

A PRACTICAL MODEL OF COPYRIGHT ECONOMICS WITH INTERMEDIARIES

PETER DICOLA

ABSTRACT. After streaming, what should copyright economics look like? The standard model of copyright economics used by courts and other policy makers – the “protect what would otherwise be a public good from copyists” model – is not general enough. It is particularly poorly suited for the era of internet streaming, because it assumes that the producers of creative works are vertically integrated with the ultimate retailers of those works. In this paper I argue for a new, rough-and-ready framework for use in formulating copyright policy. Rather than assuming a vertically integrated producer–retailer, my proposed model disaggregates the producer and the retailer, which I refer to more generally as an intermediary. One central feature of the streaming era is the growth and resulting bargaining power of technological intermediaries, including the large internet platforms and various popular streaming services. My proposed model of copyright economics allows for variation in the bargaining power of intermediaries, rather than implicitly assuming their power away. Making this shift in our mental model – our metaphor – for how copyright works changes the way we think about copyright law’s economic functions. It also suggests the need to refocus copyright policy on the important dynamics between producers and intermediaries.

1. INTRODUCTION

A key assumption in the standard model of copyright economics is that copyright owners are the relevant decisionmakers in the market for copies of their works. Copyright protection affords owners the opportunity to charge a price greater than the marginal cost of making and selling a copy. Without question, the standard model depicts the copyright owner as facing constraints in their choice. They face the costs of creating the work itself – the “first copy” – and the costs of producing subsequent copies. They also face a demand curve. That demand curve economists conceive has a lot of information embedded in it. It implicitly reflects not just consumer preferences in the abstract, but also consumer preferences given the existing, known set of business models for distributing and delivering copies of works for consumption. Somewhere in that demand curve, in other words, lurk the intermediaries that copyright owners almost always contract with to distribute, retail,

and otherwise deliver copies. In the operation of the standard model, however, the copyright owners are able to observe the demand curve (and all the information it contains) and then choose their optimal price.

But do copyright owners actually set prices? In many important contexts within many creative industries, the answer is no. Or let me put it this way: copyright owners do not set prices by themselves. Rather, an intermediary often sets the price or the price is the result of a complex negotiation between the copyright owner and an intermediary. For example, in a licensing negotiation almost two decades ago at the birth of the iTunes Store, Apple got its way and set the price of a downloaded song at \$0.99. The record labels wanted a higher price of \$1.99 – notice that initially there was no compromise at all, as though all the bargaining power in the negotiation sat with Apple and Steve Jobs (Knopper, 2009). Later, the iTunes Store moved to a three-tier approach of \$0.69, \$0.99, or \$1.29. Still the high-end price did not approach the record labels’ desired price, even for the latest songs from the most popular artists. Album prices similarly started at \$9.99 for a single album, more or less universally, but then came to vary somewhat more.

Of course, the record labels are not monolithic; different executives may have had different price preferences not reported in the journalistic accounts. Plenty of executives may have agreed with the business and marketing case for the \$0.99 price. Moreover, the record labels are not necessarily candid; one could imagine that the stories we have represent a disingenuous public-relations performance of wanting a higher price. Why would the record labels put on such a show? Possibly to placate recording artists, their managers, and their lawyers, who want the record label to maximize revenue (see Plant 1934). It might have been important to keep the peace in the industry for the record labels to appear to have been dragged kicking and screaming to Apple’s preferred \$0.99 price point. But even if we concede all these possibilities, the story of music-download pricing is not a story about copyright owners setting prices unilaterally. Intermediaries participated, perhaps decisively, in the price-setting.

There are other examples from the music industry. If you talked to a person on the copyright-owning side about their licensing deal with YouTube in the wake of the failed

secondary-liability litigation in the early 2010s, you would have heard a story about Google having them over a barrel. Even when Taylor Swift challenged Spotify she didn't challenge YouTube. No one can seriously maintain that even the sound recording copyright owners in the music industry have unfettered pricing power. Moreover, in the U.S., the musical work copyright owners are constrained by law from setting their own prices freely.

As a music-industry watcher I've known at least the outlines of the iTunes-pricing story for almost 20 years. But what I've come to wonder in recent years is this: why doesn't the standard model of copyright incorporate or at least accommodate that story? Why pretend that copyright owners have the power to set their price?

In many entertainment-industry contexts, not just in the music industry, copyright owners don't really set prices at all. They are constrained by the nature of the market they are selling to. For example, book publishers these days negotiate with Amazon and other electronic platforms when they determine retail prices for e-books (Dwyer 2014). For a while I thought of this as an internet phenomenon, but that isn't right, either. Copyright owners have long faced markets in which they cannot set prices. In the U.S., record labels dealt with Wal-Mart as a major retailer while book publishers dealt with Borders and Barnes & Noble. With this in mind, I want to propose that we replace or augment the standard model to incorporate licensing negotiations between copyright owners and intermediaries.

1.1. Aiming for a practical model. In light of these important facts, our most basic model of the economics of copyright should not have as its central dynamic the setting of price by a copyright owner. To be clear, I am talking here about the simple, core model of copyright economics that you might teach to an undergraduate course in applied economics. And that model assumes away the negotiation between copyright owners and intermediaries.

I am focused on that particular, broad-brush model because I believe it is the only economic model that most legislators, judges, regulators, lawyers, lobbyists, and copyright owners themselves have in their heads. The standard model is the metaphor that drives

copyright law and policy. What I'm saying in this paper is that the standard model of the economics of copyright is not as relevant as a model relied upon for policy ought to be, because one of its core assumptions takes the model fully away from the most relevant economic dynamic – the interaction between copyright owners and intermediaries to determine prices and even product offerings. I think the model that lawmakers use should instead center around this crucial interaction between copyright owners and intermediaries. That goal is what I mean to convey with the phrase “practical model” in the title of this article. I think we should aim to refine the rough, basic model that guides people with power when they make important decisions about copyright law and policy.

Now, in this paper, I am going to present economic models expressed in algebra, take derivatives, and provide a running numerical example that generates lots of not-round numbers. I build from the standard model and make one adjustment, augmentation, or relaxation of an implicit assumption at a time. But this can quickly generate complicated models. So I must admit that the economics in this paper is not going to jump off the page and serve practical purposes without additional work or explanation. I have tried to make the paper accessible, but this paper itself is not written for a general audience. Instead, I hope that this paper provides a foundation and enough working-through of the economics to support a new basic model that, after much development, conversation, and critique, could be used to explain, analyze, and perhaps even inform copyright policy.

A model that applies to all copyright industries requires abstraction and will not match all institutional details. For example, the music industry actually features not just one type of copyright owner but two intertwined types of copyright owners, music publishers and record labels. A more realistic (and much more complex) model would account for this three-way structure. But I have sought to keep the complexity of the models manageable. My main goal is to bring today's technological intermediaries fully into view in discussions of copyright economics.

1.2. A descriptive model, not a normative one. To allow for comparison among the standard model, which lacks intermediaries, and a progression of proposed models

that do feature intermediaries, I will propose measuring the copyright owner's profits as a percentage of the monopoly profits that a vertically integrated monopolist would earn. In this metric, the standard model sets the scale at 100 percent. In models with intermediaries, copyright owners achieve a lower percentage than that – as we will see, potentially much lower. Thus the central theme of the paper is that copyright owners reap fewer profits than the standard model would lead one to predict. This partly a matter of industrial organization economics and partly just a matter of accounting, since there are at least two entities instead of just one vertically integrated entity. Copyright scholars know this already; my point is that the standard model should be revised to reflect it.

An important clarification about my project is that I am offering a descriptive model, not a normative one. Pointing out that copyright owners lack unilateral pricing power is just a general, stylized statement of fact. I do not intend it in any way as a veiled complaint or as a tacit celebration. I am certainly not arguing that copyright owners should or should not have the power to choose the monopoly price and reap monopoly profits as a result. To my mind, there is nothing normatively correct about allocating monopoly profits to a copyright owner. My point in proposing a metric based on the percentage of monopoly profits earned by the copyright owner is solely descriptive; it allows for apples-to-apples comparisons between models. I think it is important and highly relevant for us to understand who ends up with what financial rewards. But whether the real-world allocation of profits among copyright owners and intermediaries is normatively good or should be altered is a complex question involving many different philosophical values. I think the descriptive model of this paper provides information that is highly relevant to that policy question, but I am not arguing for any particular policy or taking the side of either copyright owners or intermediaries.

I am, on the other hand, arguing for a new understanding of existing policy and what it can achieve. The models I propose here that go beyond the standard model demonstrate something important about copyright protection. In these models, the benefits that flow from copyright protection go to both copyright owners and intermediaries. As a result, policy proposals that strengthen copyright, depending on their design, might not benefit

copyright owners dollar for dollar, pound for pound, or Euro for Euro. Copyright owners only get a fraction of the benefits. Moreover, in most industries and under most contractual norms, creators only get a fraction of what the copyright owners collect. Thus, when intermediaries are taken into account, we have even less reason to think that “stronger ” copyright policy is putting much of the money in the pockets of creators.

That last point brings up a possible extension that is beyond the scope of this paper. I am arguing for a disaggregation of producers and intermediaries – to distinguish between copyright aggregators like record labels and book publishers and, on the other hand, the technology companies that deliver entertainment goods to consumers. But another key disaggregation could be between creators and copyright owners. There are, of course, interesting economics of how contracting between creators and copyright owners can vary both across and within industries (which is the subject of much excellent prior work in this journal and others). I don’t, however, discuss those economics here. I will refer to producers as a single entity, “copyright owners,” even though they are really aggregations themselves. Still, as you read the paper, you might keep in mind that as the fraction of profits going to the copyright owner diminishes, only a fraction of *those* profits are going to the human creators.

2. DISAGGREGATING PRODUCERS AND INTERMEDIARIES

In this section, I will first describe what I consider to be the standard model of copyright economics. I will do this as parsimoniously as possible, because I am sure that readers are familiar with it. I just want to highlight a few features and make a few critical comments along the way. Next, I will explain why we ought to alter the standard model – change its focus, augment what it is capable of addressing, and hopefully keep it close to earth. Finally, I will discuss the most fundamental adjustment necessary to update the standard model in a way that will be practically useful: reject the implicit assumption of vertical integration of creation, production, distribution, and delivery and replace it with a flexible approach to the intermediaries in each entertainment industry. In short, I argue that we ought to disaggregate producers and intermediaries.

2.1. The standard model. What I am calling the standard model of copyright economics is a story about a public good made private. Without protection for intangible goods, the story goes, an arbitrarily large number of copyists will enter the market and compete the price down to marginal cost. Or is it a sufficiently large number of copyists? Maybe just one copyist would have this effect? We don't seem to pay much attention to the specifics of copyists, just their destructive potential when it comes to the price that the producer of an intangible good would otherwise like to charge. Copyright law thwarts all the copyists (or is it merely a sufficient number of them?) and leaves the pricing to the producer of the intangible good, now known to us as the copyright owner. On the margin, many producers will now earn economic profit from sales of copies, or at least break even, when considering both the variable costs of distributing copies and the first-copy costs.¹ The prospect of doing epsilon better than breaking even is enough, according to the standard incentive theory, for the producer to choose to generate the creative work.

Notice that the public-good story, or more descriptively the convert-public-to-private-good story, is a price theory story. The creator-producer needs the power to set the price some distance above marginal cost. Each one wants to more than cover the first-copy costs. For the largest possible number of creator-producers to achieve this (and thus create their works and release them for sale to the public), the price each should choose is the monopoly price.

