

Attention Oligopoly*

Andrea Prat
Columbia University

Tommaso Valletti
Imperial College London

December 16, 2019

Abstract

We model digital platforms as attention brokers that have proprietary information about their users' product preference and sell targeted ad space to retail product industries. Retail producers – incumbents or entrants – compete for access to this attention bottleneck. We discuss when increased concentration among attention brokers results in a tightening of the attention bottleneck, leading to higher ad prices, fewer ads being sold to entrants, and lower consumer welfare in the product industries. The welfare effect is characterized in terms of patterns of individual usage across platforms. A merger assessment that relies on aggregate platform usage alone can be highly biased.

1 Introduction

Monopoly typically reduces quantity and increases prices compared to more competitive market structures. This is why we have merger control. There is a broad consensus in competition policy when it comes to dealing with horizontal mergers, where competition authorities typically define markets in terms of substitutability in response to price changes (Motta, 2014). Enforcers seek to understand how the merger will impact on prices, quantities, and other variables that affect consumer welfare.

But how do we proceed when the merging companies are digital platforms? As with many two-sided markets, the situation is more complex. Users characteristically do not pay and hence there is virtually no price variation on the user's side to measure substitutability.¹ Platforms then monetize on the side with the advertisers: here, a more standard analysis can be run, but often the advertising

*We thank Mark Armstrong, Matt Backus, Cristina Caffarra, Emilio Calvano, Giacomo Calzolari, Allan Collard-Wexler, Matt Gentzkow, Bruno Jullien, Massimo Motta, Patrick Rey, Mike Riordan, Szymon Sacher, Sara Shahanaghi, Giancarlo Spagnolo, Jean Tirole, and audiences at Berkeley, European Commission, King's, Leicester, Madrid, NYU, Oslo, Tokyo, Toulouse and Yale for useful suggestions.

¹Some commentators actually argue that consumers do “pay with their data”; see, e.g., the speech of EC Competition Commissioner Vestager https://ec.europa.eu/commission/commissioners/2014-2019/vestager/announcements/making-data-work-us_en.

market is described as very wide (comprising the entire online advertising world). This seems to us quite far away from the engineering structure currently in place in the online world to track individual users and hyper-target them with personalized advertising.

Because of these difficulties, many authorities have fallen back on functional market definitions on the users' side. Competition authorities, for instance, cleared both the \$1 billion acquisition of Instagram by Facebook in 2012 and the \$19 billion acquisition of WhatsApp also by Facebook in 2014. The FTC investigation of the Instagram acquisition focused on camera apps and found that Facebook would not have had a monopoly in photo taking and sharing apps if the deal closed, as it faced competition from other photo sharing apps, such as Camera+ and Hipstamatic.² The merger was also cleared by the UK OFT, which observed that, at the time, Instagram did not generate revenues from advertising and therefore was not an important competitor of Facebook for advertisers.³ For both authorities, the starting point of the investigation was the market for camera apps. Analogously, the WhatsApp merger assessment started from the analysis of the market for messaging apps.⁴

The alleged failure of competition authorities to police tech firms is one of the phenomena that inspired a group of law scholars and policy analysts, often referred to as the New Brandeis School, to question one of the core principles of modern antitrust: the consumer welfare criterion. After discussing the current state of antitrust enforcement, Lynn (2013) concludes: "We should cease to use the 'Consumer Welfare' frame to guide enforcement of our antimonopoly policies. We should update and restore the 'competition within open markets' frame put into place beginning a century ago, to protect our liberties, our democracy, our communities, our sovereignty, and our stability." Khan and Vaheesan (2016) and Wu (2019) also argue that it is time to replace the consumer welfare standard with a new paradigm.

The present paper makes a contribution to this debate by remaining within the consumer welfare paradigm and providing a quantifiable argument for blocking mergers between free-to-use digital platforms. While the New Brandeisians are right that there has been under-enforcement by antitrust authorities in digital markets,⁵ this does not automatically imply we should abandon the consumer welfare paradigm. Perhaps the correct response is to strengthen and extend it.

In this paper, we model digital platforms as *attention brokers* (Wu, 2019). Attention brokers attract users to their platform and induce them to spend time on it and use it for practical and social purposes. User attention is valuable for two reasons. First, usage data provides the attention broker

²<https://www.ftc.gov/news-events/press-releases/2012/08/ftc-closes-its-investigation-facebooks-proposed-acquisition>

³<https://www.gov.uk/cma-cases/facebook-instagram-inc>

⁴http://ec.europa.eu/competition/mergers/cases/decisions/m7217_20141003_20310_3962132_EN.pdf

⁵A report, written in March 2019 for the UK Treasury, notes that in the past 5 years alone companies such as Amazon, Apple, Facebook, Google, and Microsoft combined have made 250 acquisitions. Virtually none of these mergers were notified voluntarily, or investigated by the relevant authorities. See <https://www.gov.uk/government/publications/unlocking-digital-competition-report-of-the-digital-competition-expert-panel>

with proprietary data on the search and communication activity of users. Machine learning techniques may then be applied to infer real-time consumption preferences of individual users (Agrawal et al., 2018). The platform may learn that the user is currently interested in a particular product, say a new refrigerator, or a service, say a plumber (Milgrom and Tadelis, 2019). Second, the attention broker can now sell targeted advertising space to firms that supply the product the user is interested in (e.g., refrigerator makers, or local plumbers, which we refer to as the *retail product industry*). These ads are very valuable because they are directed to consumers who are interested precisely in that product and because these consumers' attention is already captured by the platform. This is the value proposition of attention brokers and it explains why companies such as Facebook or Google are so valuable.

What is an attention broker in practice? Attention brokers in our model are defined by their ability to: (i) obtain information about the preferences of *individual* users; and (ii) target ads to *individual* users. The set of attention brokers is changing over time with changes in technology and behavior. Currently, platforms like Facebook or Google fit to some extent our definition. Instead, more traditional media like TV or newspapers are still mainly unable to achieve (i) and (ii), though they might in the future, in which case our results will apply to them as well.⁶

How does market power among attention brokers affect consumers, even in zero-price platform markets? This paper argues that increased concentration among attention brokers can lead to reduced entry – and hence higher prices and less product variety – in retail product industries.

In a nutshell, the argument runs as follows. A monopolistic attention broker has an incentive to create an *attention bottleneck* by reducing the supply of targeted advertising. If an attention broker reduces the number of ads it sells, it will reduce the number of retail firms that have access to consumers, thus increasing their market power. This bottleneck strategy can generate higher total profits for the retail industry that are partly captured by the platform through higher total ad revenue. However, under standard conditions, this supply reduction hurts consumers who face less choice and higher prices. In this sense, indeed in our model consumers end up paying with their “attention data”.

