Social Media and News: Attention Capture via Content Bundling

Alexandre de Cornière and Miklos Sarvary*

January 4, 2018

Abstract

The growing influence of internet platforms acting as content aggregators is one of the most important challenges facing the media industry. We develop a simple model to understand the impact of content bundling by a social platform. In our model consumers can access news either directly through a newspaper's website, or indirectly through a platform, which also offers social content. Even though the platform shares revenues with newspapers whose content it publishes, content bundling harms newspapers. Its effect on news quality and news consumption depends on the media market structure and on whether the platform can personalize the content bundle.

Keywords: User-Generated Content (UGC), Media Competition, News Quality

*Alexandre de Cornière (alexandre.de-corniere@tse-fr.eu) is Assistant Professor of Economics at the Toulouse School of Economics, Université Toulouse I Capitole, and Miklos Sarvary (miklos.sarvary@columbia.edu) is the Carson Family Professor of Business at Columbia Business School. We thank Charles Angelucci, Alessandro Bonatti, Anna D'Annunzio, Benjamin Ho, Bang Nguyen, Markus Reisinger, Martin Peitz, Paul Seabright, Sandro Shelegia, participants of the 4th Annual BU Platform Symposium, 2017 Petralia Sottana Workshop, the 15th Annual Media Economics Workshop and the 2017 Net Institute conference for helpful comments and discussions. We are grateful to the Net Institute (www.netinst.org) for financial support. All errors are ours.

1 Introduction

With hundreds of millions of daily active users, a few large social networks have become the dominant online media outlets for most people. The largest among these, Facebook has reached over two billion active members across the globe who, on average, spend about an hour each day on the platform. In line with its significant consumer attention share, Facebook is estimated to capture about $37 billion of advertising revenues in 2017 corresponding to 16% of worldwide digital advertising. Other successful social platforms include Tencent's WeChat in China and VKontakte in Russia.

If, in their early days, social networks were mostly used as a way for users to share personal stories and pictures, their role has progressively evolved into one of content aggregation: an important share of the content displayed on their websites is produced by third-party publishers, who use the platforms as an alternative to their own website to reach consumers.

The news industry in particular has been affected by this change: studies show that more than 50% of consumers use social media as a source of news, and 14% as their main source (Gottfried and Shearer, 2016; Mitchell et al., 2017; Reuters, 2016). Facebook has recently surpassed Google as the main external source of traffic to newspapers' websites (Alpert, 2015; Constine, 2016).[1] This situation is a double-edged sword for publishers: social platforms provide the opportunity to reach a wider audience, yet newspapers worry about the growing power of platforms, for fear of losing their privileged relationship with readers, and eventually most of their revenues (Thompson, 2017).

The new role of social media in the news industry has recently been the subject of heated controversies. Platforms have been accused of fostering echo chambers, or of not doing enough to prevent the spread of fake news. Some of their critics argue that platforms should be held responsible for the content displayed on their websites.[2]

In this paper, we focus on another key feature of social platforms as it relates to news consumption, namely their ability to curate the content that consumers are exposed to, by providing a mix of *user-generated content* (UGC) and of professional content produced by third parties (in particular news organizations), a practice we refer to as *content bundling*. A major strategic choice faced by the platforms then concerns the design of the "newsfeed",

[1]For some news providers, Facebook's dominance is even more pronounced. For example, Buzzfeed, a leading online publisher valued at close to $1.5 billion derives 75% of its traffic from Facebook.

[2]For research about these issues, see for instance Bakshy, Messing, and Adamic (2015) or Allcott and Gentzkow (2017).

i.e. the relative prominence of UGC and news on the platform.[3]

We develop a simple model of competition for attention between a social platform and a newspaper that allows us to make two contributions. First, we shed light on a platform's incentives to use content bundling as a strategy to capture consumers' attention, to the point where it becomes a gatekeeper. Our second contribution is then to draw out the implications of content bundling on the news industry, in particular with respect to news quality, newspapers' profits and consumers' news consumption. We show in particular that content bundling is likely to harm newspapers, and that its effect on quality depends on the newspapers' market structure.

In our basic framework, a social platform and a newspaper (or publisher), both advertising-supported, compete for consumers' attention. The newspaper produces news stories and maintains a website which only offers news content. The social platform relies on its users to produce *user-generated content* (UGC), such as personal stories or pictures. On its website, alongside UGC, the platform can also show news stories produced by the publisher. We refer to the strategy of showing a mix of news and UGC as content bundling. Under content bundling, the platform gives a share of its revenues to the publisher.

Consumers have limited attention, and are heterogeneous in their demand for news, that is in the share of their attention they would like to devote to news content. The demand for news also depends on its quality, which is the result of an investment by the newspaper. Consumers can freely allocate their attention across the two websites, but, when on the platform, have to consume the bundle that is offered to them.

By bundling news and UGC, the platform diverts news consumption away from the newspaper's website, and onto its own. We show that such a strategy is profitable as long as enough consumers have a positive demand for news.

We compare the equilibrium to a benchmark without content bundling, i.e. in which the platform only displays UGC. We find that the newspaper's profit is lower with content bundling. Even though the platform showing news is not necessarily bad for the newspaper, the latter is harmed by content bundling because the platform *strategically* chooses the mix of content it offers so as to maximize its revenue, which comes at the expense of the newspaper's revenue. Content bundling also causes news consumption to be distorted upwards. Whether quality increases or decreases depends on the relative magnitude of a softening effect (negative - the newspaper's gain to attracting traffic is lower if it gets a

[3]See for instance Constine (2016) and Bradshaw (2017) for reports on how platforms such as Facebook and Snapchat have redesigned their newsfeeds.

share of the platform's revenues) and a composition effect (ambiguous - traffic may become more or less sensitive to quality under content bundling).

Our baseline model, analyzed in Sections 3 and 4, considers a single newspaper and relies on a number of simplifying assumptions. In Section 5 we discuss the robustness of our results by studying variations of the model that bring back important institutional details. Specifically, we allow for the personalization of the platform's content, the possibility for the newspaper to prevent content bundling, and the presence of large switching costs across websites for consumers.

In Section 6 we then extend our framework to allow for competition between newspapers. We again show that newspapers are made worse-off by content bundling, but we find that news quality increases.

2 Relevant literature

The paper is related to a number of literature streams, first, and foremost to the broadening literature on news/media "aggregators" (see Peitz and Reisinger (2015) for an extensive summary on this literature). As in our paper, the central question is how these intermediaries impact the consumption of news as well as the quality of content produced. On the theory side, Jeon and Nasr (2016) and Dellarocas, Katona, and Rand (2013) model aggregators as enabling consumers to find high quality news more easily. They find that the entry of an aggregator tends to increase competition among websites, leading to higher quality. The impact on newspapers profit depends on which effect is stronger: business stealing or market expansion. Rutt (2011) studies how the presence of an aggregator affects newspapers' choice of business model, and shows that it has different effects on the quality provided by free versus paying outlets. In George and Hogendorn (2012), the aggregator reduces the cost of multihoming for consumers. Unlike here, in these papers aggregators are non-strategic and do not produce their own content, but merely replicate the experience of a newspaper. Even though we also have a trade-off between business stealing and market expansion, our focus on social network leads us to emphasize a different set of issues.

A recent series of empirical papers examine the impact of aggregators on the news industry. Using disputes between Google News and Spanish publishers (Athey, Mobius, and Pal (2017), Calzada and Gil (2016)) or the Associated Press (Chiou and Tucker (2015)), empirical research finds that Google News increases overall news consumption. In particular, Athey, Mobius, and Pal (2017) document that this effect is mostly present for small publishers, who cannot rely on brand recognition to attract users and therefore benefit most from the aggregator. In relation to the theoretical work on aggregators, these papers suggest that the demand-expansion effect of aggregators dominates. George and Hogendorn (2013) studies the consequences of a redesign of Google News, and find that news aggregators can potentially change the composition of news consumption.