The standard model has many variations. I will put it into algebraic terms in what I hope is the simplest, most comprehensible, and most adaptable way. Let P be the price of a copy, and let Q be the quantity of copies sold. The total revenue derived from a

¹The first-copy costs are also "sunk" once spent, but the model as we're contemplating it here is from the perspective of the creator-producer not yet having spent them and in fact deciding whether to do so. I suppose the first-copy costs could also be called "fixed," and this is helpful to distinguish them from costs that vary with the number of copies, but inaccurate since there are so many decisions for a creator to make that influence the cost of making a creative work. Then again, I suppose one could say that any variation in the creative choices that produce a work render the work a different work entirely with a different demand function, etc., and so in that sense the first-copy cost of a specific creative work is fixed because only one set of choices would have produced that exact work in all its attributes and features.

copyrighted work, over the entire copyright term of the work, is then $P \cdot Q$.² Let F denote the first-copy costs while $C(Q)$ is the cost of goods sold, so that $\frac{C(Q)}{Q}$ is the average cost of a copy and $C'(Q)$ denotes the marginal cost of a copy. Then the accounting identity we have is

$$\Pi_{CO} = P \cdot Q - C(Q) - F \quad (1)$$

$$\pi_{CO} = P \cdot Q - C(Q) \quad (2)$$

where uppercase Π_{CO} is total economic profit and lowercase π_{CO} is the copyright owner's operating profit, i.e. the profit before taking account of first-copy costs. I will most frequently discuss the latter concept of operating profit. Because this is a model with certainty, we can treat P as $P(Q)$, the inverse demand function, and think of the copyright owner as choosing Q to maximize π_{CO} . That optimization produces the following condition that implicitly defines the optimal price P^* :

$$P^* = C'(Q) - P'(Q) \cdot Q \quad (3)$$

The optimal price is greater than marginal cost because $P'(Q)$ is negative. This suggests that the other helpful and familiar way to think of the monopolist's optimal solution is to express price as a markup over marginal cost:

$$\frac{P^* - C'(Q)}{P^*} = \frac{1}{\epsilon} \quad (4)$$

where ϵ is the elasticity of demand, $-\frac{dQ}{dP} \cdot \frac{P}{Q}$. Since $\epsilon > 0$ by this definition, the monopolist's chosen price is indeed a markup over marginal cost. The size of the markup increases as the elasticity of demand decreases.

The real thrust of the standard model is that $\pi_{CO}(P^*)$ is greater than zero. This allows the copyright owner to make the comparison between operating profits π_{CO} and first-copy costs F . The copyright owner looks ahead with rational expectations and the benefit of perfect information. If it sees that the game is worth the candle – that is,

²Prices and quantities could vary over time, of course, so it would be more general to write P_t and Q_t and then sum up from $t = 0$ to $t =$ (the number years until the author's death + 70 years) to reflect the copyright term. But the focus of the models in this paper won't be on changes over time, and so I have suppressed the time subscripts and assumed constant prices.

if it sees that $\pi_{CO} \geq F$ – then it chooses to produce the work. All the while in the background of this algebra of the standard model is the assumption that copyright law solves the public good problem, taking the copyright owner out of the price-taking situation of perfect competition and into the price-setting situation. The opportunity to choose P^* , the optimal price,³ begets the opportunity to evaluate whether the project is worthwhile.

Although the above equations hold for many different functional forms, it is common to think of the case with linear demand and constant marginal cost. For this paper, it will be helpful to have a specific running example. Say that the demand function is $Q = 200 - 2P$ and thus the inverse demand function is $P = 100 - \frac{1}{2}Q$. Assume further that $C(Q) = 10Q$ so that $C'(Q) = 10$. The copyright owner chooses $P^* = 55$ and $Q^* = 90$, earning a profit $\Pi_{CO}^* = 4,050 - F$. As long as $F \leq 4,050$, the copyrighted work in this example will be created and distributed.

For purposes of comparing the standard model to the altered models I will propose later in the paper, it is worth noting that the ratio of profits achieved to the maximum profits possible is 100% and that the share of those profits garnered by the copyright owner is also 100%.

2.2. The need to alter the standard model. The standard model, of course, features a bunch of assumptions that don't always hold. First, the creator-producer must be choosing the price, but in many entertainment-good markets it cannot. Second, the creative work must have no substitutes, but any entertainment good has at least some infinitesimal substitutability with any other entertainment good. So many times the choice of a price is either not within the creator-producer's power or would never reach the monopoly price due to the presence of substitute goods.

Let me explain in broad strokes what I'm thinking of. There are three market structures or scenarios I have in mind. First, there is the powerful-platform model, illustrated by my iTunes story. Second, there is the dispersed, non-integrated distributor model, illustrated

³Technically, of course, we had the copyright owner choose Q^* and thus P^* , which we could do because this is a model with certainty and the relationship between price and quantity in consumer demand is a deterministic, two-way relationship. In other words, working with an inverse demand function doesn't change the fact that the copyright owner is choosing a price along with a quantity.

by retailers of hard copies of books, video games, movies, TV, or music. Third, there is the single-client model, illustrated by commercial photographers. In the first two structures, intermediaries

In the first sort of structure, an intermediary platform has most of the power to set the price. The monopoly or oligopoly in distribution and delivery, not the one in production of creative works, is calling the shots. The rise of “The Four”⁴ has rendered standard copyright economics into a toy model from a bygone era. The reason is that the sales of creative works is but a small fraction of a distributor’s revenue. The distributor-platform business, and here I’m thinking of Amazon most keenly, is in the business of selling paper towels and cat food and lawn chairs. Music, TV, movies, and books are loss leaders to maintain customer loyalty and to harvest customer data. Not only are copyright owners not setting the prices – the platform isn’t even setting prices with regard to the creative industry in and of itself. It’s a component part of a much larger business. We should have noticed when, say, Wal-Mart became the largest retailer of music in the late 1990s. That was probably going to change the economics of copyright!

Even in the second structure or scenario, in which a dispersed set of competitive retailers handle delivery to consumers, copyright owners have relatively little power to set prices. That’s because they seldom operate in a situation of monopoly, but in a situation of monopolistic competition. Back in the Copyright Wars of the 2000s, responding to the critiques of copyright as a government-imposed monopoly with all the deadweight loss that implies, some defenders of the copyright system rightly pointed out that markets for creative works were not actually monopolies.⁵ Instead, the markets looked like monopolistic competition out of a Hotelling (1929) locations-on-a-line-segment model or a Salop (1979) locations-on-a-circle model. The market for recorded music, for instance, featured a small range of standard prices – say, \$9.99, \$12.99, \$15.99, and \$17.99 at a record store like Coconuts – into which every album from *Sgt. Pepper* to that of a one-hit wonder padded with 11 weak songs was priced at one of those four levels. That’s not monopoly pricing.

⁴“The four” are Amazon, Apple, Facebook/Meta, and Google/Alphabet (Galloway 2017). Particular industries have their own hugely important platforms, e.g. Netflix in video or Spotify in music.

⁵See Michael Abramowicz (2004) and Christopher Yoo (2004).

The core reason for this is well-known; there's some degree of substitution, however, weak, between *Sgt. Pepper* and the one-hit-wonder's terrible record, and the same is true for any pair of albums, any pair of books, or any pair of movies. That substitution effect, that non-zero cross-price elasticity, constrains prices.

Even with a large range of distributors – say, an ecosystem of large retail chains mixed with independent retailers – the price of a book, DVD, or compact disc had a limited range with discrete choices within that range, usually corresponding to reasonably round numbers. Pricing is done by humans, for humans. Even with a rough behavioral-economics mindset, we can say that consumers expect a comprehensible price in most markets. This constrains prices as well. For these reasons, monopolistic competition is often the better type of model to have in mind for mass-market sales of copies of creative works. It always was, even before the internet.

The third market structure or scenario I have in mind is the one-buyer, one-seller sort of transaction. Copyright owners could set prices for this kind of market. The photographer may have offered a price based on a standard menu. There might be an hourly rate, a fixed fee, or a combination. But oftentimes the prices will be subject to a negotiation rather than being set for a mass market. These are individual, private transactions, in which the terms may never become public, so that others cannot use information about the terms. Consider the situation in which a corporate client hires a photographer to take photographs that the client will use in its annual report brochure. Moreover, the terms of the interaction will be subject to a negotiation, including whether there will be a transfer of copyright or, if not, what the scope and duration of the client's license to use the photographs will be. The standard price-theory model is not a good metaphor for this situation. The photographer as integrated producer-distributor of copyrighted works here is negotiating the price with one consumer. Initially, of course, the parties had other choices about who to work with, and throughout their negotiation they maintain outside options. But again the better metaphor for this situation isn't mass-market pricing but bilateral negotiation. Other licensing situations – movie rights for books, sample licenses for music, etc. – have important features in common with my example.

Certainly other structures and scenarios are possible; I don't mean for this roughly described taxonomy to cover the universe of possibilities. Rather, I want to argue merely that the three structures I've laid out here are widespread enough to be important. And then I want to emphasize that the copyright owner is either not choosing prices at all, choosing prices under very tight constraints, or negotiating prices rather than setting them for a mass market. Those two propositions are enough to necessitate a relaxation of the assumption that copyright owners set prices.

2.3. The fundamental change: model intermediaries explicitly. The important alteration in the standard model that is called for is to stop assuming a vertically integrated creator-producer-distributor-retailer. I advocate distinguishing between two kinds of agents: copyright owners and technological intermediaries, or just intermediaries for short.

Copyright owners could be the creators themselves, or among the creators, but the copyright owners may also have acquired the copyright by contract. The modifications to the standard model I discuss here will not focus on such acquisition and licensing contracts between, say, recording artists and record labels or authors and book publishers. The focus instead is on the copyright owner's interaction with the technological or media intermediaries who will distribute the work to the public for consumption.

If we think of the copyright owner as moving first, since it must create or acquire the copyrighted work before it can be delivered to the public, then the intermediary moves second. But to specify the model it makes sense to work backwards from the intermediary's decision about how to deliver the copyrighted work to consumers. In other words, just as in standard game theory, the copyright owner's interaction with the intermediary will be based on what both parties expect to happen in the later interaction between the intermediary and the consuming public. Given that, the intermediary's interaction with consumers vis à vis a particular copyrighted work will look similar to that of the vertically integrated copyright owner in the standard model.

One difference is that the intermediary will pay a licensing fee, denoted by $L(Q)$. There are many potential relationships between copyright owners and intermediaries and many forms that the contracts between them can take. This licensing fee could be a flat fee, or buyout, in which case it would be a constant function and would not actually vary with Q . It could be a royalty per copy sold by the intermediary, and thus a linear function of Q . It could be a percentage of the intermediary's revenue, $L(Q, P)$, but because consumer demand in this model is a deterministic function of price, we can think of P as inverse demand, treat it as $P(Q)$, showing that L can ultimately be expressed as a function of Q alone.

For simplicity of illustration in this section of the paper, I will begin the analysis by specifying that $L(Q)$ is linear, with licensing fees being proportional to the quantity of copies sold and no lump-sum fee, such that $L(Q) = \ell Q$. In later subsections, I will relax that assumption and consider more flexible contractual forms.

On the cost side, for simplicity the technological intermediary incurs no variable costs. Instead, the intermediary incurs only the sunk costs of building the infrastructure for distributing copyrighted works, notated as I . The copyright owner still incurs variable costs $C(Q)$; think of these as the costs of distribution and delivery through a customer-facing interface. As in the standard model, the copyright owner faces the first-copy costs F , which are treated as sunk at the time of the pricing decisions this model is concerned with.

The intermediary's total profit from sale of a particular copyrighted work is then given by:

$$\Pi_{TI} = P(Q) \cdot Q - L(Q) - I \quad (5)$$

The intermediary's operating profit, which is what I will work with more often (but which, one should remember, must exceed sunk costs I), is given by:

$$\pi_{TI} = P(Q) \cdot Q - L(Q) \quad (6)$$

And finally the copyright owner's operating profit from the sale of a particular copyrighted work is now given by:

$$\pi_{CO} = L(Q) - C(Q) \quad (7)$$

which will, as before, have to exceed first-copy costs F in order for the copyright owner to produce the work. Notice that despite the subscripts on the two agents' respective cost functions there is no subscript on Q or P . Both agents may incur costs that vary with the quantity sold. And, importantly, we should not think of price as being under either agent's sole control. The copyright owner is not making a unilateral decision about price, but neither is the intermediary, because that would simply be introducing the opposite problem to the fundamental problem that my model addresses. Thus, in my approach the price consumers face is the result of a negotiation between copyright owners and intermediaries. Possibly the copyright owner will have all the bargaining power, possibly the intermediary will, but more likely each will have some degree of bargaining power.