This access reduction strategy does not always work: it depends on how concentrated user attention is. Suppose the retail product industry is composed of well-known larger firms and lesser-known smaller entrants. If consumers use many platforms, entrants enjoy a strategic advantage: they can inform the user about their product in more than one way and they only pay for the ads they buy. Individual platforms face a temptation to increase the supply of ads, some of which will be bought by entrants. When that happens, the retail industry will be competitive and both product prices and ad prices will be low. Instead, if the attention of the consumer is controlled by a limited number of platforms, an exclusionary strategy is easier to carry out: ad supply is low and captured mostly by incumbents. Entry is pre-empted and both product prices and ad prices are high. A merger between platforms can thus increase market power in the retail industry, to the detriment of consumers. Platform

⁶Our model also allows for the presence of traditional mass media that sell non-targeted ads alongside attention brokers.

concentration implies a reduction in the supply of ads in order to create, and extract, pre-emption rents in the retail industry.

A corollary of this argument is that the right measure of platform concentration is at the level of each *individual* user. In a world where platforms obtain personal information and can tailor ads to user, what matters is the number of platforms retail product firms can use to reach a particular user. Thus, a meaningful concentration index for attention brokers cannot be built out of aggregate market share. We use a numerical example to show that such a measure is a very inaccurate reflection of the welfare cost incurred by consumers.

The paper is structured as follows. While our main point can be made in a general setting with an arbitrary number of incumbents, entrants, and ads, and with ads that can be imperfectly informative, we prefer to begin with the most parsimonious model, which will allow us to illustrate the intuition behind the results. The general case is left for the last section.

Section 2 introduces the model with three layers of actors: non-strategic consumers, attention brokers (digital platforms), and retail producers. Producers sell goods to consumers and can advertise their production through digital platforms. There are multiple digital platforms. Consumers differ in their platform usage and may multi-home among platforms (usage is not affected by which ads they are shown – that is what makes them non-strategic). Platforms use competitive selling mechanisms to allocate ads to firms, where firms can buy non-targeted ads from traditional mass-media or targeted ads from digital platforms. The latter have two advantages: they have better knowledge about product preferences of their users and they can sell tailored ads targeted to users that want a particular product. Producers may be incumbents or entrants. The products of incumbents are more familiar to consumers than those of entrants. Any type of advertising can help entrants close this informational gap. We assume that a consumer’s surplus is higher if he is aware of the entrant’s product, but this reduces the incumbent’s revenue. Every platform used by a particular individual sells the right to run targeted ads to that individual through a competitive selling mechanism.⁷

In this environment, each consumer may be seen as an individual market. Section 3 begins our analysis with a characterization of the ad selling equilibrium in each one of these markets (Proposition 1). The probability that the consumer becomes aware of the entrant’s product increases with the number of independently owned platforms utilized by that consumer. Aggregating across consumers, we express overall consumer welfare on the basis of platform usage patterns (Proposition 4). Under some simplifying assumptions, consumer welfare is a decreasing function of a simple concentration index:

$$\sum_J \frac{m_J}{n_J},$$

⁷We thus take the view that advertising is more informative for entrants’ products than for incumbents’ (for a review of the economics literature on advertising, see Bagwell, 2007).

where m_J is the share of consumers who use a certain subset of platforms J and n_J is the number of independently-owned platforms in that subset. Welfare is therefore higher if consumers use multiple platforms because it is more expensive for the incumbent to keep consumers uninformed about the entrant's product.

In Section 5, we then characterize the welfare effect of a merger between two existing platforms (Proposition 6). For every consumer at any point in time, the relevant market is defined by the set of attention brokers that: (i) know that consumer's current preferences; and (ii) are able to target ads to that consumer. We therefore show that knowing platform usage shares is not sufficient to compute the effect of a merger: one must also know how platform usage overlap is distributed across consumers. A numerical example highlights that a regulator who tries to predict welfare effects on the basis of usage shares only can face a large level of uncertainty (of the order of 50% of total entrant-related surplus). Instead we employ US usage data of three major social media platforms, Facebook, Instagram, and Twitter, to illustrate that it is possible to build a concentration index based on our measure.

Section 6 discusses the difference between targeted platform ads and non-targeted mass media ads. We show that the latter are more likely to be employed by larger, mainstream retail industries. Instead, targeted platform ads are more likely to be chosen by niche industries – therefore the merger effects of digital platforms are also more likely to be felt in such niche industries. Section 7 conducts robustness checks considering alternative ad selling specifications and a model with multiple incumbents, multiple entrants, multiple ads per platform, and imperfectly informative ads. Finally, Section 8 concludes.

Our paper identifies the attention bottleneck effect as a theoretical possibility, but is it just an intellectual curiosity? While only a thorough empirical analysis can answer this question, there are two observations to make. First, the potential market is large: global spending on internet advertising is predicted to rise above \$350 billion in 2021 and will constitute over half of total ad spending. Digital advertising is also highly concentrated, with two firms – Google and Facebook – commanding nearly 60% of the market in the US. These two firms are also capable of deploying hyper-targeted ads.

Second, our paper may explain a somewhat puzzling phenomenon. Companies spend large sums for brand keyword ads: a famous brand (e.g., Coca Cola) advertises for keywords containing their brand, despite organic links would already be likely to appear on top of the page (e.g., www.coca-cola.com). Blake et al. (2015) show that those ads have a precisely estimated zero benefit. Why are companies seemingly wasting money? Our model offers a rationalization: in the pre-emption equilibrium, the incumbent is buying ads not to bring in more business but to avoid a potential entrant from stealing business. Indeed, Simonov et al. (2018) show that brand keyword advertising produces a small number of additional clicks (1-4%) but avoids the large loss of clicks (18-42%) that the famous brand would face if a competitor had gotten hold of that ad.⁸

⁸See also Athey and Nekipelov (2014) and De Carolis and Rovigatti (2019). Our model would imply that this bidding

Naturally, there are a number of other merger effects – positive and negative – we abstract from. The action in our model comes only from the possibility of creating an attention bottleneck. Platforms could be on very different users’ markets (e.g., a social network such as Facebook and an instant messaging service such as WhatsApp) and still related to the same user’s attention. Of course, if the platforms were competing directly against each other (currently or prospectively) there would be additional effects to be considered. Also, a merger in our model can only affect the amount of attention spread across platforms, not the effectiveness of an ad. Again, in reality the ability to combine datasets across platforms might generate additional insights that would change the ability to conduct targeted advertising that would also need to be assessed in a real-life case. Still, the attention bottleneck effect we identify will be present in those more complex environments too.

