed collections rather than individual journals, there is no issue about basing its library prices on historical individual subscriptions. *JSTOR* is a non-profit enterprise and sets relatively low prices to libraries, and pays relatively low fees to participating journals for distribution rights.

- EDLIN, A., AND D. RUBINFELD (2004): “Exclusion or Efficient Pricing? The ‘Big Deal’ Bundling of Academic Journals,” *Antitrust Law Journal*, 72(1), 119–158.
- FANG, H., AND P. NORMAN (2008): “Optimal Provision of Multiple Excludable Public Goods,” mimeo.
- GINSBURGH, V., AND I. ZANG (2007): “Bundling by Competitors and the Sharing of Profits,” *Economics Bulletin*, 12(16), 1–9.
- JEON, D.-S., AND D. MENICUCCI (2006): “Bundling Electronic Journals and Competition Among Publishers,” *Journal of the European Economic Association*, 4(5), 1038–1083.
- JEON, D.-S., AND J.-C. ROCHET (2007): “The Pricing of Academic Journals: A Two-Sided Market Perspective,” mimeo.
- MCCABE, M., AND C. SNYDER (2005): “Open Access and Academic Journal Quality,” *American Economic Review*, 95(2), 453–458.
- (2006): “The Economics of Open-Access Journals,” mimeo.
- NEVO, A., D. RUBINFELD, AND M. MCCABE (2005): “Academic Journal Pricing and the Demand of Libraries,” *American Economic Review*, 95(2), 447–452.