Now, what can happen in this model? What are the sources of conflict? In this basic setup, both the copyright owner and the intermediary are making decisions based on the potential to sell copies of a single copyrighted work. The copyright owner and the intermediary have some reason for their preferred prices to converge. They are splitting the pool of revenue from sales to consumers, and so they both have the incentive to generate a large pool of revenue. On the other hand, the parties are maximizing profits, not revenue. Each has its own cost function, which drives a wedge between the copyright owner's marginal cost curve and the intermediary's marginal cost curve. Generally speaking, either agent's marginal cost curve could be greater at any given quantity Q . And of course there is direct conflict over the licensing fee; the intermediary prefers L to be as small as possible while the copyright owner prefers L to be as large as possible.

We can see the divergence in interests in a different way by analyzing the optimization equations. Set $L(Q) = \ell Q$, so that $L'(Q) = \ell$, and $L''(Q) = 0$. The intermediary takes ℓ as given, and chooses Q^* to maximize $P(Q)Q - \ell Q$. The first-order condition for the optimum, $Q^*(\ell)$, is

$$P'(Q^*)Q^* + P(Q^*) = \ell \quad (8)$$

Write this as $P'(Q^*)Q^* + P(Q^*) - \ell = 0$, and apply the implicit function theorem to the left-hand-side to obtain:

$$\frac{dQ^*}{d\ell} = -\frac{-1}{P''(Q^*)Q^* + 2P'(Q^*)} = \frac{1}{P''(Q^*)Q^* + 2P'(Q^*)} \quad (9)$$

Since the denominator is just the second-order condition for the optimal choice of Q , we have the very logical condition that $\frac{dQ^*}{d\ell} < 0$.

On the other hand, the copyright holder chooses ℓ^* to maximize $\ell Q - C(Q)$, but where it is recognized that $Q = Q^*(\ell)$. Therefore, the first-order condition for an optimal licensing parameter is

$$Q^* + \ell^* \frac{dQ^*}{d\ell} - C'(Q^*) \frac{dQ^*}{d\ell} = 0 \quad (10)$$

Write this as:

$$Q^* + \frac{dQ^*}{d\ell} (\ell^* - C'(Q^*)) = 0 \quad (11)$$

Since $\frac{dQ^*}{d\ell} < 0$, this requires $\ell^* > C'(Q^*)$.

Alongside these incentive constraints, to borrow the vocabulary of principal-agent models, we also have two participation constraints, one for each entity. The copyright owner requires $L(Q) - C(Q) > F$ and the intermediary requires $P \cdot Q - L(Q) > I$ in order to create and disseminate the copyrighted work, respectively.

Now, going back to the intermediary's FOC, we now see that the level of output is set such that

$$P'(Q^*)Q^* + P(Q^*) = \ell^* > C'(Q^*) \quad (12)$$

Comparing with the coordinated solution, where optimal output is Q^c such that $P'(Q^c)Q^c + P(Q^c) = C'(Q^c)$, we can easily compare the coordinated solution, Q^c , with the uncoordinated (double marginalization) solution, Q^* . Since the second-order condition of the intermediary is assumed to hold (*i.e.* its payoff is concave in Q), it happens that $P'(Q)Q + P(Q)$ is decreasing in Q . The uncoordinated solution sets $P'(Q)Q + P(Q)$ at a higher number than does the coordinated solution, so we immediately see that $Q^* < Q^c$. And since the demand curve is decreasing, it also happens that $P^* > P^c$.

This analysis demonstrates the divergence in interests for copyright owners on the one hand and intermediaries on the other. The copyright owner in this model is concerned

with obtaining marginal licensing revenue equal to marginal cost, while the intermediary is concerned with obtaining a consumer price above that.

2.3.1. *Double marginalization.* Modeling vertical integration or the lack thereof is familiar in industrial organization economics. The famous “double marginalization” result involves a supplier and a retailer who are arranged as sequential monopolists – a pair of bottlenecks. The firms cannot contract in any way other than an exchange of money for goods at a wholesale price. If the parties can’t coordinate on a different contract and do not vertically integrate, this dynamic produces a socially inefficient result and, more importantly, a worse result from the each of the individual perspectives of the supplier, the retailer, and the consuming public. The equations above explain why. Without some more sophisticated contractual mechanism, the supplier charges the monopoly wholesale price to the retailer, but the retailer takes the wholesale price, treats that as its marginal cost, and then turns around and charges a monopoly price to consumer from that perspective. The resulting price is even higher than the vertically integrated monopolist’s price would be, the quantity sold is even lower, and everyone loses out.

Going back to the numerical example from the end of Section 2.1, we can see how the arithmetic of double marginalization works out. Recall that consumer demand $Q = 200 - 2P$, inverse demand $P = 100 - \frac{1}{2}Q$, and the copyright owner’s costs are $C(Q) = 10Q$. By assumption, the form of the license is restricted to $L(Q) = l \cdot Q$ where l is a per-copy rate. The copyright owner bases its actions on what it knows with certainty the intermediary is going to do. The intermediary will set its marginal revenue, $\frac{d}{dQ}Q \cdot P = \frac{d}{dQ}[Q \cdot (100 - \frac{1}{2}Q)] = 100 - Q$ equal to the marginal cost it experiences, the per-copy licensing rate l . The copyright owner takes this demand function of the intermediary as given. The copyright owner will then choose l to maximize its profits, $\pi_{CO} = (l - C'(Q)) \cdot Q = (l - 10) \cdot (100 - l)$, which gives the first-order condition $\frac{d\pi_{CO}}{dl} = 100 - 2l + 10 = 0$, meaning the copyright owner chooses $l = 55$. That per-copy licensing rate means that the intermediary will set a consumer price of $P = 77.5$ and a quantity sold of $Q = 45$. The copyright owner makes a profit $\pi_{CO} = (55 - 10) \cdot 45 = 45^2 = 2,025$ and the intermediary makes a profit

$\pi_{TI} = (77.5 - 55) \cdot 45 = 22.5 \cdot 45 = 1,012.5$. The firms' aggregate profit is only 75% of the vertically integrated monopoly level, and the copyright owner garners only $66\frac{2}{3}\%$ of the aggregate profit. This leaves the copyright owner with only half the profit of the standard model described above in Section 2.1. Meanwhile, consumers are also worse off from the higher price and lower quantity.

Double marginalization is a troubling outcome and an analytically useful result. It helps us see how, in the disaggregated model of copyright economics I am presenting here, the copyright owner and the technological intermediary might make pricing decisions that harm each other and harm both of them in aggregate. But we should also recognize that this outcome doesn't match reality. We are not seeing prices for copyrighted works in the stratosphere; we are not seeing restricted quantities; and we are not seeing consumers losing out. This points us toward what does seem to be going on – and highlights the importance for policy makers of focusing on the bargaining that occurs between copyright owners and intermediaries.

2.3.2. More flexible contracts. In contrast to the conditions of the double marginalization result, copyright owners and technological intermediaries definitely engage in contractual negotiations, developing sophisticated legal instruments to govern the sale of copyrighted works (or subscription to catalogs containing them, as I'll discuss in Section 3.2). (Touve and DiCola, 2014). In some cases, as in the television and film industry, copyright owners have achieved vertical integration by developing their own direct-to-consumer distribution service, e.g., Disney Plus. In other cases, copyright owners have taken an equity stake in a technological intermediary, e.g., the major record labels' stake in Spotify, although I do think even that deal stops well short of vertical integration. In any event the economic dynamic to consider is whether integration has been achieved, if so to what degree, and if not what bargain the copyright owner and the intermediary have struck.

To keep things general and flexible, I would simply suggest thinking of the copyright owners and the intermediaries as splitting the surplus. Economists have long shown that a two-part tariff structure for the license can achieve this, where the form of our licensing

function $L(Q)$ is $L(Q) = \lambda + C'(Q) \cdot Q$, where λ is a flat fee and the royalty rate is then set equal to marginal cost. (This type of approach is also known as the “franchise fee” approach.) The copyright owner would charge as high a flat fee as it can get, but regardless of the level of λ , the intermediary now has the incentive to set the retail price for consumers at the vertically integrated monopolist’s price. Together, the copyright owner and the firm are maximizing joint profits. The flat fee λ determines how much of the surplus each party gets.

Another solution to the double marginalization problem is for the copyright owner to set a maximum retail price in a practice known as resale price maintenance. Given that copyright owners and technological intermediaries are known to bargain over (and disagree over) the consumer retail price, it seems plausible that some entertainment industries have solved the problem in this way. The contract might take the form of a maximum retail price coupled with either a discounted wholesale price charged by the copyright owner or, equivalently, a royalty paid to the copyright owner on each sale.

Yet another solution is known as nonlinear pricing, under which the licensing rate itself varies with the quantity sold. Designed optimally, this can also induce the intermediary to choose the same retail price as a vertically integrated monopolist.

It is easy to apply any of these schemes to the running numerical example, but for concision we’ll consider resale price maintenance. Suppose the copyright owner and the intermediary will agree for the retail price to be $P = 55$, same as in the vertically integrated case. The aggregate profit to be split is then $\pi_{CO} + \pi_{TI} = 4,050$. The breakdown between the two parties’ profits will depend on their bargaining power and the bargaining process. In this simple model, the Nash bargaining solution would be to split the additional surplus generated by solving the double marginalization problem 50–50.⁶ Each party’s outside option, we assume, is simply the double-marginalization result. It is as though the parties are contemplating a renegotiation of the simple licensing contract but are free to walk away and leave the simple licensing contract $L(Q) = l \cdot Q$ in place.

⁶For a recent discussion of the application of bargaining models, including the Nash bargaining model, to the music industry’s royalty rates, see U.S. Copyright Royalty Judges, 2022.

To calculate the Nash outcome, recall that the outcomes under double marginalization were $\pi_{CO} = 2,025$ and $\pi_{TI} = 1,012.5$. Aggregate profit had fallen 1,012.5 short of the vertically integrated monopolist level. Thus, with resale price maintenance, we might expect a bargained result for the copyright owner of $\pi_{CO} = 2,025 + \frac{1012.5}{2} = 2,531.25$, an improvement over double-marginalization problem. The intermediary correspondingly would receive $\pi_{TI} = 1,012.5 + \frac{1,012.5}{2} = 1,518.75$, also an improvement. This time, aggregate profit is 100% of the vertically integrated monopoly solution while the copyright owner's share is down to 62.5%. But of course different bargaining results, different shares, and different profit levels for the copyright owner are entirely possible. As long as the copyright owner achieved a share of $(50 + \varepsilon)\%$, it would be worthwhile to the copyright owner to avoid double marginalization.

No matter which of the solutions the copyright owner adopts to address the problem posed by the lack of vertical integration, the key takeaways are that the copyright owner and the intermediary will bargain with each other and that they will divide the surplus. Notice that even if the two firms jointly maximize profits through one of these ideal contracts, the copyright owner is actually only receiving its bargained-for share of monopoly profits – not 100 percent of those profits, as the standard model implicitly assumes. This emphasizes how important it is to consider two distinct margins once we depart from the standard model: (1) the ratio of aggregate profits to vertically-integrated-monopoly profits and (2) the share of profits garnered by the copyright owner.