1.1 Literature

Attention markets are central to recent cases and in policy works (Evans, 2017; Wu, 2019). Wu (2019) in particular has introduced the idea that social media companies are first and foremost attention brokers: they capture the attention of their users and sell it to advertisers. To our knowledge, no formalizations of the product market implications of mergers between attention brokers are available.

A main building block in our model is related to pre-emption rents in the retail product industry (which platforms try to possibly extract via their merger). Pre-emption refers to the opportunities for firms with market power to maintain their monopoly power. There is an extensive literature in this respect. Gilbert and Newbery (1982) study a monopolist’s incentive to maintain its monopoly power by patenting new technologies before potential competitors, and this activity can lead to patents that are neither used nor licensed to others (“sleeping patents”). Other examples include spatial location models with pre-emption (Eaton and Lipsey, 1979) and brand proliferation (Schmalensee, 1978). The common idea is that a monopolist will spend more on some scarce resource (e.g., R&D leading to a patent, or enter a location that thus becomes unavailable for a rival) to preserve monopoly rents.⁹ This happens when entry results in a reduction of total profits below the joint-profit maximizing level, which is quite a general result (Gilbert and Newbery, 1982). Asker and Bar-Isaac (2014) analyze a dynamic game where an incumbent manufacturer is currently selling to multiple retailers and can use various forms of vertical restraints such as resale price maintenance and royalty rebates. The incumbent is trying to prevent entry of another firm, which would lead to lower profits. In their paper the incumbent can be viewed as “bribing” retailers to keep industry monopoly – something that happens in our paper too, has competitive implications in retail product markets. AdWords (run by Google) is also an environment where bids are hyper-targeted that would fit our setting.

⁹In our model, scarcity is created on users’ attention, which can then be monetized by incumbent product advertisers. There is a long-standing debate in economics as to whether preventing competitors from presenting informative advertising results in higher prices for advertised products (and harms consumers); see Kaldor (1950), and Telser (1964).

albeit in a different setting.

Our paper is more broadly related to the literature on two-sided markets initiated by the seminal works of Rochet and Tirole (2003), Anderson and Coate (2005) and Armstrong (2006).¹⁰ The extant literature looks at the externalities between the two sides. When one side of the market is represented by advertisers, the main research focus is the link between the advertising price and the nuisance that ads generate on audiences. Advertisements are placed by monopoly producers that want to reach as many customers as possible. On the other side of the market, audiences are typically not overlapping between platforms (in the literature jargon, consumers “single-home”, while advertisers “multi-home”). The literature also normally deals with competition between duopoly platforms.

The single-homing assumption for consumers means that there is automatically an “attention bottleneck” built in to these models. In our work, we simplify one side of the market, as we take the distribution of consumers on platforms as given.¹¹ We also further simplify as we do not consider externalities between advertisers and consumers. Our focus is instead on the competitive externality between firms that can obtain ads from online platforms where consumers can multi-home.

At a theoretical level, our sequential ad auctions can be seen as games played through agents (Prat and Rustichini, 2003). Multiple principals (the producers in this paper) offer conditional monetary transfers (auction bids here) to multiple agents (the platforms) in an attempt to influence their decisions (the allocation of ad spaces). The key result of that literature is that the set of subgame perfect equilibria often does not contain an equilibrium that maximizes the surplus of all principals and agents. In the case at hand, the surplus-maximizing outcome corresponds to incumbent monopoly, so the presence of multiple platforms may preclude that outcome. This is bad news for the total profit of producers and platforms, but, if we assume that incumbent monopoly is not in the interest of consumers, then this result is good news for consumer surplus. This result is driven by the presence of multiple agents: If there is only one agent, the menu cost literature has shown that there always exists an equilibrium that maximizes the surplus of all principals and the agents (Bernheim and Whinston, 1986). While there are a number of key technical differences, the present result builds on an intuition that is similar to Prat and Rustichini (2003): strong two-sided competition among principals and agents makes it hard to reach the collusive equilibrium. Obviously, the current paper is quite different: while it uses a game played through agents as one of its building blocks, the set-up is much richer because it includes multiple consumers and multiple industries.¹² The present model has also a connection to gatekeeping

¹⁰Caillaud and Jullien (2003) study a related model of competition between information brokers and focus on indirect information externalities. See Rysman (2004) for an empirical application.

¹¹In the two-sided markets literature, because a platform’s profits from advertising are roughly proportional to “eye-balls”, platforms then have a strong incentive to subsidise access by consumers to the platform. This is natural justification for the free access consumers have to these platforms, which is a starting point of our modelling ingredients.

¹²There is also a technical difference. Prat and Rustichini (2003) use simultaneous first-price auctions, while this paper uses sequential second-price auctions. Moreover, Prat and Rustichini consider a much more general set-up but only

in politics: Diermeier and Myerson (1999) model a legislative process where lobbyists bribe legislators to pass a law and asks how different institutional arrangements affect the total amount of bribes.

The present paper is loosely related to the vast literature on foreclosure (see, e.g., Rey and Tirole, 2007, and Whinston, 2006). The similarity is that our set-up can be interpreted as a situation with a downstream industry (the producers) who require the input supplied by an upstream industry (the platforms) in order to access final consumers.¹³ However, our focus is quite different. The foreclosure literature mostly studies the effect of different – and potentially quite complex – vertical contracting environments, while leaving the interaction between the downstream industry and consumers in a reduced form. For instance, a classic question is whether exclusivity clauses should be outlawed, or whether vertical integration can lead to foreclosure. Because of the complexities of contracting under externalities, this theoretical literature has often focused on special market structures such as monopoly either upstream or downstream, or on competing vertical chains, where each supplier deals only with distinct retailers. Instead, we focus on a stable and simple vertical contracting problem, online ad auctions, while the focus of the analysis is on the heterogeneity of consumers with respect to their relationship to the downstream industry: namely all the action arises from the fact that different consumers use different sets of platforms. This simplified contract setting also allows us to deal quite naturally with multiple online platforms upstream and competing retailers downstream.

Eliaz and Spiegler (2019) also model targeted advertising on digital platforms. They characterize the optimal ad display rule and advertising fee in a situation where advertisers have private information about the quality of their match with consumer types.

We consider multiple platforms in order to conduct a meaningful analysis of platform mergers. Overlaps of consumers among platforms are crucial for our findings in the merger analysis. As said above, this is assumed away by the extant literature with single-homing consumers, with a few exceptions. Ambrus et al. (2016) and Anderson et al. (2018) consider overlapping viewerships, focusing again on externalities (positive from audiences to advertisers, and negative the other way around) and not on the product market implication of selling ads to those consumers.¹⁴ Audience overlap is central to Prat's (2018) media power definition, but that model is concerned exclusively with the influence of news sources on the electoral process, while here we are interested in the influence on markets, not politics.