Our work specifically focuses on social networks as news intermediaries, the major difference being that these platforms also host user-generated content (UGC) that directly competes with the content of publishers (see Luca (2015) for a summary of the economics literature on UGC). This is relevant because, increasingly, it is such platforms (as opposed to search engines) that generate traffic to news content. Yildirim, Gal-Or, and Geylani (2013) study the effect of UGC on the horizontal competition between news providers, but they do not consider the presence of an endogenous intermediary as we do. Theoretical research on UGC and social networks specifically is scarce and focuses mostly on network formation.[4]

In our model, the platform allocates consumers' attention by choosing the mix of content that it displays. In this respect it is similar to a search engine, which allocates traffic through its ranking and design (see de Cornière and Taylor (2014) or Burguet, Caminal, and Ellman (2015)). However, in these papers the intermediary enjoys an exogenous bottleneck position: consumers have to use the search engine to find content. In contrast, our mechanism is one where the "within platform" allocation of attention (i.e. content bundling) determines how consumers allocate their attention between the platform and the newspaper. The gatekeeping role of the platform thus emerges endogenously.

Our framework assumes multi-homing but we abstract away from the core concern of the multi-homing literature applied to media, namely that it may lead to inefficient (duplicate) advertising when an advertiser is present on multiple publishers (see, Ambrus, Calvano, and Reisinger (2014), Athey, Calvano, and Gans (2017), and Anderson, Foros, and Kind (2016) for a detailed treatment of this issue). As Alaoui and Germano (2016), we also assume that consumers are time constrained in their consumption of media and our results resonate to theirs in that competition between content suppliers (including the social network) distort consumers' media consumption. However, we focus on consumers' time allocation across qualitatively different content providers and we abstract away from the editorial process of publishers when multiple topics are present.

Finally, our paper is related to the literature on bundling (see e.g. Nalebuff (2003) for an overview). The reason for content bundling differs from some of the standard explanations of bundling, such as price-discrimination or leverage of market power. Here,

[4]See, for example Bala and Goyal (2000) and Jackson and Wolinsky (1996) for earlier models, and Jackson (2010) for a review. See also Zhang and Sarvary (2015) who consider local network effects.

the platform bundles its own content with that of its rival in order to capture more attention from consumers.

3 Baseline model

We consider a model where consumers can consume two kinds of content: news and user-generated content (UGC). To convey the core intuitions we make a number of simplifying assumptions that we relax in Sections 5 and 6. Specifically, we start with a monopolist newspaper (indexed by 1), who must invest $c(q)$ to produce news stories of quality q, where c is increasing and convex. User-generated content is produced by users of a monopolist social platform (indexed by 0), at no cost for the platform. UGC quality is exogenous.

Consumers have heterogeneous preferences regarding content. A consumer of type θ who consumes a quantity x of news (of quality q) and y of UGC derives a utility $U(x, y, q, \theta)$, non-decreasing in x and y. We assume that $U_{x,\theta} \geq 0$,[5] i.e. that high types have a larger marginal utility for news content. News quality increases the marginal utility of news consumption: $U_{x,q} > 0$. However, this effect is weaker for higher levels of quality: $U_{x,q,q} \leq 0$.[6] We assume that θ is distributed according to a continuous c.d.f. F, of density f, on a support $[\underline{\theta}, \overline{\theta}]$. For simplicity we assume that the distribution of types has no atoms. In the baseline model, θ is a consumer's private information. We relax this assumption in Section 5, when we allow the platform to personalize consumers' newsfeed.

Consumers have an *attention constraint*: $x + y \leq 1$.[7] For a given quality q, a type θ consumer's demand for news $\hat{x}(\theta, q)$ is the solution to

$$\max_{x,y} U(x, y, q, \theta) \quad \text{s.t} \quad x \geq 0, \quad y \geq 0 \quad \text{and } x + y \leq 1.$$

From our assumptions, $\hat{x}(\theta, q)$ is non-decreasing in both its arguments. Moreover, we assume that $\hat{x}(\theta, q) > 0$ for all $\theta > \underline{\theta}$, and that $\hat{x}(\theta, q)_{q,q} \leq 0$. Similarly, $\hat{y}(\theta, q)$ is the demand for UGC. We assume that consumers have no outside option, so that the attention constraint is always binding and $\hat{y}(\theta, q) = 1 - \hat{x}(\theta, q)$.

[5] $U_{x,\theta}$ is the cross derivative of U with respect to x and θ.

[6] $U_{x,q,q}$ is the third-order partial derivative. At this point, we impose no restriction on the sign of $U_{x,q,\theta}$, that is, we do not specify whether high types' or low types' demand for news is more sensitive to quality.

[7] Empirical evidence supports the notion that people allocate a finite time across different media outlets (see, for example, Boik, Greenstein, and Prince, 2016). Moreover, the main results hold in a model where the total time $t_0 + t_1$ is endogenous and generates an opportunity cost $K(t_0 + t_1)$. Such a model is less tractable though, in particular under competing publishers.

Example: For the sake of illustration, we sometimes use the following utility function:

$$U(x, y, q, \theta) = (\alpha(q + \theta) + \beta\theta q) \ln(x) + y \tag{1}$$

The associated demand for news is $\hat{x}(\theta, q) = \max\{\min\{\alpha(q + \theta) + \beta\theta q, 1\}, 0\}$.[8] We then assume that $\alpha \geq 0$, and that the parameters of the model are such that $\frac{\partial \hat{x}}{\partial \theta} \geq 0$ and $\frac{\partial \hat{x}}{\partial q} \geq 0$. When $\beta > 0$, high types' demand for news is more sensitive to quality than low type's demand. We refer to this as the *linear model*. Special cases include $(\alpha, \beta) = (1, 0)$ (what we refer to as the *additive model*, particularly convenient to obtain closed form solutions) and $(\alpha, \beta) = (0, 1)$ (the *multiplicative model*).

Even though consumers have preferences over contents, they cannot directly choose which content they consume. Instead, they allocate their unit of attention across two websites: one operated by the newspaper, and one by the platform. While the newspaper's website can only offer news content, the key feature of our model is the platform's ability to display news from the newspaper alongside its own UGC. Such *content bundling* is a strategic choice: the platform decides the share λ of news that consumers are exposed to when they visit its website. If a consumer spends t_0 units of time on the platform's website, he therefore consumes a quantity $t_0(1 - \lambda)$ of UGC, and a quantity $t_0\lambda$ of news (on top of the news he gets directly from the newspaper's website).

Websites are purely advertising-supported. We normalize the monetary value of one unit of attention by a consumer to one.[9] Thus, when a consumer spends t_1 units of time on the newspaper's website (what we call direct traffic), the newspaper generates direct revenues of t_1. The newspaper also derives revenues from indirect traffic, i.e. from the news stories that consumers are exposed to while on the platform's website. More specifically, we assume that if the platform shows a share λ of news and if a consumer spends t_0 units of time on its website, the newspaper's indirect revenue is $t_0\lambda(1 - \phi)$, where $\phi \in [0, 1]$ is the share of news-related ad revenues that the platform keeps for itself. The platform's revenue is then $t_0(1 - \lambda + \lambda\phi)$.

One can interpret advertising revenue sharing between the platform and the newspaper either as explicit payments between them, or as capturing the idea that the consumer divides his attention between the two. In the baseline model, the revenue sharing rule $(\phi, 1 - \phi)$ is exogenous, but we relax this assumption in Section 5.3.