2.4. Implications of disaggregating copyright owners and intermediaries. Adding an intermediary agent to the standard model of copyright, even in this simple setup with one copyright owner and one intermediary, lays the foundation for considering other market structures. But even in the simplest form of this model, as described here, our attention is usefully redirected. Instead of thinking about a copyright owner choosing a price, we must instead concentrate on a potentially multi-dimensional, multi-stage bargain between the copyright owner and the intermediary. The two agents may contract over the consumer-facing price P as well as the functional form and parameters of the licensing fee $L(P)$.

All I want to establish at this stage is that this bargaining process and its outcome are what should be of highest interest. There are many possibilities for how the bargaining game between the copyright owner and intermediary may proceed, and there could be large, consequential differences in what different specifications predict. For example, the parties could have incomplete information about each other's cost functions; for another, they could have different expectations about consumer demand. My purpose here is neither to present just one of those possibilities nor to explore all of them. None of that is necessary to make the point to my intended audience for this model: legislators, judges, lawyers, lobbyists, et al. I would rather stick to the simplest possible microeconomics to convey that policy makers should think about copyright owners and intermediaries as separate entities – and indeed as opposed entities, whose interests diverge for quite general and fundamental economic reasons, even before one puts a specific structure on the various functions (demand, cost, and licensing).

3. THE NATURE OF THE NEW INTERMEDIARIES

The first major step in departing from the standard model of copyright economics was to disaggregate the copyright owner and the intermediary. The next major step, described in this section, is to try to capture two important features of the new intermediaries based on the technologies developed over the last few decades. The first feature is that many intermediaries are conglomerates, selling a vast variety of goods well beyond copyrighted works and entertainment products. The second feature is that the way intermediaries disseminate copyrighted works to the public has shifted in a way that complements the conglomerate nature of the intermediaries. The move to streaming, in short, enhances the advantage that technological intermediaries reap from selling both copyrighted works and a range of other consumer goods.

This section aims to sketch as simple a microeconomic model of these phenomena as possible, to add realism to our revised model of copyright economics without (hopefully) too much complexity. To make the model more digestible for the model's target audience, i.e. to maintain practicality, I will continue to track how outcomes relate to the outcome of

the (vertically integrated) monopolist in the standard model. Once again, a key takeaway will be that copyright owners are likely to achieve only a fraction of that outcome.

3.1. Conglomerates Instead of Single-Focus Intermediaries. Now let's expand the model by having the intermediary in more than one line of business. Let us recognize that intermediaries, especially today, are often conglomerates, especially the dominant firms like Alphabet/Google, Amazon, and Apple. The easiest example to keep in mind for this section of the paper, in my opinion, is Amazon. Its executives have been explicit about entering entertainment industries for the purpose of enhancing the firm's overall retail business, which now extends to just about any product one can imagine. My working assumption in this section is that it matters to Amazon's business strategy that it sells paper towels, lawn furniture, clothing, and toys as well as books, movies, TV shows, and music (see Owsinski 2016). Similarly, it matters for other conglomerates that they have other lines of business besides entertainment goods that are helped by their sale of entertainment goods. The basic model of copyright economics should reflect this.

The way in which an intermediary finds synergy in selling both entertainment goods and (many) other goods has multiple facets. The entertainment goods might attract consumers to the same retail platform. Having the opportunity to purchase entertainment goods through the intermediary's stores or online interface might help retain consumers. Some intermediaries will seek to implement at least mild incentives for lock-in through loyalty programs, saved preferences, or even just knowledge of how to navigate the catalog of goods. Intermediaries today, unlike the intermediaries of old, collect much more data about consumers, even offline but especially online. Any bit of learning about a consumer's habits and preferences might enhance the intermediary's ability to sell other goods to that consumer, with a wider range, more consistency, and a greater share of wallet.

All these facets of the relationship between the intermediaries of today and their consumers could be modeled individually in great detail. But this will get complicated quickly on a theoretical level, and could require proprietary data to test and calibrate. In this paper I am seeking a practical model to guide copyright practice and policymaking. So

my suggestion is to model the relationship between entertainment goods sold through a technological intermediary and all the other goods that intermediary sells as a complementarity between the goods from the intermediary's perspective. In particular, I would treat this as a one-way complementarity: an increased quantity of copyrighted works sold has a positive effect on the number of other goods sold. With just that assumption, we will be able to see the possibility of downward pressure on the price of entertainment goods and the profits that copyright owners earn.

The copyright owner's optimization is the same as above in Section (2.3). The copyright owner sets $L'(Q) = C'(Q)$, equating its marginal licensing revenue with its marginal cost, i.e. the marginal post-creation costs. As before, both parties will have a participation constraint that their operating profits must cover their fixed costs (F for the copyright owner's first-copy costs, I for the intermediary's infrastructure costs).

To model this conglomerate effect as simply as possible, I suggest that one think of a composite good, Z , that doesn't have any connection to entertainment goods other than being sold by the same intermediary in the same location (whether that location is virtual, actual, or both). I will call the price of the composite good A so that the conglomerate's revenue from non-entertainment goods is $A \cdot Z$. (This choice of notation is my little device to keep reminding you of the Amazon example.) The cost function for other goods is notated as $\Gamma(Z)$. The intermediary's operating profits are now given by:

$$\pi_{TI} = P \cdot Q + A \cdot Z - L(Q) - \Gamma(Z) \quad (13)$$

where P , Q , and $L(Q)$ represent the same variables as before. Also, we continue to require $\pi_{TI} > I$ for the intermediary to operate. The copyright owner's problem remains the same as in Section 2.3. The key assumption is that A and Z , the (composite) price and quantity of other goods sold by the intermediary, are not independent of Q . Rather, both are functions of Q . In particular, to model complementarity, we assume that $A'(Q) > 0$ and $Z'(Q) > 0$. The quantity of entertainment goods sold shifts demand for the intermediary's other offerings higher. The resulting first order condition resolves to the following:

$$P^* = L'(Q^*) - P'(Q^*) \cdot Q^* - A'(Q^*) \cdot Z - Z'(Q^*) \cdot [A - \Gamma'(Z)] \quad (14)$$

The left-hand side along with the first and second terms of the right-hand side of this expression are the same as those that appear in equation (8) in Section 2.3 above. By assumption, the third term is negative, since $A'(Q) > 0$. The fourth term is also negative, with $Z'(Q) > 0$, as long as $A \geq \Gamma(Z)$, that is, as long as the intermediary is not selling other goods at a loss, in aggregate (recall that Z is a composite good). Knowing that the sign of those terms is negative tells us that being a conglomerate gives the intermediary incentive to set a lower price P for entertainment goods to achieve a greater quantity sold Q .

An expansion of the running numerical example can illustrate this. Keeping the same parameters as before for the market for the copyrighted work (i.e., continuing to assume $Q = 200 - 2P$ and $C'(Q) = 10$), the key new parameters are the size of the intermediary's market for other goods, the elasticity of demand in that market, and the magnitude of the cross-elasticity between Q and Z .⁷ Building on the previous section, I will assume for purposes of this section that the licensing contract specifies a maximum retail price and provides for the copyright owner to collect 62.5% of the profits from selling Q .⁸ Recall that we are also assuming that the parties do not bargain over the intermediary's revenue from the sales of Z . Part of the idea behind that assumption is that the intermediary is large and opaque enough to withhold any information about the rest of their business beyond retailing copyrighted works.

If we specify the demand function $Z = 1,000 - 10A + 0.1Q$ and a constant-marginal-cost function $\Gamma(Z) = 25Z$, then the intermediary no longer wants to choose the vertically-integrated-monopoly price and quantity (which were $P = 55$ and $Q = 90$) but instead chooses $P^* \approx 50$ and $Q^* \approx 100$.⁹ The intermediary is facing a tradeoff. It wants to

⁷Recall that for simplicity I am assuming that the complementarity is asymmetric; increased sales of Q increase sales of Z , but not vice versa. This really amounts to assuming that the effect of Q on Z is larger than the reverse, but it keeps the math simpler.

⁸Recall that this 62.5% figure is based on the Nash bargaining outcome. The disparity stems from the differing outcomes for the copyright owner and the intermediary under double marginalization, which I am treating as the parties' outside option.

⁹The reason is that, under the set of assumptions described in the text, $\pi_{TI} = \frac{3}{8}(P \cdot Q - C(Q)) + A \cdot Z - \Gamma(Z) - I$. To get Z and A in terms of Q , one can use the first-order condition with respect to Z , which is $\frac{\partial \pi_{TI}}{\partial Z} = A'(Z) \cdot Z + A - \Gamma'(Z) = 0$. Substitute in $A'(Z) = -0.1$, $\Gamma'(Z) = 25$, and $A = 100 - 0.1Z + 0.01Q$ (the latter of which is obtained by solving the specification for Z for A). This produces $Z = 375 + 0.05Q$ and $A = 62.5 + 0.005Q$. Moving to the other first-order condition with respect to Q , we have $\frac{\partial \pi_{TI}}{\partial Q} = \frac{3}{8}(P + P'(Q) \cdot Q - C'(Q)) + A'(Q) \cdot Z + [Z'(Q) \cdot (A - \Gamma'(Z))] = 0$.

increase the quantity Q because that is acting as a demand shifter, increasing both the price A and the quantity Z of other goods sold on the intermediary's platform or in the intermediary's stores. But it doesn't want to harm profits from Q , i.e. sales of the copyrighted work, too much, because it's still receiving 37.5% of the profits from such. As a result, the intermediary chooses a price-quantity combination that is different, but not drastically different. With these figures, the total operating profit from sales of the copyrighted work declines to about 4,000. With the $\frac{5}{8}/\frac{3}{8}$ split the copyright owner gets 2,500 (instead of 2,531.25) and the intermediary gets 1,500 (instead of 1,518.75). There is downward price pressure on the copyrighted work but it is certainly not being treated as a loss leader in this particular example. In fact, under these particular assumptions the intermediary is effectively insisting on making a significant operating profit from sales of the copyrighted work, too.

Of course, this numerical example depends heavily on the specific parameters, especially the degree of complementarity between Q and Z . It also depends on the assumptions that the licensing bargain takes a specific form and that it is no different in light of the conglomerate nature of the intermediary. I held these constant for the numerical example to allow an apples-to-apples comparison. But the more complicated nature of the intermediary's business could certainly alter the nature of the licensing contract between the two parties. For example, rather than a maximum retail price, the copyright owner might seek to secure a specific price (i.e., the price a vertically integrated monopolist would charge). In return, the intermediary might garner a greater share of the profits from sales of the copyrighted work. So this example might be subject to renegotiation. But either way, the copyright owner will fare worse than the case in which the intermediary was not a conglomerate selling a complementary composite good.

3.2. Modeling Intermediaries' Consumer Products and Services. One complication in adapting the standard model to account for intermediaries, especially in the

Substitute in the values of $A'(Q)$, Z , $Z'(Q)$, and A that we just obtained, as well as the value of $\Gamma(Z)$, to obtain: $\frac{3}{8}(100 - \frac{1}{2}Q) + \frac{3}{8}(-\frac{1}{2} \cdot Q) - \frac{3}{8} \cdot 10 + 0.005(375 + 0.05Q) + 0.05(62.5 + 0.005Q - 25) = 0$. Collecting terms, we have $(37.5 - 3.75 + 1.875 + 3.125 - 1.25) - (0.1875 + 0.1875 - 0.00025 - 0.00025)Q = 0$, or $37.5 = 0.3745Q$, meaning $Q \cong 100.133$ and $P \cong 49.933$, to be more exact.

streaming area, is that the sellers of entertainment products are not selling copies anymore but services that allow access to a large library of works on demand. See Towse (2020) for a discussion of the economics of streaming.

3.2.1. *Entertainment delivered in a new form.* The standard model is about selling copies: a one-time purchase that permits a consumer to enjoy a copyrighted work whenever they like. The paradigmatic case involves a physical copy that embodies the work: a codex, a plastic disc, a tape, a cartridge, or a print. Such physical copies would themselves be owned as personal property. The shift to digital changed this. Even though digital copies have a physicality to them, stored as magnetic ones and zeroes, the typical legal regime for digital copies changed from ownership to licensing. See, for example, software licenses or iTunes Store purchases.