Media mergers have been analyzed, e.g., by Anderson and McLaren (2012) who consider their impact on media bias. There also exists empirical specific work on mergers in two-sided markets (e.g., Chandra and Collard-Wexler, 2009, on Canadian newspapers; Jeziorski, 2014, on US radio stations).

prove an existence result, while here restricting attention to a more specialized model will allow us to characterize the equilibrium set.

¹³Strictly speaking, in our framework the input is only needed by the entrant producer. The incumbent producer can still hoard it for the only purpose of preventing entry.

¹⁴See also Athey et al. (2018) for a model of advertising where consumers multi-home across media outlets. Competition for advertising slots sold by a monopoly platform is analyzed by de Corniere and De Nijs (2016).

We are not aware of empirical studies on mergers between online platforms.

2 Model

To keep the exposition simple, the core of the paper uses a “toy model” with just one incumbent and one entrant, who can each show at most one ad. Moreover ads are assumed to be perfectly informative. None of these assumptions is crucial for our main result: Section 7 will remove all of them and consider a general setting with multiple incumbents and entrants, multiple ads per platform, and imperfectly informative ads. However, for now we explore the intuition behind the result in the simplest possible setting.

There are a continuum of potential buyers of total mass 1 and a set K of retail industries, each of which produces a product k . Every consumer is interested in buying exactly one of the products. Let γ_k be the share of consumers that want product k . Without loss of generality assume that k is ordered so that γ_k is nonincreasing in k .

There are two types of advertising channels. First, there are mass media outlets that sell a large number of non-targeted ads seen by all consumers. Each mass media ad costs a flat price denoted by a . Second, there is also a set M of digital platforms (henceforth platforms). Let m_J denote the mass of users that use the subset J of platforms (and only those). Obviously, $\sum_J m_J = 1$. These platforms differ from mass media because they know what product k each of their users wants and they can run targeted ads that differ from user to user.

Every retail industry k is composed of an incumbent and an entrant. Buyers interested in product k are familiar with the incumbent’s product but they are unaware of the entrant’s product. The only way they can learn about the entrant’s product is if they see an ad about it, either from mass media or from platforms. By seeing an ad, a consumer becomes perfectly informed about the advertised product. For each individual buyer, if the buyer is aware of the incumbent only, the buyer’s expected payoff (gross payoff minus price) is u_1 , the incumbent’s expected payoff (price minus production cost) from that buyer is π_1 , and the entrant’s expected payoff is 0. If a buyer becomes aware of the entrant’s product, the buyer’s expected payoff is u_2 , the incumbent’s payoff is π_2 , the entrant’s expected payoff is π_E . Assume that entry benefits consumers, $u_1 < u_2$, and that $u_1 + \pi_1 < u_2 + \pi_2 + \pi_E$. We will analyze two cases: when entry increases total profit ($\pi_1 < \pi_2 + \pi_E$) and when entry decreases total profit ($\pi_1 > \pi_2 + \pi_E$). Note that in both cases entry is efficient from the perspective of both consumer and total welfare maximization.^{15,16}

¹⁵Under the utilitarian definition of social welfare, in the incumbent monopoly case total welfare is $u_1 + \pi_1$, while total welfare in the duopoly case is $u_2 + \pi_2 + \pi_E$.

¹⁶All payoffs are expressed per consumer. Aggregate values are obtained by integrating over consumers. For instance, the total profit created by a mass s of consumers that generate individual profit π is $s\pi$, etc.

Timing is as follows. First, consumer preferences are randomly determined. Second, every firm chooses whether to buy a non-targeted mass media ad at the flat rate. Then, each platform runs a second-price auction for each one of its users. So, for each user in the subset J , there are $\#J$ auctions. We assume that for each user the order of those auctions is drawn randomly and that there is complete information about the whole auction sequence. The entrant and the incumbent participate in every auction and the winner gets to display an ad to that user. The auctions are run independently across users and goods.

2.1 Discussion of Modeling Choices

The model above is highly stylized. It is an attempt to capture in the most parsimonious way the pre-emption phenomenon we are trying to describe.

Some assumptions – second-price auction, one incumbent, one platform, one ad per platform, perfectly informative ads – are made for analytical convenience as they allow for a relatively simple closed-form characterization of equilibrium. Section 7 will relax each one of them and show how our pre-emption condition can be extended.

We discuss below the major explicit and implicit assumptions that we have made:

- *Knowledge and Targetability.* The model assumes that there are attention brokers that offer perfectly targetable ads and have perfect knowledge about their users. This assumption is meant to capture an idealized situation where digital platforms collect large amounts of browsing data and process it through machine learning, perhaps the limit point of the long-term trajectory of technological progress, as described by Larry Page, one of the founders of Google, in 2000: “Artificial intelligence would be the ultimate version of Google. The ultimate search engine that would understand everything on the web. It would understand exactly what you wanted, and it would give you the right thing. We’re nowhere near doing that now. However, we can get incrementally closer to that, and that is basically what we work on.”¹⁷

This assumption brings two major advantages. First, it makes the model tractable: without it, we would have asymmetric information between platforms and/or sellers which would make the analysis laborious and opaque. Second, we see some value in characterizing Larry Page’s asymptotic scenario, especially in a fast changing world. Future research could revisit our set-up and extend the analysis to a world where platforms have asymmetric information.

Since in our model each user can receive targetable ads, and this allows the retail product firms to target that consumer, we also implicitly assume that product markets are defined at the same individual level (e.g., product firms can engage in first-degree price discrimination). This corresponds, for instance, to the price discrimination case analyzed by Armstrong and Vickers (2019).