[8] In the rest of the paper we drop the max and min operators to ease notations, but demand for news and for UGC should always be thought of as being between 0 and 1.

[9] A priori, we have no reason to assume that either website is more efficient at advertising. We discuss this assumption below.

Timing and equilibrium: The timing is as follows: at $\tau = 1$, the newspaper chooses a quality q, publicly observed, and incurs the cost $c(q)$. We view q as a long-term strategic choice. At $\tau = 2$, the platform chooses the share of news λ it shows to its users. At $\tau = 3$, consumers observe λ and choose $t_0(\theta, q, \lambda)$, the time they spend on the platform as a function of their type, of the quality of news and of the platform's content mix. We look for subgame-perfect equilibria.

4 Equilibrium analysis

4.1 Benchmark: UGC-only newsfeed

As a benchmark, we start with the case in which the platform cannot bundle news content alongside UGC (i.e. $\lambda = 0$).

After observing q, consumers choose how much attention to allocate to the platform and to the newspaper. Because the platform only offers UGC, and there are no costs associated to switching from one media to the next, consumers can consume their desired mix of content. A consumer of type θ then spends $\hat{x}(\theta, q)$ on the newspaper site, and $1 - \hat{x}(\theta, q)$ on the platform. The total time spent on the newspaper's website is therefore

$$T_1(q, \lambda)|_{\lambda=0} = \int_{\underline{\theta}}^{\overline{\theta}} \hat{x}(\theta, q) dF(\theta). \tag{2}$$

Profit is concave in q, and the optimal quality for the newspaper, denoted $\tilde{q}$, is the solution to $\max_q T_1(q, 0) - c(q)$, that is $\tilde{q}$ solves:

$$\frac{\partial T_1(\tilde{q}, 0)}{\partial q} = c'(\tilde{q}). \tag{3}$$

Let $\tilde{\pi}_1$ be the newspaper's associated profit. We now turn to the analysis of the game where the platform can freely choose λ, and proceed by backward induction.

4.2 Consumers: allocation of attention with content bundling

At $\tau = 3$, if news quality is q, a consumer of type θ would like to consume a quantity $\hat{x}(\theta, q)$ of news. By spending t_0 units of time on the platform, and $1 - t_0$ on the newspaper, he gets a quantity of news, $x(t_0, \lambda) = t_0\lambda + (1 - t_0)$ and a quantity of UGC, $y(t_0, \lambda) = t_0(1 - \lambda)$.

If $\lambda \geq \hat{x}(\theta, q)$, the consumer's demand for news is more than satisfied by the platform alone. Such a consumer then decides to spend all his time on the platform, $t_0(\theta, q, \lambda) = 1$. Consumers such that $1 > \hat{x}(\theta, q) > \lambda$ can achieve their optimal content mix by spending

$t_0(\theta, q, \lambda)$ on the platform such that

$$t_0(\theta, q, \lambda)(1-\lambda) = \hat{y}(\theta, q) \Leftrightarrow t_0(\theta, q, \lambda) = \frac{\hat{y}(\theta, q)}{1-\lambda}.$$

Finally, if $\hat{x}(\theta, q) = 1$, the consumer allocates all his attention to the newspaper, that is $t_0(\theta, q, \lambda) = 0$.

We denote by $\hat{\theta}_1(q, \lambda)$ the solution to $\hat{x}(\theta, q) = \lambda$, i.e. the largest type who does not visit the newspaper, and by $\hat{\theta}_2(q)$ the smallest solution to $\hat{x}(\theta, q) = 1$, i.e. the lowest type who does not visit the platform. We sometimes omit the arguments and simply write $\hat{\theta}_1$ and $\hat{\theta}_2$. Summarizing, we obtain Lemma 1 below.

Lemma 1. ***(Optimal allocation of attention)*** *When the newspaper is of quality q and the platform shows a share λ of news content, a consumer of type θ allocates a share $t_0(\theta, q, \lambda)$ of his attention to the platform, where*

- $t_0(\theta, q, \lambda) = 1$ *if* $\theta \leq \hat{\theta}_1$,
- $t_0(\theta, q, \lambda) = \frac{1-\hat{x}(\theta,q)}{1-\lambda}$ *if* $\theta \in \left(\hat{\theta}_1, \hat{\theta}_2\right)$,
- $t_0(\theta, q, \lambda) = 0$ *if* $\theta \geq \hat{\theta}_2$.

In the benchmark where $\lambda = 0$, consumers allocate a share $\hat{x}(\theta, q)$ of their attention to the newspaper. When $\lambda > 0$, that share is lower, because part of the demand for news is already satisfied by visiting the platform. More generally, any increase in λ shifts attention from the newspaper to the platform, a point we elaborate on when we discuss the choice of λ. While this effect does not directly affect consumers such that $\theta > \hat{\theta}_1$, whose consumption of news is still $\hat{x}(\theta, q)$, it introduces a consumption distortion on lower types, who, even though they stop visiting the newspaper, end up consuming too much news relative to what they would like ($\lambda > \hat{x}(\theta, q)$).[10] Figure 1 illustrates the link between demand for news and news consumption under content bundling.

Note also that, for a given quality q, total news consumption increases with λ. Indeed, a consumer's news consumption is $\max\{\hat{x}(\theta, q), \lambda\}$. Whether equilibrium news consumption increases is more ambiguous, as it will depend on the effect of λ on q, an effect we examine below.

[10] "Too much news" does not mean that consumers are forced to consume news content that brings them negative utility. They enjoy the news content, but would prefer UGC instead.

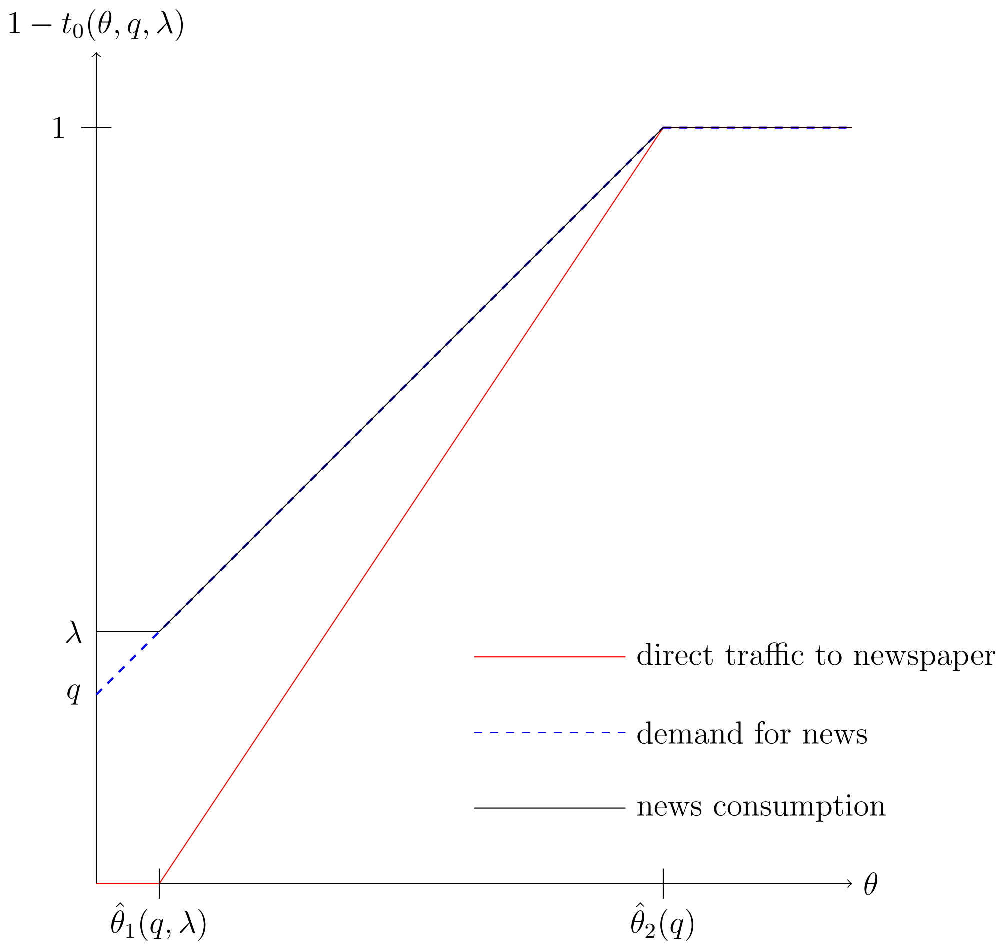

Figure 1: Individual demand for news and direct traffic to newspaper under content bundling when $\hat{x}(\theta, q) = \theta + q$.