Now, the standard model was always kind of eliding or papering over a distinct business model that some copyright owners have employed for over a century. Selling copies means exploiting the reproduction and distribution rights of copyright. But licensing to movie theaters or radio stations, for example, means exploiting the public performance right instead. Here, the copyright owner is using an intermediary, but the consumers don't walk away from their interaction with the intermediary with a copy in hand. The intermediary's business model is based on ticket sales, advertising, or something else besides retail of copies. The standard model isn't really telling that story. It's possible but awkward to think of a movie theater chain as akin to a DVD purchaser. And, in keeping with my theme in this paper, the standard model suppresses the role of the intermediary. So even just to keep up with the full range of 20th Century business models, we might want to augment or alter the standard model.

To try to keep up with 21st Century business models, specifically streaming on demand, we need to adjust the standard model. Streaming is a stronger substitute for selling copies than movie theater going or radio listening ever was. It involves subscriptions, which satellite radio also does, but the scale of the market for video and music subscriptions vastly outpaces anything previous to the last decade-plus of streaming. (Video games and

books have had some success with a subscription model, but not as much. Visual art really isn't the same.)

So what should it look like, to model the copyright owner's participation in the streaming economy? Do we need to make additional alterations to accommodate what a streaming intermediary does that differs economically from what a retail intermediary does? A model that attempted to capture every nuance of the current suite of business models that entertainment companies employ would be unwieldy. But I think it is possible and advisable to make a few parsimonious adjustments that will set us up to better understand the current reality of entertainment-industry intermediaries.

3.2.2. Counting plays instead of copies. As with any business model, there are lots of choices about how to capture subscriptions. Here, I've tried to make choices that maintain as many parallels with the sale of copies as possible, only making enough changes to capture the essence of subscription economics while maintaining primary focus on what I am arguing is the central dynamic of copyright economics: copyright owners negotiating with intermediaries. To capture the shift from selling copies to selling subscriptions that provide access to a library of works, of which the copyright owner's work is but one, the existence of other copyright owners and the existence of multiple copyrighted works looms larger than that fact does with the models discussed up to this point. The model nonetheless still focuses on the decisions of one intermediary and one copyright owner.

Moving to a subscription-based business model means that there is not a simple retail price for a good that consumers either buy once or don't buy during the copyright term (which, as a reminder, is the relevant time period I have been using throughout). Now the price consumers pay is about gaining and maintaining access to the copyrighted work in question (and, implicitly, gaining and maintaining other works that have been licensed to the intermediary's subscription service).

The intermediary's operating profit function is given by:

$$\pi_{TI} = P \cdot S + A \cdot Z - L(S, N) - \Gamma(Z) \tag{15}$$

To notate the shift we are contemplating here, instead of Q for quantity of copies sold, we have S for the quantity of subscriptions sold. More precisely, S is the number of subscription-months sold over the copyright term. (I am referring to months just because monthly subscription payments happen to be most common; subscription-years would also be a fine way to think of it.) If we think of S as containing that durational component, then P now has the intuitive interpretation as the monthly subscription fee. Just as the quantity of copies sold was a deterministic function of price in all the variations of the model above, S is a deterministic function of P here.

From Section 3.1 we retain the idea that the intermediary, as a conglomerate, experiences a complementarity between its product related to providing access to copyrighted works and all its other products. Specifically, we have the price A and the quantity Z of the composite good both increasing in subscriptions S . This section builds on the altered model by assuming that the complementarity between subscriptions and all other goods is even stronger than the complementarity between copies and all other goods. The model here provides one way to capture the important business development that streaming models convince consumers to engage extensively, often daily, with intermediaries' apps, interfaces, and platforms. The opportunities to build customer loyalty and learn about consumer preferences are even more vast than before. To foreshadow one of the results of this section, that increase in complementarity between S and Z (compared to Q and Z) will provide incentive for the intermediary to charge an even lower subscription price.

The other adjustment to the intermediary's optimization problem is that the licensing fee is now a function of two variables: subscriptions S and another new variable N which represents the expected number of plays of or listens to the copyrighted work that the consumers actually engage in. (The *actual* number of plays can be thought of as stochastic, but I am defining N as the expectation over the distribution of actual plays achieved.) The copyright owner's licensing revenue increases in both S and N ; the former increases the notional revenue pool for copyright owners and the latter helps determine the particular copyright owner's share. It may help to think of the variable N as really N_i , indexing for a specific copyright owner, but for now we can suppress the subscript.

The copyright's owner's operating profit function is given by:

$$\pi_{CO} = L(S, N) - C(N) \tag{16}$$

In addition to the licensing fee L depending on two variables, S and N , just as in the intermediary's profit function, there is one other difference in the copyright owner's profit function compared to what came before in Section 3.1. The copyright owner doesn't have a direct role in selling subscriptions to the whole library or catalog of available works. So the copyright owner's promotion and marketing costs – that is, whatever costs come after creation of the work – are not a function of S but only a function of N . In other words, $C(N)$ represents the cost of achieving an expectation of N streams of a copyrighted work, after the work has been created by expending the sunk first-copy cost F .

These specifications of the profit functions of each party actually makes it really easy to see the divergence in interests between the copyright owner and the intermediary. Not only do they have different revenue and cost functions, but their respective cost functions depend on different variables. In fact, each party is really exercising choice over a different variable: the intermediary is choosing the number of subscriptions S (and therefore the subscription price P) and the copyright owner is choosing the number of plays or streams N . To be clear, both the intermediary and the copyright owner are also negotiating over the form of the licensing fee $L(S, N)$. But as before, I am thinking of the licensing negotiation as a prior stage before each party optimizes what it can.

For the copyright owner, things are similar to before, but the copyright owner can affect the expected number of plays N through its expenditures $C(N)$. When the copyright owner chooses a particular level of $C(N)$, say $C(N) = k$, we assume that k is associated with a particular level of N . Assuming that this cost function $C(N)$ is known, then, choosing a particular level of spending is equivalent to indirectly setting N .

$$\frac{\partial L(S, N)}{\partial N} = C'(N) \tag{17}$$

Once the structure of the licensing contract is set, the copyright owner is only choosing N (and doing so indirectly, as explained above). In other words, the number of subscriptions is modeled as somewhat out of its control after the licensing negotiation. But the copyright owner might have bargained for contractual restrictions on the intermediary's behavior. Depending on the specific form of the license, many interesting dynamics could play out.

Turning to the intermediary's problem, the intermediary is choosing S , the number of subscriptions, and thus also choosing P because in this model the intermediary is a monopolist. (To repeat, the intermediary is also making another choice by negotiating the form of the licensing fee $L(S, N)$, but that decision between the parties happens before optimization.) Optimization for the intermediary produces the following condition:

$$P^* = \frac{\partial L(S, N)}{\partial S} - P'(S) \cdot S - A'(S) \cdot Z - Z'(S) \cdot [A - \Gamma'(Z)] \quad (18)$$

which mirrors the condition for the intermediary's preferred price fairly closely. The main difference from the perspective of the intermediary's optimization problem is that the marginal cost of licensing is now a partial derivative because L is a function of two variables. By assumption, the intermediary is not controlling the number of plays. Rather, it can only choose the partial effect of the number of subscriptions S on the licensing fee. Another way of putting this is that the intermediary has a large role in setting the subscription price, but in this model does not affect the number of plays. A more complicated and sophisticated model would incorporate on-platform advertising (e.g. banner ads, inclusion in playlists or viewing suggestions, or app notifications), but for my purposes here I have kept the model (relatively) simple.

For purposes of an illustration that comes as close to an apples-to-apples comparison as possible, I will try to adjust the running numerical example I have been using in the most minimal and straightforward way I can. Imagine first, counter to what I suspect and have assumed to be the case, that the level of complementarity between subscriptions S and the intermediary's composite good Z is the same as before, so that $Z = 1,000 - 10A + 0.1S$. This means the intermediary wants to set $S = 100$ and $P = 50$, analogous to its incentives in Section 3.1.

We can then choose $L(S, N)$, N , and $C(N)$ such that the subscription-based model of this section matches the joint profits produced by the sales-of-copies model of the previous section. Specifically, let $L(S, N) = (S \cdot P) \times 0.0007N$, i.e. subscription revenue times a small fraction times the number of streams. Then let $C(N) = 0.0025N^2 - 1.5N$, so that the copyright owner's variable costs are now quadratic and the marginal cost curve is no longer flat but slopes upward, and let $N = 1000$. Under these initial assumptions, subscription revenue is still 5,000, the copyright owner's variable costs are still 1,000, the intermediary's operating profits from retailing copyrighted works are still 1,500, and the copyright owner's operating profits are still 2,500 (i.e. 3,500 in licensing revenue minus the 1,000 in costs). The payoffs for each party are the same as in the sales-of-copies example, although the dynamics that produce them are different. Again, the goal here is to make an apples-to-apples comparison of what happens if the level of complementarity increases.

Toward that end, now suppose the level of complementarity between subscriptions S and the composite good Z is twice as high, so that $Z = 1,000 - 10A + 0.2S$. Under these parameters, the intermediary will choose $S \approx 111$ and $P \approx 44.5$. This would reduce subscription revenue to 4,939.5. That slight drop in the copyright owner's marginal revenue from licensing, which is proportional to subscription revenue, causes the copyright owner to choose a slightly lower $N \approx 991.5$ and earn a slightly lower operating profit of $\sim 2,457.8$. The intermediary's operating profit, just from subscriptions to stream the copyrighted work, increases to $\sim 1,511$, the reason being that subscription revenue drops less than the licensing revenue the copyright owner has forgone by spending less on marketing and promotion. Thus, the share of profits enjoyed by the copyright owner, under these assumptions and parameters, anyway, has declined slightly from 62.5% to 61.9%. Meanwhile, joint profit across both firms from selling access to the copyrighted work is down from 98.8% of the vertically-integrated-monopoly level (without subscriptions) to 98.0%.

If the level of complementarity between S and Z were even greater, say five times higher than previously, with $Z = 1,000 - 10A + 0.5S$, then the intermediary chooses $S \approx 145$

and $P \approx 27.5$. Subscription revenue drops to 3,987.5. The copyright owner chooses $N = 858.25$. The copyright owner's operating profit drops to ~ 1841.5 . The intermediary's operating profit from subscriptions alone increases to $\sim 1,592$. The copyright owner's share of profits has declined to 53.6%. And the joint profit from selling access to the copyrighted work has declined to just 84.8% of the vertically-integrated-monopoly level (without subscriptions).

The specifics of the example are artificial, to be sure, but this illustrates the two firms' incentives within the model I am proposing. As the complementarity between the technological intermediary's other products and subscriptions to access copyrighted works increases, the price of subscriptions drops, the copyright owner's share of profits decreases, the joint profit of the two firms from selling access to copyrighted works decreases, and the copyright owner's resulting operating profit has decreased as well.

4. COMPETITION ON ONE OR BOTH SIDES OF THE NEGOTIATION

So far, the augmentations and alterations of the standard model in this paper have involved a bilateral monopoly: one copyright owner negotiating with one intermediary. Moreover, the models are open to the interpretation, at least, that the copyright owner is licensing just one work. But, as mentioned in Section 3.2, once one starts thinking about a subscription model, there is pressure on that one-copyright-owner and one-work interpretation. Even if there is one copyright owner, that firm is licensing a catalog of works for which the firm has aggregated the copyrights. But of course it is also the case that there is more than one copyright aggregator in each of the major entertainment industries. Meanwhile, there is also competition among multiple technological intermediaries. Even the large internet platforms that have delved into one or more entertainment markets compete with each other for subscribers.