¹⁷<http://www.artificialbrains.com/google>

- *Knowledge about the Set of Platforms.* The results in Sections 3 and 4 do not require the platforms to know whether a specific consumer is using other platforms. This is because the mechanism is given: a second-price auction. Where that knowledge matters is in Section 5 when we discuss the effect of a merger. In order to decide whether to sell ads as a bundle or not, a firm should know how many other platforms are available. We do not know whether that is true in practice. Obviously, an imperfect but accurate measure would still be sufficient to make a decision that is optimal for most consumers.
- *One Incumbent, One Entrant.* The model can also be extended beyond the assumption that there is one incumbent and one entrant; see Section 7.2. The pre-emption condition is unchanged, but it is more challenging to provide a closed-form characterization of the competitive equilibrium because an equilibrium in pure strategies may fail to exist and there is no general characterization of mixed-strategy equilibria outside specific cases (Szentes and Rosenthal, 2003).
- *One Ad per Platform.* We assume that each platform has only one ad to sell. Section 7.2 analyzes this extension.
- *Full Informativeness of Ads.* Seeing one ad – whether from a generalist channel or a digital platform – is sufficient to become aware of the entrant’s product. One possible extension of this paper is to assume some degree of inattention on the part of consumers, so that exposure to an ad does not guarantee knowledge. This would create scope for the platforms to put more than one ad for sale. Section 7.2 relaxes this assumption.
- *Second-Price Auctions.* We assume that each platform sells its ad through a second-price auction. In practice, platforms employ n -price mechanisms, like Google’s AdWords. In Section 7.1, we show that our main result is robust to alternative selling mechanisms used. We consider a situation where, instead of using second-price auctions, platforms can make any take-it-or-leave-it offer.
- *Random Order of Auctions.* We assume that platforms hold their ad auctions sequentially. If auctions were simultaneous, we would face the kind of non-existence of pure-strategy equilibria discussed in Szentes and Rosenthal (2003). Given that auctions are held sequentially, our equilibrium characterization holds for any platform order. We choose a symmetric random order as the most agnostic option.
- *Set of Digital Platforms.* As mentioned in the introduction, the set of digital platforms that serve a particular individual includes those platforms that: (i) have information about the preferences of that individual; and (ii) are able to target ads to that individual. This set arguably includes social media like Facebook and search engines like Google if the individual uses them. It does not probably include other digital media, like online newspapers or stream services, because either

(i) or (ii) or both fail. At the current state of technology, those outlets are best represented as more traditional “mass” media, like television and newspapers who sell ads to a bulk of generic “eyeballs” who are all shown the same ad. Of course, if technology were to change in the future, the set of digital platforms may change too and our model would apply to whatever the relevant sets are.

We discuss (i) and (ii) immediately below.

- *Difference between Targeted and Non-Targeted Ads.* The model posits a stark difference between media platforms that have perfectly targetable ads and media platforms that have non-targeted ads. The truth is that all platforms have some information about their users and some leeway to target the ads: even newspapers have some sense of who is more likely to read a particular section and target ads accordingly (e.g., hotel ads in the travel section). However, digital platforms have a threefold advantage: they have access to user behavior data, which provides them with accurate information about user preferences, they can customize ads individually, and they can sell them individually. Currently, only some social media platforms and search engines are able to achieve this triple advantage. In the future, the set of platforms with this capability may increase. Our paper applies to whichever set of platforms has this targeting capability.
- *No Information Synergies between Platforms.* Although this restriction is an immediate consequence of the knowledge and targetability assumption above, it should be highlighted separately. Our idealized platforms already know everything about their users. Hence, a merger between two or more platforms cannot increase their knowledge base or enhance their ability to target individual customers. Future empirical research should try to assess how less than omniscient platforms might make information gains or reduce information processing costs if they merge.
- *Non-strategic Consumers.* In our model social platform usage patterns are unaffected by advertising. Our consumers choose to use a certain set of platforms for pure consumption value, without taking into account that they may receive useful information about products or they may be charged different prices depending on the set of platforms they utilize. While a fully rational consumer should weigh these factors before checking his or her Facebook page, introspection suggests that myopia is not an entirely unrealistic assumption. Future research could add a platform usage selection stage to the present model.
- *Similar Industries.* We assume that u_1 , u_2 , π_1 , and π_2 are the same across different industries. The model could be easily extended to industry-specific values.
- *Reduced-Form Payoffs.* The payoffs of firms, platforms and consumers are expressed in reduced form as u_1 , u_2 , π_1 , π_2 , and π_E . Recall that these payoffs are expressed per consumer. They can be microfounded in a number of ways – under the knowledge and targetability assumption above.

For instance, assume the consumer derives utility V_I from the incumbent's product and utility $V_E \in (V_I, 2V_I)$ from the entrant's product. If the consumer is only aware of the incumbent, the incumbent will charge a price V_I and the consumer will buy from the incumbent. If the consumer is also aware of the entrant, he will buy from the entrant at price $V_E - V_I$. This yields $u_1 = 0$, $u_2 = V_I$, $\pi_1 = V_I$, $\pi_2 = 0$, and $\pi_E = V_E - V_I$, which satisfies all the assumptions above. In particular, entry increases consumer welfare and total welfare but decreases industry profits. While this is the simplest microfoundation, one could also allow the consumer to buy multiple items of both products or endow him with stochastic preferences.¹⁸

- *Incumbent/Entrant Product Familiarity.* The model can be extended beyond our extreme assumption that all consumers are aware of the incumbent's product and unaware of the entrant's product. There may be segments of consumers that are unaware of the incumbent's product, aware of the entrant's product, or both.
- *Focus on Consumer Welfare rather than Total Welfare.* In line with standard competition policy practice, our key metric will be consumer welfare. One could instead focus on total welfare, which also includes platform and producer profits. A set of standard assumptions would guarantee that consumer welfare and total welfare go in the same direction.

3 Consumer Segment Analysis

In this section and the following two, we focus exclusively on targeted platform ads and abstract from generalist mass media ads (assuming that their price a is so high that none of the entrants buy them). We will consider that possibility in Section 6.

In this section, we analyze the sale of targeted ads in the last stage of the game. We focus on a particular industry k (hence, for notational simplicity, we shall drop the index k in this section).

We begin by characterizing the equilibrium of an auction in segment J . Let $n_J = \#J$.

Let us also introduce four equilibrium payoff variables:

- R : the expected platform revenue in a segment with n_J platforms, namely the expected revenue of a platform before the auction order is decided. As the platforms are symmetric, this is given by the total revenue divided by n_J .
- Π_I : the net payoff of the incumbent in a segment with n_J platforms, namely her expected gross payoff (π_1 or π_2 , depending on the equilibrium) minus whatever she pays to the platforms.
- Π_E : the net payoff of the entrant in a segment with n_J platforms, namely her expected gross payoff (π_E or 0, depending on the equilibrium) minus whatever she pays to the platforms.

¹⁸This microfoundation is similar to the model in the discrimination case considered by Armstrong and Vickers (2019).

- U : the payoff to consumers, either u_1 or u_2 depending on the equilibrium.

The relevant equilibrium concept is subgame perfection with the refinement that the two firms do not offer bids that constitute weakly dominated strategies in a static sense. Namely, if for a given auction V is a player's continuation payoff from winning the auction and v is the player's continuation player from losing the auction, the player cannot bid strictly more than $V - v$.