4.3 Platform: optimal content bundling

Suppose that news quality is q. If the platform displays a share λ of news content, the total amount of attention that it receives is

$$T_0(q, \lambda) = \int_{\underline{\theta}}^{\overline{\theta}} t_0(\theta, q, \lambda) dF(\theta). \tag{4}$$

Each unit of attention generates a revenue $(1 - \lambda + \lambda\phi)$, so that the platform's profit is $\pi_0(q, \lambda) = (1 - \lambda + \lambda\phi) T_0(q, \lambda)$.

The platform's trade-off is the following: by showing more news content (increasing λ), the platform can receive more of the consumers' attention, by the logic discussed in the previous subsection. However, showing more news leads to lower advertising revenue per-unit of attention. The next proposition is our first main result:

Proposition 1. *The platform always shows some news content:* $\lambda^* > 0$.

Proof. The derivative of the platform's profit with respect to λ is

$$\frac{\partial \pi_0(q,\lambda)}{\partial \lambda} = -(1-\phi)T_0(q,\lambda) + (1-\lambda(1-\phi))\frac{\partial T_0(q,\lambda)}{\partial \lambda}. \tag{5}$$

Using Lemma 1 and equation (4), we have

$$\frac{\partial T_0(q,\lambda)}{\partial \lambda} = \frac{\partial \hat{\theta}_1}{\partial \lambda} f(\hat{\theta}_1) - \frac{\partial \hat{\theta}_1}{\partial \lambda} \underbrace{\frac{1-\hat{x}(\hat{\theta}_1,q)}{1-\lambda}}_{=1} f(\hat{\theta}_1) + \int_{\hat{\theta}_1}^{\hat{\theta}_2} \frac{1-\hat{x}(\theta,q)}{(1-\lambda)^2} dF(\theta). \tag{6}$$

Evaluating the derivative of the profit at $\lambda = 0$, we thus get

$$\frac{\partial \pi_0(q,0)}{\partial \lambda} = \phi T_0(q,0) - F\left(\hat{\theta}_1(q,0)\right). \tag{7}$$

By our assumption that $\hat{x}(\theta,q) > 0$ for all $\theta > \underline{\theta}$, we have $\hat{\theta}_1(q,0) = \underline{\theta}$. Because F is atomless, we thus obtain $\frac{\partial \pi_0(q,0)}{\partial \lambda} = \phi T_0(q,0) > 0$. This proves the result. ■

The intuition for Proposition 1 is the following. When $\lambda = 0$, consumers get all their news from the publisher. By slightly increasing λ, the platform displaces part of this news consumption towards itself, thereby increasing the attention it receives, attention valued at ϕ. The potential cost of doing so is to show news to some users who would have preferred to see UGC. Because λ is very small, this effect is negligible.

The result is sensitive to our assumption that the mass of consumers who want to consume no news is zero. Indeed, if such a mass was large enough, the cost of serving news instead of UGC to these "news drop-outs" could overcome the benefit from increased attention by news consumers. However, note that the result would also holds if the mass of consumers whose demand for news is zero was small enough.

Examples: To better understand some of the forces that determine the optimal choice of λ, we use the linear model (see Equation 1) assuming that the distribution of types is uniform on $[0,1]$. Using Lemma 1 to obtain individual demands, and integrating these demands over the set of types, we find that the total time spent on the platform is:[11]

$$T_0(\lambda,q) = \frac{1+\lambda-2\alpha q}{2(\alpha+\beta q)}. \tag{8}$$

It follows that $\lambda^*(q) = \alpha q + \frac{\phi}{2(1-\phi)}$. The first term ($\alpha q$) corresponds to the demand for news of the lowest type ($\theta = 0$). Because all consumers want to consume at least a quantity αq of news, the platform must choose $\lambda \geq \alpha q$. Beyond this quantity, the platform's optimal

[11]Here we only write T_0 for cases where $\lambda \geq \alpha q$, which is always true in equilibrium.

strategy depends on the share ϕ of revenues it captures when it shows news: for large values of ϕ the platform has an incentive to show a lot of news content to its users.

4.4 Newspaper: choice of quality

Besides understanding the strategic incentives of the platform to provide news content to its users, we seek to assess the effects of content bundling on the news industry, i.e. on newspaper's profit and choice of quality. The newspaper's profit is

$$\pi_1(q,\lambda) = T_1(q,\lambda) + (1-\phi)\lambda T_0(q,\lambda) - c(q) \equiv R_1(q,\lambda) - c(q),$$

where $R_1(q,\lambda)$ denotes the newspaper's advertising revenues. We assume that the primitives are such that this profit is concave in q.[12] Similarly, define $R_0(q,\lambda) \equiv (1-\lambda(1-\phi))T_0(q,\lambda)$, which represents the platform's revenues. Notice that $R_0(q,\lambda) + R_1(q,\lambda) = T_0(q,\lambda) + T_1(q,\lambda) = 1$ for any (q,λ).

In period $\tau = 1$, acting as a Stackelberg leader, the newspaper knows that the platform will choose $\lambda = \lambda^*(q)$. Its objective function is thus

$$\pi_1(q,\lambda^*(q)) = R_1(q,\lambda^*(q)) - c(q) = 1 - R_0(q,\lambda^*(q)) - c(q).$$

Because $\lambda^*(q)$ maximizes $R_0(q,\lambda)$, the envelope theorem implies that $\frac{d\pi_1(q,\lambda^*(q))}{dq} = \frac{\partial \pi_1(q,\lambda^*(q))}{\partial q}$. Using the notation $\lambda^* = \lambda^*(q^*)$, the newspaper's first-order condition then writes

$$(1-(1-\phi)\lambda^*)\frac{\partial T_1(q^*,\lambda^*)}{\partial q} = c'(q^*). \tag{9}$$

Comparing (3) and (9), one can distinguish two effects of content bundling by the platform: a softening and a composition effect. The softening effect corresponds to the smaller return to a marginal increase in direct traffic T_1, from 1 (in the benchmark) to $1-(1-\phi)\lambda^*$. When the platform bundles content, the newspaper collects a share of its revenues, and increasing T_1 is less valuable. The softening effect reduces the incentives to invest under content bundling.