This section proceeds as follows. First, I will consider the situation in which there is still one copyright owner but there are multiple intermediaries. Then, I will consider the reverse situation of multiple copyright owners but one technological intermediary. Finally, I will offer some thoughts about the most complex situation with multiple intermediaries

and multiple copyright owners. Throughout this section, my approach will be to use the most basic tools possible to suggest the range of possible outcomes and how policy makers might develop a rough hypothesis about what is going on in a particular entertainment industry.

4.1. Multiple intermediaries, one copyright owner. The main effect of the copyright owner having multiple intermediaries to bargain with is that the copyright owner is going to do better. I will assume that the intermediaries cannot collude directly (which would risk violation of the antitrust laws). Their inability to coordinate is going to weaken their hand, not just individually but collectively. Exactly how well the copyright owner will do in the model depends, of course, on the specific assumptions one makes.

For example, suppose we go back to the setting of Section 2.3, except the one copyright owner is now upstream of two technological intermediaries. If the competition between the intermediaries is in the nature of Bertrand competition over price (rather than quantity), then just these two intermediaries will compete the surplus on their side of the negotiation down to nothing. In that case it is easy to see that the copyright owner will do well, achieving the vertically-integrated-monopoly result – and without going to the trouble of vertically integrating.

Perhaps a more interesting and seemingly realistic approach is to think of the now two intermediaries competing over quantity supplied to the market, i.e. Cournot competition. Each vies to deliver a certain number of copies to the market. Again, go back to the model of Section 2.3 about avoiding double marginalization. (Later in this section I will discuss the augmented model with conglomerates and with the shift to subscriptions and streaming.) Recall that the double-marginalization result depended on assuming that the contract between the copyright owner and the intermediary was limited to the copyright owner selling copies wholesale. No lump-sum fee or agreement to a maximum retail price was possible. Moreover, we are assuming that the copyright owner in this model is selling into a market, not engaging in bilateral negotiations with each intermediary; otherwise,

one intermediary might seek to pay a premium to obtain exclusive distribution rights to the copyrighted work.

Under those restrictions, with Cournot competition we have each of the two intermediaries best-responding to: (a) the other intermediary's quantity supplied and (b) the per-copy licensing rate l , which the copyright owner chooses before the intermediaries set their quantities; call them Q_1 and Q_2 . The market price for copies, P , is a function of aggregate supply $\bar{Q} = Q_1 + Q_2$, but each firm only controls its own quantity.

The intermediaries' operating profit functions are:

$$\pi_{TI1} = P(\bar{Q}) \cdot Q_1 - l \cdot Q_1 \quad (19)$$

$$\pi_{TI2} = P(\bar{Q}) \cdot Q_2 - l \cdot Q_2 \quad (20)$$

And the copyright owner's operating profit function is:

$$\pi_{CO} = l \cdot \bar{Q} - c \cdot \bar{Q} \quad (21)$$

which includes the assumption that the copyright owner's marginal cost is a constant c . Based on these assumptions, the first-order conditions for intermediary #1 and intermediary #2 become best-response functions, because each firm's first-order condition includes both its own quantity supplied and the other intermediary's quantity supplied. The fact that price P is a function of aggregate quantity $\bar{Q} = Q_1 + Q_2$ causes this.

Returning to the running numerical example, adjusted so that there are two intermediaries and inverse demand is given by $P = 100 - \frac{1}{2}(Q_1 + Q_2)$, we see a different result than before. The competition among the intermediaries means that the double marginalization problem is less acute. Rather than a single intermediary choosing $Q = 45$ and $P = 77.5$ as before, the two-intermediary result is that $Q_1 = Q_2 = 30$, so that $\bar{Q} = 60$ and $P = 70$. Each intermediary only makes a profit $\pi_{TI1} = \pi_{TI2} = 450$, while the copyright owner's profit is $\pi_{CO} = 2,700$. Collectively, the firms have achieved $88\frac{8}{9}\%$ of the

vertically-integrated-monopoly surplus (up from 75% when there was only one intermediary), and the copyright owner's share of that surplus is 75% (up from $66\frac{2}{3}\%$). From there, one can safely extrapolate this well-known result in industrial organization economics: as the number of intermediary firms increases, Cournot competition approaches the perfect competition/Bertrand result.

If the form of licensing contract is not restricted, the copyright owner can do even better. The monopoly price and quantity for copies of the copyrighted work can be achieved through, say, a lower per-copy licensing rate and a lump-sum "franchise fee." But the copyright owner will likely get an even larger share of this bargain to the firm's aggregate first-best solution than before, because its outside option is better and also better relative to how the two intermediaries fared. (The ultimate result depends on one's specific assumptions about the bargaining game played among the three firms, which is complicated by the fact that the two intermediaries may not coordinate.)

We can extend this Cournot framework to incorporate the other augmentations presented in Section 3 of this paper. For example, consider the adjustment that the intermediaries might each be conglomerates, each selling composite goods Z_1 and Z_2 , which composite goods we suppose are not in direct competition with one another. If, analogous to the model of Section 3.1, we have $Z_1 = 1000 - 10A_1 + 0.1Q_1$ and $Z_2 = 1000 - 10A_2 + 0.1Q_2$, then the intermediaries' chosen quantities will increase to $Q_1 = Q_2 \approx 32.5$, so that $\bar{Q} \approx 65$ and $P = 67.5$. The reason is that the complementarity of Q_1 and Q_2 with Z_1 and Z_2 , respectively, gives each intermediary incentive to increase their quantity supplied. In terms of profits from the sale of copies of the copyrighted work, each intermediary is doing worse than in the non-conglomerate version of this model. Their operating profits from sales of the work decrease from 450 to 406.25. But that loss is more than made up for by enhanced sales of the composite good.

This is an even better starting point for the copyright owner; it is closer to collectively achieving the vertically-integrated-monopoly price and quantity. Because we are comparing this to an equilibrium in which the aggregate supply was even lower than the monopoly level, the intermediaries' conglomerate nature pushes them in the right direction. Even if

a more flexible contract than just a per-copy licensing rate is not possible, the total surplus is 92.3% of the vertically-integrated-monopoly level and the copyright owner's share is 78.3%. Again, the main point of this section is just to confirm what one would expect: competition among intermediaries is generally advantageous for the copyright owner.

One could also extend the subscription model of Section 3.2 to a setting in which two or more intermediaries compete. I won't present a specific model here, because it is so sensitive to the specific functional forms and parameters assumed for the licensing contracts $L(S, N)$ and the copyright owner's marketing-and-promotion cost function $C(N)$. But I will offer a few observations. The nature of competition in the subscription world is more complex. The intermediaries are each seeking to be a non-exclusive licensee of the copyrighted work. One key issue in designing the model is whether the copyright owner has any influence on the subscription price and quantity, or whether the choice will be up to each intermediary. Once the model has this level of complexity, it is tempting to shift from a deterministic model – where the copyright owner's expenditures $C(N)$ guarantee a certain number of plays N , which then feeds directly into the copyright owner's licensing fee – to a probabilistic model in which the copyright owner takes a chance on advertising. My guess, for now, is that attempting to capture this reality renders the model impractical for my purposes here. In this paper, I think the main point is that competition among intermediaries is likely to enhance both the percentage of the vertically-integrated-monopoly surplus achieved and the copyright owner's share of that surplus.

4.2. Multiple copyright owners, one intermediary. The next task is to consider the reverse situation: multiple copyright owners negotiating with one intermediary. For some real-world motivation for this scenario, think of copyright owners of television, film, music, and video games negotiating with YouTube. Similar to the previous section, I will build first on the model from Section 2.3, changing it from a model of an upstream monopolist in copyrighted works and a downstream monopolist in distribution technology to a model with upstream duopolists negotiating with that downstream monopolist.

This task turns out to require more substantial adaptation of time-tested, off-the-shelf models than the models I have described in the previous sections. One possibility would be to adapt models of monopsony. But the vast majority of those models focus on the labor market. And, more limiting for my purposes here, they describe only the situation with an arbitrarily large number of sellers of labor and one buyer of labor. My interest here, as with much of the preceding analysis in this paper, has more to do with industrial organization economics rather than labor economics; that is, it has to do with several copyright owners negotiating with one large intermediary. There is some very interesting literature on buyer power in distribution or “countervailing power” (for surveys, see Inderst and Mazzarotto 2008; Snyder 2005). Experts in that topic seem to acknowledge, astutely, that there could be much more work done on this topic, that there are multiple dynamics and institutions that could come into play, and that as a result simple predictions about the resulting prices, quantities, contracts, and so on, are often not feasible.

One unsatisfactory (but slightly informative) way to capture my ideal scenario – two copyright owners, one intermediary – would be to assume that the copyrighted works in question were perfect substitutes. In that setting, there is a market price for copyrighted works. The two copyright owners are suppliers, but by assumption what they are supplying is fungible. In the absence of a way for the copyright owners to commit to restricting supply (e.g. through strategically self-imposed capacity constraints), they would face Bertrand competition due to the strong incentive to undercut each other, and they charge only marginal cost as their licensing fee. That latter scenario renders the copyright owners unable to cover their first-copy costs F_1 and F_2 . That might be happening in the real world, but one would only observe it as works not created.

Continue with the notion of perfect substitutes. If the copyright owners could somehow restrict supply or otherwise commit to not undercutting one another’s licensing fees, they could land at the Cournot outcome, keeping their licensing fees above their marginal cost. The monopolist intermediary would then charge a monopoly margin over those licensing fees. In short, this would be the reverse of the scenario described in Section 4.1. In the now-familiar running example, suppose we keep the demand for both (in this model)

fungible copyrighted works at the same level as before, so that the relevant, apples-to-apples comparison is how the two copyright owners fare collectively. The copyright owners would charge $l_1 = l_2 = 40$ (rather than the monopoly level of 55) and each earn operating profits of $\pi_{CO1} = \pi_{CO2} = 900$. Because of double marginalization, the intermediary will ultimately charge $P = 70$, with $Q = 60$, and earn profits of 1800 (having paid a licensing fee of 40 per copy). The three firms have together generated $88\frac{8}{9}\%$ of the vertically-integrated-monopoly outcome, but here the copyright owners collectively receive only a 50% share. Multiplying those two shares, the copyright owners earn $44\frac{4}{9}\%$ of the operating profits that a vertically integrated monopolist would. But this model, again, is assuming an extreme degree of substitutability between copyrighted works.

As I mentioned, it is unsatisfactory to assume that, despite the existence of two copyright owners producing two copyrighted works, the two works are perfect substitutes. To move beyond this, I will adapt the famous circular model of product differentiation from Salop (1979). Contrary to most uses of the circular model or Hotelling's (1929) linear model, I am not in this paper interested in entry and exit. Rather, I am interested in the short- and medium term dynamics of prices, quantities, and the resulting profits, since the latter is what matters for my particular focus within copyright policy. So instead of moving towards a model in which the copyright owners' location choices are endogenous (as in Abramowicz 2004, Yoo 2004), I will consider a model in which the locations are fixed and the spatial aspect is there to represent consumer substitution among copyrighted works and to allow a straightforward comparison with the preceding models in this paper.

Suppose, then, that consumers are arrayed uniformly on a circle to represent their ideal points, i.e. the place in "product space" that would provide them with the most utility. Demand for each copyrighted work is a function of the distance along the circumference between a given copyrighted work's location on the circle and the consumer's location. Each consumer buys only one copyrighted work. I will continue to assume that the copyright owner's marginal costs are constant. And for purposes of this analysis, the intermediary is not a conglomerate and is not selling subscriptions but rather retailing copies.