Proposition 1 *Let*

$$\bar{n} = \frac{\pi_1 - \pi_2}{\pi_E}. \quad (1)$$

(a) *If $n_J < \bar{n}$, we have incumbent **monopoly**: all ads are sold to the monopolist and every platform has a revenue π_E .*

(b) *If $\bar{n} + 1 > n_J > \bar{n}$, we have **entry with positive profit**: at least one ad is sold to the entrant and the expected platform revenue is strictly positive.*

(c) *If $n_J > \bar{n} + 1$, we have **entry with zero profit**: at least one ad is sold to the entrant and the expected platform revenue is zero.¹⁹*

The payoffs in the three cases are reported in the table below

		(a)	(b)	(c)
<i>Platform</i>	R	π_E	$\frac{\pi_1 - \pi_2 - (n_J - 1)\pi_E}{n_J}$	0
<i>Incumbent</i>	Π_I	$\pi_1 - n_J\pi_E$	π_2	π_2
<i>Entrant</i>	Π_E	0	$n_J\pi_E - (\pi_1 - \pi_2)$	π_E
<i>Consumer</i>	U	u_1	u_2	u_2

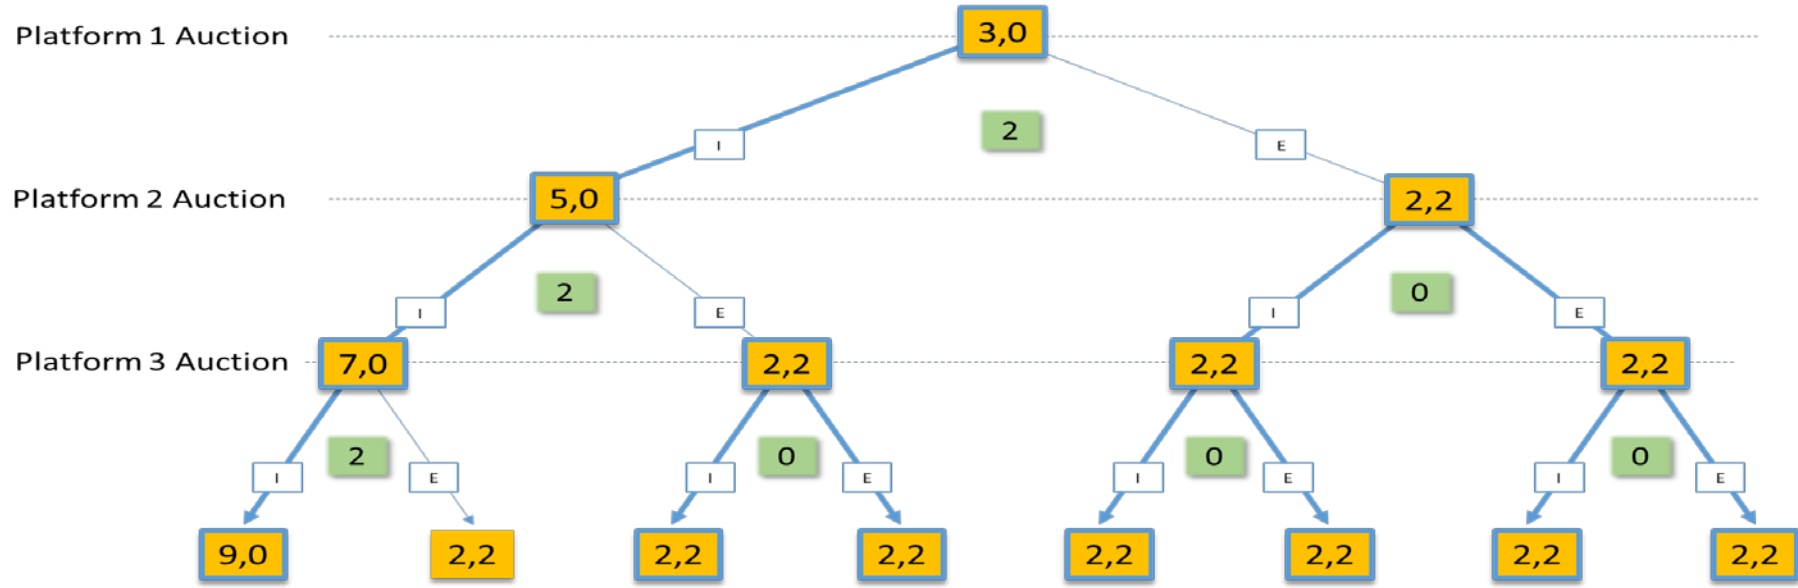
Proposition 1 is best understood through the example in Figure 1. Assume there are three independently-owned platforms. The profit of a successful entrant is $\pi_E = 2$. The profit of an incumbent who ends up in duopoly is $\pi_2 = 2$. The profit of an incumbent who manages to defend her monopoly varies in the three cases: $\pi_1 = 9$ in Example 1, $\pi_1 = 7$ in Example 2, and $\pi_1 = 5$ in Example 3. The value of π_1 is the only difference between the three plots. Each plot represents a game tree.

Each node corresponds to an auction: there are three levels, one for each auction. The two numbers to the left of each node represent the equilibrium continuation payoff starting at that node (for instance, (3,0) to the left of the top node in case (a) means that the continuation payoff at the beginning of the game – thus of the whole game – is 3 for the incumbent and 0 for the entrant; the pair (5,0) just below and left of (3,0) is the continuation payoff after the first auction has been won by the incumbent). The numbers to the left of the end node are the primitive payoffs as described above. The numbers below each node represent the equilibrium bid for that auction. On-equilibrium branches are