The composition effect works as follows: under content bundling, direct traffic to the publisher $T_1(q,\lambda^*)$ only comes from consumers such that $\hat{x}(\theta,q) > \lambda^*$, whereas in the benchmark direct traffic $T_1(q,0)$ comes from all consumers. The number of consumers who adjust their viewing pattern following an increase in q is thus smaller under content bundling. However, because the time spent on the newspaper's website by these consumers is $1 - t_0^*(\theta,q,\lambda^*) = \frac{\hat{x}(\theta,q)-\lambda^*}{1-\lambda^*}$, it is more responsive to an increase in q than under the

[12] Concavity holds for instance in the linear model with uniform distribution of types.

benchmark, where $1 - t_0^*(\theta, q, 0) = \hat{x}(\theta, q)$. The overall sign of the composition effect, and therefore the effect of content bundling on news quality, is ambiguous in general. We can say more when we adopt a linear specification:

Proposition 2. *Suppose that $\hat{x}(\theta, q) = \alpha(\theta + q) + \beta\theta q$, and that θ is uniformly distributed over $[0, 1]$. Then:*

(i) If $\beta \leq 0$, quality goes down under content bundling compared to the benchmark.

(ii) If $\beta > 0$ and $\phi \to 1$, quality is higher under content bundling.

Proof. When $\hat{x}(\theta, q) = \alpha(\theta + q) + \beta\theta q$, the total time spent on the newspaper's website is $T_1(q, \lambda) = \frac{2(\alpha + q(\alpha+\beta)) - (1+\lambda)}{2(\alpha+\beta q)}$, and therefore $\frac{\partial^2 T_1(q,\lambda)}{\partial q \partial \lambda} = \frac{\beta}{2(\alpha+\beta q)^2}$. If $\beta \leq 0$, we have $\frac{\partial^2 T_1(q,\lambda)}{\partial q \partial \lambda} \leq 0$, which implies that direct traffic to the newspaper is less sensitive to quality under content bundling compared to the benchmark. This means that both effects (softening and composition) go in the same direction, and quality is unambiguously lower when the platform bundles content than under the benchmark. If $\beta > 0$ then T_1 becomes more sensitive to increases in q under content bundling, and so the two effects go in opposite direction. When ϕ is large enough (e.g. $\phi \to 1$) then the softening effect vanishes (as the newspaper gets no revenue from indirect traffic), and only the composition effect remains, which leads to a higher quality than under the benchmark. ■

Intuitively, under content bundling, low types consume news mostly from the platform. Direct traffic to the newspaper (T_1) is therefore mostly composed of high types. When $\beta < 0$, direct traffic is thus less sensitive to q than under the benchmark, whereas the opposite holds when $\beta > 0$. When the sensitivity of direct traffic is higher under content bundling (i.e. when $\beta > 0$), the effect also need to be strong enough so as to offset the softening effect. When ϕ is close to 1 the softening effect is small: the newspaper does not value indirect traffic enough for competition to be relaxed.

Newspaper profits

While in equilibrium news quality may increase or decrease, we find that the newspaper's profit unambiguously declines with content bundling by the platform:

Proposition 3. *The newspaper's profit is lower under content bundling than under the benchmark.*

Proof. Because λ is chosen optimally by the platform, we have, for any q, $R_0(q, \lambda^*(q)) > R_0(q, 0)$. This is true in particular for $q = q^*$: $R_0(q^*, \lambda^*) > R_0(q^*, 0)$. Since $R_0(q, \lambda) + R_1(q, \lambda) = 1$, the previous inequality rewrites $R_1(q^*, \lambda^*) < R_1(q^*, 0)$. Substracting $c(q^*)$ from each side, we get $\pi_1(q^*, \lambda^*) < \pi_1(q^*, 0)$. By revealed preferences, we know that $\pi_1(q^*, 0) \leq \pi_1(\tilde{q}, 0)$, which implies that $\pi_1(\tilde{q}, 0) > \pi_1(q^*, \lambda^*)$. ■

Even though content bundling by the platform may soften competition and increase total news consumption, it cannot benefit the newspaper. The reason is that λ is chosen optimally by the platform to increase its revenue, which mechanically reduces the newspaper's revenue. The potential saving on costs is never enough to compensate this loss.

5 Extensions

This section explores three extensions to the basic model, still assuming a monopolist newspaper. In the first, we allow the platform to offer personalized content to each of its users. In the second we assume that consumers face large switching costs and cannot multihome. In the third, we allow the newspaper to remove its content from the platform, and look at a bargaining game between the two firms.

5.1 Personalized newsfeed

In the previous analysis, the platform does not have the ability to customize the mix of content it offers to each consumer. In practice however, a firm like Facebook offers different mixes to different users, leveraging the considerable amount of data it has gathered about them. We now introduce personalization to our model by assuming that the platform can observe consumers' types and can condition λ on both q and θ.

The timing is thus as follows: at $\tau = 1$, the newspaper chooses q. At $\tau = 2$ the platform observes q and θ, and chooses $\lambda(\theta, q)$. At $\tau = 3$, consumers optimally allocate their attention between the newspaper and the platform. Let $\hat{q}$ be the equilibrium quality in this case. We have the following proposition:

Proposition 4. *When the platform can personalize the newsfeed:*

1. *The platform chooses $\lambda(\theta, q) = \hat{x}(\theta, q)$.*
2. *Consumers allocate all their attention to the platform.*
3. *News quality is lower than under the benchmark: $\hat{q} < \tilde{q}$.*

Proof. Given θ and q, the platform clearly wants to offer $\lambda(\theta, q) = \hat{x}(\theta, q)$: showing less news would induce the consumer to allocate some of his attention to the newspaper, while consuming the same amount of UGC. Showing more news would not increase the time spent on the platform, but would reduce the profitability of this time. Consumers then find it optimal to allocate all their attention to the platform. To see that quality is lower than under the benchmark, note that, for a given q, the overall consumption of news is $\int \lambda(\theta, q) dF(\theta) = \int \hat{x}(\theta, q) dF(\theta) = T_1(q, 0)$ (see expression (2)). Because all the news consumption occurs on the platform, the newspaper's profit is then $(1 - \phi)T_1(q, 0) - c(q)$, instead of $T_1(q,0) - c(q)$ under the benchmark. The marginal return to investment is then lower than under the benchmark. ■

Perfect personalization allows the platform to monopolize consumers' attention, turning the newspaper purely into a content supplier to the platform. For a given quality level q, consumers' utility is maximized. However, by lowering the returns to investment compared to the benchmark, this leads to a decreased quality of news.

The assumption of perfect personalization allows us to get a clean result, but is a strong one. An alternative way to model personalization would be to assume that the platform can partition its customers into subgroups. For instance, suppose that the interval is partitioned into K intervals $I_k = [\theta_{k-1}, \theta_k)$, with $\underline{\theta} = \theta_0 < \theta_1 < ... < \theta_K = \overline{\theta}$. The platform only observes which interval I_k each consumer belongs to. Then for each interval the platform chooses a λ_k. For each interval I_k, the analysis of consumers' behaviour as well as of the platform's newsfeed design mirrors that of subsections 4.2 and 4.3. As K goes to infinity, we would converge to perfect personalization.

5.2 Single-homing consumers

An important modeling choice that we make in this paper is to assume that the only source of friction is that users cannot choose what content they consume while on the platform. In particular, we ignore another potential source of friction, namely the existence of switching costs between websites, which could deter consumers from consuming their optimal mix of content. In this subsection we test the robustness of our results by assuming that consumers incur large switching costs and are constrained to visit only one website (i.e. to singlehome).

If the platform offers a share λ of news, a consumer has a choice between consuming a mix $(x,y) = (\lambda, 1-\lambda)$ on the platform and a mix $(1,0)$ on the newspaper's website. The platform therefore attracts all the consumers of type θ such that $U(\lambda, 1-\lambda, \theta, q) \geq U(1,0,\theta,q)$. To analyze this model we make some further assumptions: we assume that θ is uniformly distributed on $[0,1]$, that $\hat{x}(\theta,q) = \alpha(\theta+q) + \beta\theta q$ (if between 0 and 1), and that the utility function takes the form $u(x,\theta,q) = v(|x - \hat{x}(\theta,q)|)$, where v is a non-increasing function of the difference between actual and desired consumption of news.[13] We refer to this as the linear-uniform model.