With this setup, one can port over almost all the parameters from the running numerical example of this paper. Let the circumference of the circle be 200, which is our measure of the maximum number of potential consumers. Previously we had inverse demand $P = 100 - \frac{1}{2}Q$. To have each differentiated work generate that same area under the demand curve, we now have inverse demand $P_{i,j} = 100 - |d_{i,j}|$ where $d_{i,j}$ represents the distance from copyrighted work i (here, $i = 1, 2$) to consumer j . (Note that, unlike the Cournot example in this section, the total surplus is now twice as large, since each copyrighted work has an equivalent demand curve to what has come before.) The marginal cost for each copyright owner is still a constant, $c_1 = c_2 = 10$. We can now consider different scenarios as to how closely the two copyright owners locate near each other, which represents the degree of substitution between them.

If the two copyright owners are located on opposite sides from each other – call them positions 0 and 100 around our circle of circumference 200 – then no consumers at all are diverted from one work to the other in equilibrium. As a monopolist would, our two monopolistically competitive copyright owners choose $l_1 = l_2 = 55$. The intermediary sets $P_1 = P_2 = 77.5$ in response, taking its own monopoly markup. Total operating profit and each copyright owner’s profit is the same as before. The intermediary garners one-third of the profit from the sales of each work, as in subsection 2.3.1 on double marginalization.

Next suppose the two copyright owners are located at 0 and 50 on the 200-circumference circle; they sit at 12 o’clock and 3 o’clock, so to speak. Now, there are some consumers on the circle who value both copyrighted works more than marginal cost and must choose between them. In other words, there could be meaningful overlap in demand for the two goods. If they were vertically integrated monopolistic competitors selling directly to consumers, i.e. if they were not selling through a monopolist intermediary, the copyright owners in this model would set $l_1 = l_2 = 67.5$. In the circle model (or in a linear model, for that matter), the optimal price would increase with competition because each copyright owner is losing consumers with lower willingness-to-pay, while (by assumption) each copyright owner’s biggest fans, so to speak, remain as enthusiastic as ever.

But at these locations of 0 and 50 the presence of the intermediary prevents this result. Each copyright owner still chooses a licensing fee of 55 and the monopolist intermediary still chooses a retail price of 77.5. The double-marginalization process means that the overlap in demand for the two copyrighted works is not relevant. Each is selling to only a fraction of its potential consumers, those with the highest willingness to pay. Now, if the intermediary can negotiate a more flexible contract with each copyright owner, instead of being constrained to a per-copy licensing fee, then the intermediary might be able to implement the result that maximizes the firms' joint surplus at $P_1 = P_2 = 67.5$. Given that all three firms are earning the same profits under the double-marginalization result, perhaps the bargaining among them would reach the Nash solution. But the bilateral nature of the required negotiations, since the copyright owners cannot coordinate on price, might complicate that bargaining result.

Finally, consider the circle model with the copyright owners located even closer to each other at 0 and 25. In that event, the intermediary ends up having strong incentives to raise the price to a relatively high level, serving only the highest willingness-to-pay consumers. Through backward induction, the copyright owners choose relatively high licensing fees. I find an equilibrium at $l_1 = l_2 = 75$ and $P_1 = P_2 = 87.5$. Joint operating profits fall to 47.8% of the (two-)monopoly level, and the copyright owners collectively get 83.8% of that total, for an ultimate fraction of 40.1% of vertically-integrated-monopoly profits. This example is a good reminder that industrial organization models often involve multiple, conflicting incentives and can produce counter-intuitive results. The firms could all do much better, collectively, with a joint-surplus-maximizing price of 61.5, but the intermediary can't resist a very high price even when the licensing fees are fairly low. Again, if other contractual mechanisms are available and if the parties can agree on them through bilateral negotiations between copyright owner #1 and the intermediary and between copyright owner #2 and the intermediary, then there would be much surplus to share.

Without a doubt, this example is driven by the particular setup of the circle model and my particular assumptions about the parameters. My goal here is not, and cannot be, to

assert that this result is likely or that it matches a particular entertainment market. I just wanted to show some of the possibilities for duopoly copyright owners to prop their licensing fees up above marginal cost. The dynamics of the circle model suggest that even with more copyright owners, one could obtain this result, since by assumption as more firms enter, more consumers can purchase goods that delight them ever more. (The social tradeoff for that collective increase in utility among consumers, in Salop's model, is the fixed costs of entry for each firm, which I have been leaving to one side as I focus on operating profit.)

But if the number of copyright owners is large, and they are dealing with a monopolist intermediary, then the circle model seems like it leads us astray. The significant bargaining power of the intermediary starts to look like monopsony power. This kind of story resonates more with the example of YouTube that I mentioned at the beginning of the section – a vast number of copyright owners negotiating with one dominant technological intermediary. My aim is not to offer a single model to fit all scenarios. Instead, I think the models I have discussed in this section suggest how important the negotiations between copyright owners and intermediaries are for the financial outcomes of copyright owners. Models in industrial organization exist that can at least attempt to capture some of the key dynamics.

The models presented in this subsection assume that the products of the multiple copyright owners are either perfect or partial substitutes. None of the models, therefore, address complementary copyrighted works, including the situation of perfect complements that exists in the music industry with musical works and sound recordings. In other words, the discussion here of "multiple copyright owners" does not address the particular three-way structure of the bargaining in the music industry among music publishers, record labels, and technological intermediaries.

4.3. Multiple intermediaries and multiple copyright owners. Working with economics on the chalkboard, as I have been doing in this paper, it is important to remember that anything can happen on the chalkboard and also that anything can happen in the real

world, and very likely not what one has depicted on the chalkboard. That lesson is especially important when addressing the scenario I'd like to address in this section: multiple intermediaries and multiple copyright owners. If off-the-shelf models with two sellers (our copyright owners) and one buyer (our technological intermediary) are relatively scarce, then models of the members of one oligopoly engaging in bilateral negotiations with the members of another oligopoly are also scarce. The same goes for models of monopolistic competition in production together with monopolistic competition in retail. Monopoly and atomized competition can be complex enough. The market structures in between are all the more so.

Number of Firms	1 TI	2-5 TIs	6-15 TIs	More than 15 TIs
1 CO	Bilateral monopoly	Greater CO share likely	Approaching monopoly for CO	Monopoly for CO
2-5 COs	Greater TI share likely	Many possibilities	Many possibilities	Oligopoly for COs
6-15 COs	Approaching monopoly for TI	Many possibilities	Relatively low markups	Close to perfect competition
More than 15 COs	Monopoly for TI	Oligopoly for COs	Close to perfect competition	Like perfect competition

Table 1: Market Structure Among Copyright Owners and Technological Intermediaries

That said, and with an eye toward steering things back toward the practical, I think it is worthwhile to generate a table that lays out the various scenarios. The rows of Table 1 correspond to rough categories of possible upstream market structures among copyright owners and the columns correspond to the rough categories of possible downstream market

structure among intermediaries. The entries in the table are just a rough description of the market outcomes one might expect, based in part on the analysis in this paper and in part on general principles of industrial organization economics. Remember that by copyright owners I mean to refer to aggregators of copyrights, like book publishers, movie and TV studios, or record labels. Thus, with a category like “2-5 COs,” what I mean is that two to five copyright aggregators in an industry together have substantial market share. Of course each industry features many smaller, independent copyright-owning firms and individuals, and for some applications that extended market structure would be worth analyzing. But for my purposes here I just want to address broadly the large players who are likely to negotiate their own deals with technological intermediaries, or at least lead the way within an industry in making those deals.

In the four-by-four Table 1, the box in the top left corner corresponds to the situation addressed by the models in Sections 2 and 3. The two adjacent boxes would include the duopoly-on-one-side, monopoly-on-the-other-side models discussed in Sections 4.1 and 4.2. Even though all those boxes in or near the top left corner of the table certainly admit of a large range of possibilities, we can expect some kind of split of the surplus between the copyright owner and the intermediary.

The fourth row and the fourth column correspond to situations in which at least one side, or both sides, of the copyright owner – technological intermediary divide features more than 15 firms, such that even in Cournot competition one would expect a less than a 2% markup over marginal cost. Selling into (or buying) from such a market structure should mean that the market structure on the seller’s (or buyer’s) side largely dictates outcomes. But of course facts, circumstances, institutional details, or regulations could complicate that prediction.

The center of the table – the middle four boxes – is the region of interest for this section of the paper. And perhaps all that can be said is that there are “many possibilities,” as I have written in the table. Because this kind of interaction between market structures seems most common and realistic, I think a plausible research agenda for copyright economics would be to explore more models in this area. And, when that sort of analysis

is done, I think it would be helpful to use the twin metrics I have used in this paper to evaluate scenarios: the share of vertically-integrated-monopoly profits that the firms earn in aggregate and the share of those profits that the copyright owners receive. In this way, we can both compare different scenarios with some of the more tractable ideal types, while also keeping focus on the product of those two shares, which tells us how much operating profit the market has allocated to copyright owners. I hope by this point in the paper you are convinced that we should be more flexible than assuming all these fractions are 100%.

5. COPYRIGHT POLICY WITH VARIABLE-STRENGTH INTERMEDIARIES

In this final section of the paper, I would like to pull back from the abstract economic models and the running numerical example that I have used to illustrate them. I recognize that many members of the audience whose thinking about copyright I would like to change or at least augment are not going to engage with the economics or the math. So in this section I would like to review the key takeaways, translate some of the analysis here into plainer language, offer a way to visualize the framework, and suggest how some of this abstract analysis could be used as a tool for policy analysis. Finally, I conclude with some preliminary thoughts about how adopting a view of copyright economics with intermediaries might shape our approach to policy.

5.1. Interpretation and broader perspective for policy makers. I explained at the outset that I want to change the standard model in a way that accomplishes two goals, which we can now see are in some tension. I want the simple model we have in our heads to more realistically focus on the hugely important dynamic between copyright owners and (large, technological) intermediaries. And I want the model to be understandable. Optimization conditions with partial derivatives in them may take us well past what is digestible for my ultimate target audience for this model. So let me step back a bit before I proceed with further alterations, augmentations, and complications.

I argued in the sec:Introduction that the standard model of copyright economics is worth studying, discussing, and revising because it is what legislators, judges, regulators, lawyers, lobbyists, and copyright owners themselves have in mind when they make arguments about

copyright law and policy. Some of them may envision the class graph of a downward-sloping demand curve, a price set above marginal cost, a deadweight loss triangle, and a rectangle of operating profit to cover and exceed first-copy costs. Who knows – some people are geometric and visual thinkers. But perhaps many others don't think of the model explicitly at all. They think of the need to thwart copyists to protect copyright owners. That's important, of course – all the models in this paper assume the existence of copyright protection and the notional copyright owners in these models depend on it. But, as I have argued, the standard model leaves out the negotiations that shape so much about entertainment industries, especially today.

In light of this, I would like to offer a visual representation of what I see as the two key measures of a copyright industry that all the models in this paper can produce as predictions. It is a simple graph on a unit square with the percentage of vertically-integrated-monopoly profits on the x-axis and the copyright owner's share of profits generated on the y-axis. Multiplying these two fractions together gives us the percentage of vertically-integrated-monopoly profits that the copyright owner receives: the pile of money that the copyright owner earns to (hopefully) cover first-copy costs.

Figure 1 simply plots the percentages I reported in Sections 2.1, 2.3.1, and 2.3.2, respectively. Expressing the copyright owner's result as a share, and keeping the two component fractions of this share distinct in the two-dimensional graph, allows a story to be told. In one industry, consumer prices might be relatively high, but the copyright owner's share could be low. In another, consumer prices could be lower but the copyright owner's share could be higher. If, under whatever normative framework one is applying to copyright, one wishes to increase (or decrease) the financial incentives that copyright is providing, it would be relevant and important to know where things stand in terms of this type of diagram.