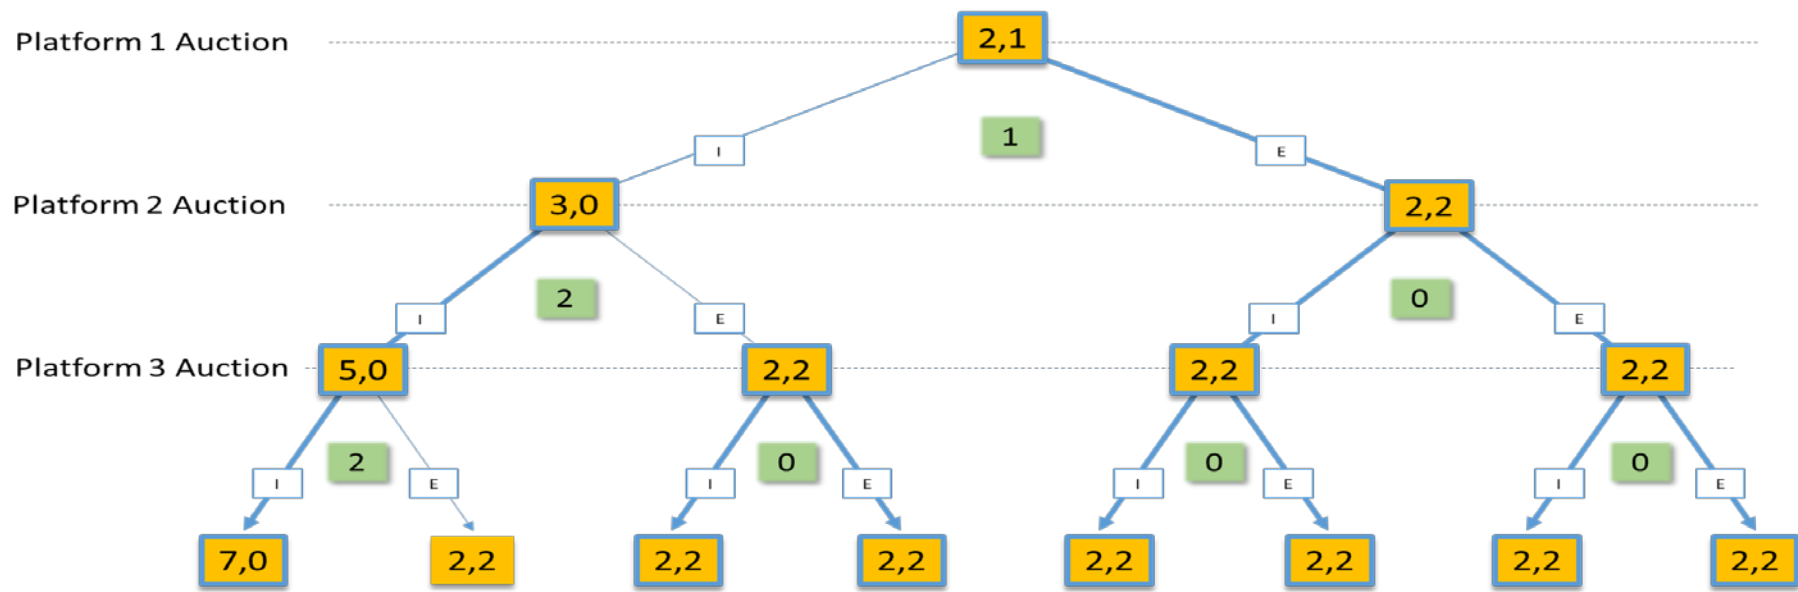
¹⁹To simplify exposition, we disregard the nongeneric case where the conditions in (a), (b), and (c) hold as equalities. When $n_J = \bar{n}$, both (a) and (b) are possible and they generate the same payoff for the platforms but not the same consumer surplus. When $n_J = \bar{n} + 1$, (b) and (c) coincide.

Figure 1

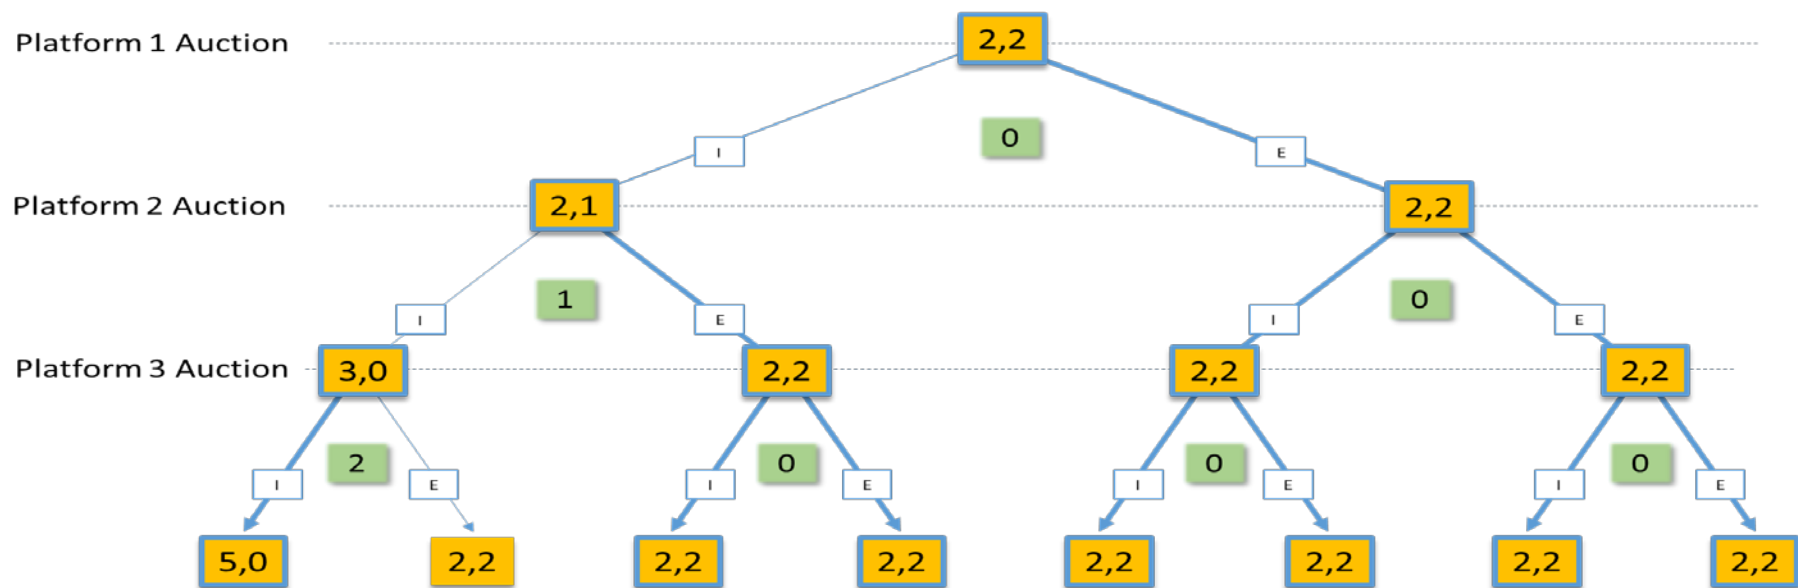
Example 1: $\pi_1 = 9, \pi_2 = \pi_E = 2 \rightarrow$ Incumbent Monopoly



Example 2: $\pi_1 = 7, \pi_2 = \pi_E = 2 \rightarrow$ Entry with Positive Platform Profit



Example 3: $\pi_1 = 5, \pi_2 = \pi_E = 2 \rightarrow$ Entry with Zero Platform Profit



- 3,0 Incumbent/Entrant continuation payoffs (net of payments in current and upcoming auctions)
- 2 Equilibrium bid in auction
- I E Auction winner (Incumbent/Entrant)
- / Off/on equilibrium path branch

denoted in bold. Note that once the incumbent has lost one auction, both players are indifferent with regards to the remaining ads, so all branches are possible equilibrium paths.