For a given λ, consumers who choose to use the platform are such that $|1 - \hat{x}(\theta,q)| > |\lambda - \hat{x}(\theta,q)|$, i.e. such that $\hat{x}(\theta,q) < \frac{1+\lambda}{2}$ (see Figure 2). The total time spent on the platform is then $T_0^{SH}(q,\lambda) = \frac{1+\lambda-2\alpha q}{2(\alpha+\beta q)}$. But this is precisely the time spent on the platform when consumers can multihome at no cost (see Equation (8)). Given that the platform's

[13]Given that $y = 1 - x$ in equilibrium, we drop y from the utility function.

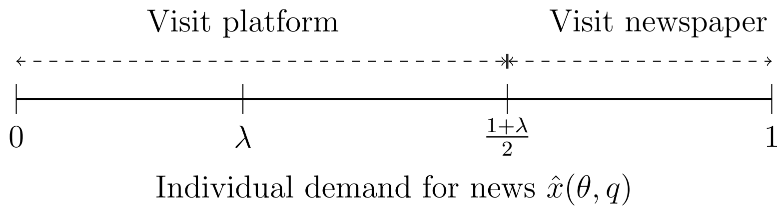

Figure 2: Consumer decision under singlehoming

profit is $(1-\lambda(1-\phi))T_0^{SH}(q,\lambda)$, the optimal λ will also be the same as in the multihoming case. A similar reasoning applies to T_1 and the optimal q.

Proposition 5. *In the linear-uniform model, the equilibrium values of λ, q and of firms' profits are the same when consumers singlehome as when they multihome.*

From Proposition 5, we can conclude that the effect of content bundling by the platform on the equilibrium choice of quality as well as on the newspaper's profit is the same as in the baseline case of multihoming (Section 4). However the implications for consumer surplus are quite different. In the baseline model, content bundling does not benefit consumers, who would have been able to consume their optimal mix by multihoming. Here on the other hand, it allows platform users to consume a positive amount of both kinds of content, whereas in the absence of content bundling they would have to choose a unique kind of content.

The equivalence between multihoming and singlehoming is of course a special feature of the linear-uniform model. Yet we believe that most of the effects that we highlighted through our baseline model would carry over: the trade-off governing the choice of λ by the platform (attracting more attention by serving less profitable content), the softening of competition due to indirect traffic revenues for the newspaper, the ambiguous sign of $\frac{\partial^2 T_1}{\partial q \partial \lambda}$, as well as the negative effect on newspaper's profit.

5.3 Newspaper opt-out

In practice, a newspaper with sufficient resources has the ability to remove its content from social platforms, or at least to make it harder for the platforms to show news. Given the adverse effect of content bundling on the newspaper's profit, here, we investigate how the ability to opt-out affects the equilibrium outcome.

Consider the following extension of our baseline model: at $\tau = 0$, the platform offers a contract of the form (F, ϕ) to the newspaper. F is a fixed payment, and ϕ is the share of the advertising revenue that the platform keeps whenever it shows some news to its consumers.[14] At $\tau = 1$ the newspaper accepts or rejects the contract, and chooses a quality q. At $\tau = 2$ the platform chooses λ if the newspaper has not opted-out. $\lambda = 0$ otherwise. At $\tau = 3$ consumers observe q and λ and optimally allocate their attention among the two websites.

Starting from $\tau = 1$, the game is the same as in our baseline model. In particular, if the newspaper rejects the contract, its profit is $\tilde{\pi}_1$. To be accepted, the contract must then deliver a payoff at least equal to $\tilde{\pi}_1$ to the newspaper. Of course the platform does not need to offer more, and so in equilibrium the newspaper is indifferent between accepting and rejecting the offer. The platform's profit is then equal to the industry profit minus $\tilde{\pi}_1$.

At $\tau = 0$, the platform therefore chooses ϕ so as to maximize the industry profit. Because the industry revenue is constant and equal to one, the profit is maximized when the cost - i.e. the quality - is minimized. One way to do so is to offer $\phi = 1$, i.e. to not share revenue with the newspaper. Indeed in that case, at $\tau = 2$, the platform finds it optimal to choose $\lambda = \hat{x}(\overline{\theta}, q)$ i.e. the highest desired news consumption for a quality q in the population, because by doing so it ensures that consumers spend all their time on its website (no consumer wants more news than what the platform offers). Unlike when $\phi < 1$, there is no cost for the platform associated with showing news, because it keeps all the revenue. The newspaper then anticipates that it will get no direct traffic no matter its quality choice, and therefore chooses to not invest in quality.

Proposition 6. *When the platform offers a contract and the newspaper can opt-out, equilibrium quality of news is minimal.*

In the next section we consider a model with multiple publishers. Among other things, this will allow us to show that newspapers' ability to opt-out is less critical in that context, because the platform can rely on a prisoner's dilemma logic and ensure newspapers' participation without having to offer fixed payments.

6 Multiple publishers

The assumption that the newspaper is the unique producer of news is clearly not innocuous. Indeed, it drives to a certain extent the "softening effect": all the news consumed through the platform come from the monopolist newspaper, who therefore has less of an incentive to compete with the platform for direct traffic. In reality, social networks may bundle many news outlets on consumers' newsfeeds.

[14] Absent the fixed payment the newspaper would always reject the offer, as per Proposition 3. This simple two-part tariffs is actually enough to maximize profit, so there is no need to study more involved schemes (e.g. contracts dependent on q).

To capture this idea, suppose that there is a continuum of symmetric newspapers on the market. Each newspaper has a mass one of traditional readers, who decide how to allocate their time between the newspaper and the social platform. When a traditional reader of newspaper i visits the platform, he is exposed to UGC and news, in proportions $1 - \lambda$ and λ. Due to the atomistic nature of the market, we assume that the news a consumer is exposed to while on the platform comes from different outlets than his usual newspaper. Newspapers are local monopolists in the sense that consumers cannot reach other newspapers directly. Advertising revenues are the same as in the baseline model. In particular, the platform captures a share ϕ of revenues when it displays news.

The timing is the following: at $\tau = 1$ newspapers simultaneously choose their quality q, at a cost $c(q)$. The quality of a newspaper is observed by the platform and by its traditional readers. At $\tau = 2$ the platform chooses the share of news it displays, λ. Consumers observe λ. At $\tau = 3$ consumers decide how to allocate their time between their usual newspaper and the platform. We look for a perfect Bayesian equilibrium where firms play a symmetric strategy and consumers form rational expectations about newspapers' quality choice (aside from their usual one, which they observe).

For tractability, we assume that $c(q) = cq^2$, that θ is uniformly distributed on $[0, 1]$ and we specify the following preferences: when a consumer reads a quantity x of news of quality q he gets a utility qx. Consuming an amount y of UGC gives the consumer a utility of $1 - \frac{(1-\theta-y)^2}{2}$. With these preferences, and assuming all newspapers have the same quality q, if the consumer could choose directly which content to consume, he would spend a share $\hat{x}(\theta, q) = \min\{\theta + q, 1\}$ of his time reading news, and $\hat{y}(\theta, q) = 1 - \hat{x}(\theta, q)$ reading UGC.