Figure 1: Copyright Owner's Result Under Vertically Integrated Monopoly, Double Marginalization, and Flexible Contracts

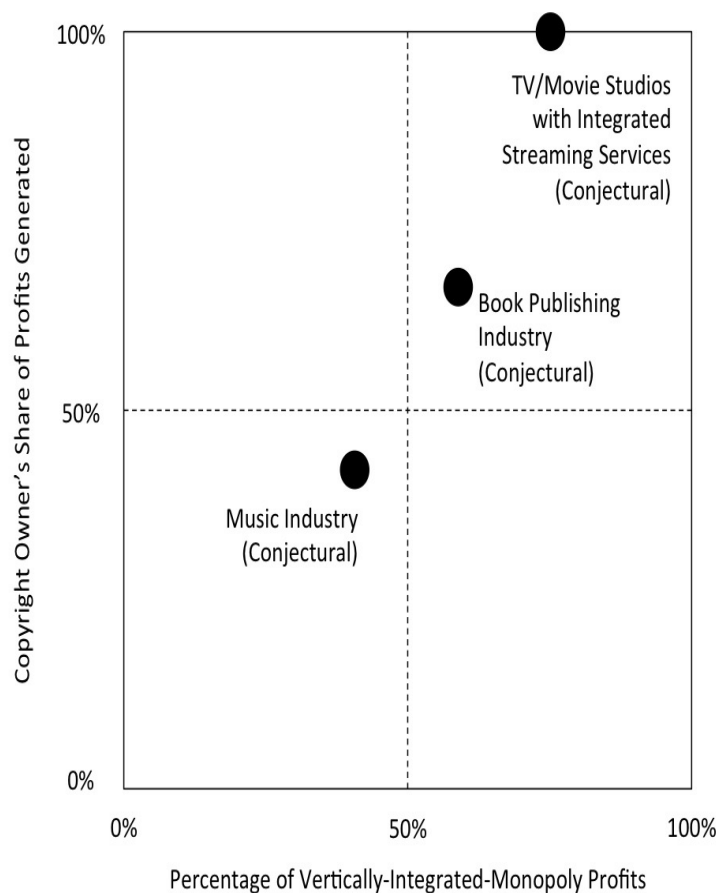


Figure 2: Conjectural Diagram of an Industry Comparison

Figure 2, on the other hand, is a conceptual diagram with purely conjectural figures. It takes the framework of this paper and this two-dimensional diagram and expresses my hope that researchers could actually collect real data to compare industries someday. My impression is that consumer prices and most likely operating profits in the music industry declined over the course of the shift from compact discs to file-sharing to licensed downloads to streaming. Moreover, intermediaries like Apple seemed to capture a substantial portion

of the surplus. So perhaps, in terms of the framework of paper, the music industry sits in the lower left-hand quadrant. The book industry, by contrast, has kept prices for books relatively high despite the shift to digital. So perhaps it remains in the (lower left-hand part of) the upper right-hand quadrant. Meanwhile, studios like Disney now control their own integrated streaming services, instead of sharing surplus with an intermediary. This purely notional graph is meant to illustrate the desirability of that state of affairs, for those with the resources, technology, organization, and circumstances to achieve it.

In addition to what I hope are helpful graphs, I would like here to summarize the main points of this paper, in a way that I hope would be quite accessible to legislators, judges, policy makers, attorneys, creators themselves, and other participants in the copyright system.

- Some copyright owners might be vertically integrated and deliver their works directly to consumers. But many other copyright owners rely on intermediaries, including intermediaries using digital, networked technology, to provide access to the copyrighted works they produce and market.
- Copyright owners and technological intermediaries might interact in a market, but more likely they will interact in bilateral contract negotiations. This makes the economics more interesting but less familiar and less predictable at the same time.
- The standard model of copyright economics assumes a vertically integrated copyright owner. To understand the important economic dynamics of the licensing interactions between copyright owners and technological intermediaries, we need new models to serve as metaphors and supply stories about what might be happening in the entertainment industries.
- Models of bilateral negotiations between sellers and buyers can be complicated, but one way to make apples-to-apples comparisons among them is to pay attention to two simple measures: (1) the percentage of vertically-integrated-monopoly profits and (2) the copyright owner's share of profits generated. Multiplying those two fractions together tells us how much money the copyright owner will earn compared to a vertically integrated monopolist.

- Many intermediaries are conglomerates, for which selling access to copyrighted works is complementary to selling many other goods. All else equal, this gives conglomerates incentive to set a lower price for access to copyrighted works. If the price was already higher than copyright owners preferred, this could benefit them. Otherwise, this effect intensifies the conflicting preferences of copyright owners and intermediaries.
- The market structures of the copyright-owning industry and the technological-intermediary industry both matter, but outcomes are probably hard to predict. Nonetheless, the nature of licensing negotiations between members of each industry should be the focus of research, because of the central role of those licensing negotiations in determining economic outcomes in the entertainment industries.
- We can learn about the state of various entertainment industries and compare their economic situations in terms of the outcomes of the bargaining between copyright owners and intermediaries.

5.2. Intermediating Copyright's Economic Functions. Technological intermediaries don't just mediate between the reading, viewing, and listening public and copyright owners. They also mediate between copyright policy and copyright owners. In other words, copyright policy works through technological intermediaries. Just as intermediaries filter, curate, and present the works that reach consumers, they filter the incentives that copyright law tries to give copyright owners. Some copyright owners have a direct line to their customers. But now, and over history, many have not. And the digital age has, in some sectors, expanded the power of the intermediaries to capture the value of providing access to copyrighted works.

My central argument in this paper is that the basic model of copyright economics should reflect the role of intermediaries. The standard model suppresses that role, implicitly assuming it away, distracting from the crucial negotiations over the licensing contracts that shape and produce prices, quantities, and profits in entertainment industries. To be sure, the standard public-good model of copyright has an important role; it has lasted and

it speaks to us for a reason. In one sense, it is a public finance approach to copyright, facilitating discussion of concepts like deadweight loss or optimal compulsory licensing rates. But perhaps the internet age calls for a shift in emphasis, perhaps only a partial one, from a public finance approach to an industrial organization approach. The models presented in this paper, gradually augmenting a model of bilateral monopoly between a producer and a retailer, are offered here not for their theoretical sophistication but to serve as metaphors and to illustrate how relevant various models of industrial organization could be to copyright economics.

Another way to see my argument is its highly descriptive or positive approach. Leaving normative considerations aside – because copyright, like any legal field, contains so many valid normative frameworks – I think it matters to all of us how much money is actually reaching copyright owners. Suppose the copyright owners in a particular industry are achieving low markups, losing share of profits to intermediaries, or both. Copyright law may have become a weak tool to allocate funding through the market to those copyright owners. That isn't in any way an argument against copyright; in the models of this paper, the negotiating leverage of the copyright owners, such as it is, depends on copyright. In other words, the models all assume copyright protection.

But just because copyright is necessary for obtaining licensing revenue from technological intermediaries does not mean it is sufficient. Other policy responses may be called for, depending on one's normative goals, so long as those goals hinge upon moving money from consumers to copyright owners and creators. I don't just mean looking to substitutes for copyright policy like grant funding, although I personally think that might be wise and just. I also mean potentially looking to other bodies of law, such as antitrust law, to regulate technological intermediaries and reign in their bargaining power against copyright owners. In addition, I wonder whether copyright law itself needs to expand and change. Copyright has long regulated communications and media firms (Wu 2004). But perhaps it should do so more explicitly, and, where appropriate, enhance (or limit) the bargaining leverage of copyright owners as they negotiate with intermediaries.

If I have convinced you that the dynamic interactions between copyright owners and technological intermediaries have always been important to copyright policy, and if you suspect that these interactions have only become more important during the internet era, then I have achieved my goal. The next steps on the theoretical side would be to make progress on applied models of conglomerate intermediaries and subscription services, beyond the sketches I have offered here, as well as models of multiple copyright owners negotiating bilaterally with multiple intermediaries. On the empirical side, despite the structural and pragmatic challenges to acquiring useful data, I hope that the models of this paper might lead to a framework simple and broad enough that rough estimates could be made of where things stand in various entertainment industries. I continue to think, perhaps more than ever, that it would be fascinating, though challenging, to measure as much as possible about the results of the licensing interactions between copyright owners and technological intermediaries.

REFERENCES

- Abramowicz, M. (2004)**, “An Industrial Organization Approach to Copyright Law”, *William and Mary Law Review*, 46; 33–124.
- Bolton, P. and M. Dewatripont (2005)**, *Contract Theory*, Cambridge, MA: MIT Press.
- Cabral, L.M.B. (2000)**, *Introduction to Industrial Organization*, Cambridge, MA: MIT Press.
- Dwyer, C. (2014)**, “Amazon and Hachette Reach a Deal on E-book Pricing”, *The Two-Way (NPR.org)*, November 13, 2014, available at <https://www.npr.org/sections/thetwo-way/2014/11/13/363794096/amazon-and-hachette-reach-a-deal-on-e-book-pricing>.
- Galloway, S. (2017)**, *The Four: The Hidden DNA of Amazon, Apple, Facebook, and Google*, New York: Penguin.
- Hotelling, H. (1929)**, “Stability in Competition”, *Economic Journal*, 39; 41–57.
- Inderst, R. and Mazzarotto, N. (2008)**, “Buyer Power in Distribution”, in *ABA Antitrust Section Handbook: Issues in Competition Policy Vol. 3*; 153–178.
- Knopper, S. (2009)**, *Appetite for Self-Destruction: The Spectacular Crash of the Record Industry in the Digital Age*, New York: Free Press.
- Nicholson, W. (1995)**, *Microeconomic Theory: Basic Principles and Extensions*, 6th ed., Fort Worth: Dryden Press.
- Owsinski, B. (2016)**, “Get Ready for the Amazon Streaming Music Disruption”, *Forbes*, June 9, 2016, available at <https://www.forbes.com/sites/bobbyowsinski/2016/06/09/amazon-streaming-music-disruption/>.
- Plant, A. (1934)**, “The Economic Aspects of Copyright in Books”, *Economica*, 1; 167–195.
- Salop, S.C. (1979)**, “Monopolistic Competition with Outside Goods”, *Bell Journal of Economics*, 10; 141–156.

- Scotchmer, S. (2004)**, *Innovation and Incentives*, Cambridge, MA: MIT Press.
- Shapiro, C. and H.R. Varian (1999)**, *Information Rules: A Strategic Guide to the Network Economy*, Boston: Harvard Business School Press.
- Snyder, C.M. (2008)**, “Countervailing Power”, in Durlauf, S.N. and L.E. Blume, (eds), *The New Palgrave Dictionary of Economics*, 2nd ed., London: Palgrave Macmillan; pp. 1188-89.
- Tirole, J. (1988)**, *The Theory of Industrial Organization*, Cambridge, MA: MIT Press.
- Touve, D. and P. DiCola (2014)**, “Licensing in the Shadow of Copyright”, *Stanford Technology Law Review*, 17; 397–459.
- Towse, R. (2020)**, “Dealing with Digital: Economic Organisation of Streamed Music”, CIPPM/Jean Monnet Working Papers, No. 01-2020, Centre for Intellectual Property Policy and Management.
- U.S. Copyright Royalty Judges (2022)** *In re Determination of Royalty Rates and Terms for Making and Distributing Phonorecords (Phonorecords III)*, Docket No. 16-CRB-0003-PR, Initial Ruling and Order After Remand, July 1, 2022, pp. 21-44, available at <https://app.crb.gov/document/download/27063>.
- Viscusi, W.K., J.M. Vernon, and J.E. Harrington, Jr. (2000)**, *Economics of Regulation and Antitrust*, 3rd ed., Cambridge, MA: MIT Press.
- Wu, T. (2004)**, “Copyright’s Communication Policy”, *Michigan Law Review*, 103; 278–366.
- Yoo, C. (2004)**, “Copyright and Product Differentiation,” *New York University Law Review*, 79; 212–280.

PETER DICOLA IS A PROFESSOR OF LAW AT NORTHWESTERN PRITZKER SCHOOL OF LAW. PROF. DICOLA CAN BE CONTACTED AT P-DICOLA@LAW.NORTHWESTERN.EDU.