Benchmark (no content bundling): Without content bundling consumers have the choice between news from their usual newspaper and UGC from the platform. Each consumer then spends a share $\hat{x}(\theta, q)$ of his time reading news. The situation is the same as in the baseline model with a single newspaper: A newspaper's profit is

$$\tilde{\pi}_1 = \int_0^1 \hat{x}(\theta, q)d\theta - cq^2 = \int_0^{1-q} (\theta + q)\, d\theta + \int_{1-q}^1 d\theta - cq^2 = \frac{1 + 2q - (1 + 2c)q^2}{2}.$$

The equilibrium quality is then $\tilde{q} = \frac{1}{1+2c}$.

Content bundling: At $\tau = 3$, suppose that a consumer's usual newspaper has quality q and that the news quality he expects to obtain while on the platform is q^*. The consumer then chooses the time he spends on the platform, t_0, so as to maximize $(1 - t_0)q + t_0\lambda q^* + 1 - \frac{(1-\theta-(1-\lambda)t_0)^2}{2}$. The solution to this maximization problem is

$t_0(\theta, \lambda, q, q^*) = \max\{\min\{\frac{1-q-\theta-(1-q^*-\theta)\lambda}{(1-\lambda)^2}, 1\}, 0\}$. Let $\hat{\theta}_1(\lambda, q, q^*)$ be the largest solution to $t_0(\theta, \lambda, q, q^*) = 1$, and $\hat{\theta}_2(\lambda, q, q^*)$ the smallest solution to $t_0(\theta, \lambda, q, q^*) = 0$.

At $\tau = 2$, the platform chooses λ to maximize its profit. Because newspapers are atomistic, λ does not depend on a single newspaper's decision. If all newspapers except a finite number play q^*, the platform receives a total amount of attention $T_0(\lambda, q^*) = \frac{1+\lambda-2q^*}{2}$.[15] Its profit is then maximized by setting $\lambda(q^*) = \min\{q^* + \frac{\phi}{2-\phi}, 1\}$.

At $\tau = 1$, suppose that newspaper i expects all other newspapers to play q^*. Its profit writes

$$\begin{aligned}\pi_1 = &\int_{\max\{\hat{\theta}_1(\lambda,q_i,q^*),0\}}^{\min\{\hat{\theta}_2(\lambda,q_i,q^*),1\}} (1 - t_0(\theta, \lambda, q_i, q^*))\, d\theta + \int_{\min\{\hat{\theta}_2(\lambda,q_i,q^*),1\}}^{1} d\theta \\ &+ \lambda(1-\phi)\left[\int_0^{\max\{\hat{\theta}_1(\lambda,q^*,q^*),0\}} d\theta + \int_{\max\{\hat{\theta}_1(\lambda,q^*,q^*),0\}}^{\min\{\hat{\theta}_2(\lambda,q^*,q^*),1\}} t_0(\theta, \lambda, q^*, q^*) d\theta\right] - c(q). \quad (10)\end{aligned}$$

The first two integrals represent direct traffic to the newspaper, i.e. traffic from its usual readers, who actually observe the choice q_i. The first integral is traffic by the usual readers who also visit the platform, while the second corresponds to usual readers who do not. The term between brackets correspond to indirect traffic, i.e. consumers who access the newspaper through the platform: the third integral corresponds to consumers who only visit the platform, while the fourth one represents consumers who also spend time on their favorite newspaper. Importantly, these consumers do not observe the actual q chosen by the newspaper, but rather form an expectation over the quality of news they expect to receive on the platform q^*, so that indirect traffic is not sensitive to q_i.

Assuming that c is large enough, the newspaper's profit is concave. In a symmetric configuration, we have $\hat{\theta}_1(\lambda, q^*, q^*) = \lambda - q^* \geq 0$ and $\hat{\theta}_2(\lambda, q^*, q^*) = 1 - q^* < 1$. The first-order condition for a symmetric equilibrium then writes

$$\begin{aligned}\int_{\hat{\theta}_1(\lambda(q^*),q^*,q^*)}^{\hat{\theta}_2(\lambda(q^*),q^*,q^*)} -\frac{\partial t_0(\theta, \lambda(q^*), q^*, q^*)}{\partial q} d\theta - c'(q^*) = 0 \Leftrightarrow &\int_{\lambda(q^*)-q^*}^{1-q^*} \frac{d\theta}{(1-\lambda(q^*))^2} = 2cq^* \\ &\Leftrightarrow q^* = \frac{1}{2c(1-\lambda(q^*))}. \quad (11)\end{aligned}$$

Comparing q^* and $\tilde{q}$, we have the following result:

Proposition 7. *In the model with monopolistic competition with additive preferences, equilibrium quality is higher with content bundling. Newspapers' profits are lower.*

[15] $T_0(\lambda, q^*)$ is given by Equation 8 with $\alpha = 1$ and $\beta = 0$.

Remember that in the baseline model with a single newspaper and additive preferences content bundling lowers equilibrium quality. The intuition for the reversal of the result in a model with monopolistic competition is as follows. First, content bundling no longer creates a softening effect: when a consumer reduces the time he spends on newspaper i's website and increases the time he spends on the platform, newspaper i does not get any indirect revenue from that consumer. Therefore, the cost for a newspaper of losing direct traffic is the same with and without content bundling. Second, with content bundling, direct traffic to newspaper i is more sensitive to q_i under monopolistic competition than under monopoly. Indeed, under monopoly, investment in quality by the newspaper also increases the quality of news that consumers get while on the platform. Under competition on the other hand, an increase in q_i makes newspaper i more attractive without changing the value consumers expect to get from the platform. Formally, we have $-\frac{\partial t_0(\lambda,\theta,q_i,q^*)}{\partial q_i}|_{\text{competition}} = \frac{1}{(1-\lambda)^2} > \frac{1}{1-\lambda} = -\frac{\partial t_0(\lambda,\theta,q)}{\partial q}|_{\text{monopoly}}$.

An interesting difference with the model with a single newspaper has to do with the possibility for newspapers to opt-out of the platform. Even though newspapers' profit is lower with content bundling (by a similar argument as under monopoly), newspapers face a prisoner's dilemma: opting-out of the platform leads a newspaper to lose indirect traffic from consumers who would not have read it otherwise, and does not allow to increase direct traffic from its usual readers. It is therefore not a viable strategy for newspapers.

To achieve tractability, we have made two strong assumptions in particular. The first is that indirect traffic to newspaper i does not depend on q_i, i.e. that the platform grants equal prominence to newspapers irrespective of their quality. If the platform were to favor high quality newspapers, then this would increase incentives to invest even further, reinforcing Proposition 7. On the other hand, if we consider that the platform can personalize consumers' newsfeed by serving them their default newspaper, then the competitive model becomes equivalent to the model with a monopolist newspaper, where personalization leads to lower content quality. Our second assumption is that newspapers do not compete head-to-head: an increase in q_i does not reduce newspaper j's direct traffic.

7 Discussion and concluding remarks

Social networks have gained tremendous importance in the last decade, claiming a significant share of consumer attention. They have achieved such prominence by leveraging network effects *and*, more recently, by successful content bundling, whereby third party content is presented in their users' "newsfeed". This strategy, in turn, has started to fundamentally transform media production and consumption, a phenomenon of general public interest given the importance of a healthy news industry. Our main contribution is to develop a simple model of competition for attention between a social platform and newspapers, allowing us to shed light both on the strategic motives for content bundling and on its implications on the news industry.

The two main results of our analysis are that content bundling allows the platform to increase the share of attention it captures, and that newspapers' profits are reduced even though news consumption may increase. Regarding the former result, an interesting phenomenon is that content bundling allows a sort of endogenous gatekeeping phenomenon, whereby consumers read news content through the platform even though they could bypass it. This is especially true when bundles can be personalized: consumers then allocate all their attention to the platform. The key driver for the profit-reduction result is that content bundling is a strategic choice by the platform, which ensures that it is always chosen so as to increase the platform's revenues.

Regarding the quality of news content, we uncover several opposing forces that make the overall effect ambiguous. Our analysis suggests that quality is more likely to decrease in situations where newspapers have more market power and when quality mostly increases low types' demand for news. With competing newspapers, quality is more likely to go up. The result that newspapers are harmed by content bundling begs the question of why they would allow it to happen, assuming they could control it. We offer two answers. First, a monopolist newspaper could indeed threaten to opt-out to induce the platform to offer it more favorable terms. However we find that in this situation the quality of the news would unambiguously go down, as the result of a collusive outcome. Second, the ability to opt-out is much more limited when multiple newspapers compete, as the situation is akin to a prisonners' dilemma.

From a welfare perspective, content distorts low type consumers' ideal mix of content. Importantly, *for a given quality*, content bundling increases overall news consumption. While this may be desirable under certain circumstances (e.g. when there are positive externalities from news consumption), the distortion reduces consumer surplus. Moreover, if news quality decreases significantly under content bundling, news consumption may actually decrease in equilibrium.

In our baseline model content bundling does not *directly* generate any efficiencies, because consumers can choose their content mix costlessly absent bundling.[16] We show that an alternative way of modelling the situation, where switching costs prevent consumers from "mixing and matching", generates the same predictions regarding profits and quality,

[16]There could be some indirect efficiencies, if the resulting incentives to provide quality were increased for example.

suggesting a certain robustness of our results. Of course the welfare implications are then quite different.

Our analysis focused on the impact of a social network on news publishers. Our model readily applies to publishers in other content domains who also seek to be present in consumers' 'newsfeed' on social media. Beyond social networks narrowly defined, the modeling framework also seems to be applicable to a broader set of interactions between multi-sided platforms and third-party 'content' providers. For example, video distribution platforms such as Netflix, Hulu or Amazon Prime Video all bundle third party content in their offering. Here, the role of newspapers is played by movie studios or TV networks who can monetize their content independently but are attracted by the platforms' captive customer base. Music streaming platforms (such as Spotify or Deezer) also share these features. While these examples still retain the core characteristic of hybrid competition between a platform and traditional content providers, some important differences subsist, for instance the fact that the platforms charge consumers for access.

Future research exploring these related, yet different environments is probably warranted as are some of the aspect that our model omitted. For example, we did not explicitly consider network effects except for the fact that they confer some market power to the platform. Indeed, it is easy to show that without market power the platform would have to set $\lambda = 0$. Similarly, we did not consider heterogeneity across newspapers and have neglected the case where newspapers' revenue model is based on subscription instead of advertising. Including these features would be fruitful avenues for future research.

References

Alaoui, Larbi and Fabrizio Germano (2016). "Time Scarcity and the Market for News". *Working Paper,* April.

Allcott, Hant and Matthew Gentzkow (2017). "Social Media and Fake News in the 2016 Election". *Working Paper, Stanford University* January.

Alpert, I. Lukas (2015). "Publishers Warily Embed With Facebook". *The Wall Street Journal*.

Ambrus, Attila, Emilio Calvano, and Markus Reisinger (2014). "Either or Both Competition: A "Two-Sided" Theory of Advertising with Overlapping Viewerships". *Working Paper, Duke University* September.

Anderson, Simon, Oystein Foros, and Hans Jarle Kind (2016). "Competition for Advertisers and for Viewers in Media Markets". *Working Paper,*

Athey, Susan, Emilio Calvano, and Josuhua S. Gans (2017). "The Impact of Consumer Multi-homing on Advertising Markets and Media Competition". *Management Science* Forthcoming.

Athey, Susan, Markus Mobius, and Jeno Pal (2017). "The Impact of Aggregators on Internet News Consumption". *Working Paper, Stanford University* January.

Bakshy, Eytan, Solomon Messing, and Lada A Adamic (2015). "Exposure to ideologically diverse news and opinion on Facebook". *Science* 348.6239, pp. 1130–1132.

Bala, Venkatesh and Sanjeev Goyal (2000). "A Noncooperative Model of Network Formation". *Econometrica* 68.5, pp. 1181–1229.

Boik, Andre, Shane Greenstein, and Jeffrey Prince (2016). *The Empirical Economics of Online Attention*. Tech. rep. National Bureau of Economic Research.

Bradshaw, Tim (2017). "Snapchat redesigns app in new effort to challenge Facebook". *Financial Times* 29 November.

Burguet, Roberto, Ramon Caminal, and Matthew Ellman (2015). "In Google we trust?" *International Journal of Industrial Organization* 39, pp. 44–55.

Calzada, Joan and Ricard Gil (2016). "What Do News Aggregators Do? Evidence from Google News in Spain and Germany".

Chiou, Leslie and Catherine Tucker (2015). "Content Aggregation by Platforms: The Case of The News Media". *NBER* Working Paper 21404.July.

Constine, Josh (2016). "Facebook puts friends above publishers in "News Feede Values" and ranking change". *Techcrunch*.

De Cornière, Alexandre and Greg Taylor (2014). "Integration and search engine bias". *RAND Journal of Economics* 45.3, pp. 576–97.

Dellarocas, Chrysanthos, Zsolt Katona, and William Rand (2013). "Media, Aggregators, and the Link Economy: Strategic Hyperlink Formation in Content Networks". *Management Science* 59.10, pp. 2360–79.

George, Lisa M and Christiaan Hogendorn (2012). “Aggregators, search and the economics of new media institutions”. *Information Economics and Policy* 24.1, pp. 40–51.

George, Lisa Megargle and Christiaan Hogendorn (2013). “Local News Online: Aggregators, Geo-Targeting and the Market for Local News”.

Gottfried, Jeffrey and Elisa Shearer (2016). “News Use Across social Media Platforms”. *Pew Research Center,* May 26.http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/.

Jackson, Matthew O. (2010). *Social and Economic Networks.* Princeton University Press, NJ.

Jackson, Matthew O. and Asher Wolinsky (1996). “A Strategic Model of Social and Economic Networks”. *Journal of Economic Theory* 71.108, pp. 44–74.

Jeon, Doh-Shin and Nikrooz Nasr (2016). “News Aggregators and Competition among Newspapers on the Internet”. *American Economic Journal: Microeconomics* 8.4, pp. 91–114.

Luca, Michael (2015). “User-Generated Content and Social Media”. *Handbook of Media Economics* Volume 1B.Chapter 12, pp. 563–92.

Mitchell, Amy, Jeffrey Gottfried, Elisa Shearer, and Kristine Lu (2017). “How Americans Encounter, Recall and Act Upon Digital News”. *Pew Research Center,* February 9.http://www. journalism.org/2017/02/09/how-americans-encounter-recall-and-act-upon-digital-news/.

Nalebuff, Barry (2003). “Bundling, tying, and portfolio effects”. *DTI Economics Paper* 1, p. 2.

Peitz, Martin and Markus Reisinger (2015). “The Economics of Internet Media”. *Handbook of Media Economics* Volume 1A.Chapter 10, pp. 445–530.

Reuters (2016). *Reuters Institute Digital News Report.* Oxford University, UK.

Rutt, James (2011). “Aggregators and the News Industry: Charging for Access to Content”. *Net Institute* Working Paper.October, #11-19.

Thompson, Ben (2017). “Why Facebook Fails as a Platform, How Facebook Can Truly Help the Media - and Itself”. *Stratechery blog.*

Yildirim, Pinar, Esther Gal-Or, and Tansev Geylani (2013). “User-Generated Content and Bias in News Media”. *Management Science* 59.12, pp. 2655–66.

Zhang, Kaifu and Miklos Sarvary (2015). “Differentiation With User-Generated Content”. *Management Science* 61.4, pp. 898–914.