Example 1 of Figure 1 corresponds to case (a) of Proposition 1. The incumbent (she) has a lot to gain from defending her monopoly. In equilibrium, she offers a bid of 2 in all three auctions. This is exactly enough to ward off the attack of the entrant (he), who is willing to offer at most 2 because he stands to gain $\pi_E = 2$. Thus, the incumbent wins all three auctions and the total revenue of all platforms is $n\pi_E = 6$.

Example 2 represents case (b), where the incumbent has a little less to gain from defending her monopoly. In fact, that gain is 5, which is less than the cost $n\pi_E = 6$ of defending her monopoly. In equilibrium, the entrant gets an ad. He bids 1 at the first auction and he wins this ad. The bid of 1 exactly offsets the amount the incumbent is willing to pay to win the first auction as she anticipates having to fork out a total of 4 in the following two auctions. The total revenue of the three platforms collapses from (a) to (b): it goes from 6 to 1. This is mainly a reflection of two facts: monopoly is worth more than duopoly to the two firms in aggregate, and the incumbent must pay a substantial part of her rent to the platforms in order to defend her monopoly.

In Example 3, which depicts case (c), the incumbent has even less to gain from defending her monopoly. The gain is now 3. In equilibrium, the entrant gets an ad and he bids nothing for it.

Proposition 1 captures a strategic asymmetry between the incumbent and the entrant. The entrant wants the consumer to become aware of his product and he can do it through *any of the available platforms*. The incumbent, who wants to make sure that the entrant's product is not known to the consumer, must prevent entrant access on *all of the available platforms*. This is a strategic advantage for the entrant and it means that he has to pay less to make his product known than the incumbent has to pay to keep the consumer in the dark. If the incumbent succeeds in keeping out the entrant, it must be that she pays π_E to all platforms – that's a total of $n\pi_E$. If the entrant succeeds in getting his product known, he will pay at most π_E to all platforms and the incumbent pays nothing.²⁰

There are two remarks on Proposition 1:

Remark 2 *If entry increases total profit ($\pi_1 < \pi_2 + \pi_E$), condition (1) always fails and hence we always have duopoly.*

Remark 3 *If condition (1) fails, total platform profit is bounded above by*

$$\pi_1 - \pi_2 - (\bar{n} - 1)\pi_E = \pi_E.$$

²⁰The monopolization condition in Proposition 1 is similar to the monopolization condition in Proposition 1 of Asker and Bar-Isaac (2014). In both cases entry occurs if the number of agents is sufficiently large with respect to the ratio between the effect of entry on the incumbent and the effect of entry on the entrant.

4 Industry Analysis

Now that we have characterized what happens within each consumer segment, we turn our attention to the whole social platform industry. We consider the following timing:

1. In every industry k , a potential entrant appears, whose gross per user profit in case of entry is π_E , a random variable drawn from a cumulative distribution F (uncorrelated across industries).²¹
2. In every segment J , a random ordering of the platforms in J is selected and a sequence of auctions is run according to that order.
3. In every segment J , payoffs to consumers, platforms, the incumbent, and the entrant are made according to the outcome of the auction.

The equilibrium of the subgame beginning in 2, for every segment J is characterized in Proposition 1. The next proposition aggregates across consumer segments and entrant strength levels.

Proposition 4 *Expected consumer surplus is given by*

$$\bar{U} = u_2 - (u_2 - u_1) \sum_J m_J F\left(\frac{\pi_1 - \pi_2}{n_J}\right).$$

Whether a particular segment J remains a monopoly depends on a comparison between the entrant's strength π_E and the number of platforms n_J used in that segment. The term

$$\sum_J m_J F\left(\frac{\pi_1 - \pi_2}{n_J}\right)$$

can be interpreted as an average concentration index across segments.

Its expression becomes particularly simple if the entrant strength level π_E is uniformly distributed. Namely, assume that F is a uniform distribution with support on $[0, M]$ with $M > \pi_1 - \pi_2$. The latter assumption implies that there is always a positive probability of entry, even in a segment with just one platform. We then have

Corollary 1 *If F is a uniform distribution, expected consumer surplus is given by*

$$\bar{U} = a - b \sum_J m_J \frac{1}{n_J},$$

where a and b are constants.

The term $\sum_J m_J \frac{1}{n_J}$ is a simple concentration index: the weighted average platform share across segments, which in turn is just the reciprocal of the average number of independent platforms serving that segment. An increase in the concentration leads to an increase in the probability that the entrant is kept out of one or more segments, in turn leading to a consumer welfare loss.

²¹In this section, equilibrium is identical in all industries. In section 6, industry size γ_k will play a role in determining whether industry k advertises on social media or mass media.

5 Merger Analysis

Suppose now that two platforms, i and j , merge. The analysis proceeds as follows. We first characterize the optimal selling strategy of the merged entity and the ad allocation, segment by segment. We then aggregate across segments and examine the overall welfare effect of the merger.

Let us begin with the optimal selling strategy of the merged entity. The new company faces three sets of segments:

- (a) Those that used neither of the two merging platforms;
- (b) Those that used only one platform;
- (c) Those that used both platforms.

In (a) and (b), nothing changes after the merger, and Proposition 1 still applies as before. All the action is in (c), where the merged entity faces a choice between continuing to sell ads on the two platforms independently or bundling them into one item. Obviously, the merged entity is under no obligation to coordinate sales between the two platforms it owns. It will do so only if it is in its interest.

We assume that in either case the order in the auction is still random.²² We first show what happens if the new company chooses to bundle ads. This result will tell us whether the company prefers bundling and what its effect is on welfare.

Proposition 5 *If the merged entity moves from unbundled ad selling to bundled ad selling, its expected revenue in segment J is:*

- *Strictly lower if J was in equilibrium (a) before the merger;*
- *Strictly higher if J was in in equilibrium (b) before the merger;*
- *Weakly higher if J was in equilibrium (c) before the merger.*

If J was in (a) or (c) before the merger, consumer welfare U remains unchanged. If J was in (b), consumer surplus declines from u_2 to u_1 .

A merger is beneficial to the merging platforms in a given segment when it allows them to increase equilibrium revenues by restricting supply. Proposition 5 shows that this happens mostly when the number of pre-merger platforms n_J was less than a unit away from the incumbent monopoly threshold \bar{n} , namely when the pre-merger situation was the case (b) in Proposition 1. In that case, the

²²If the merged entity can manipulate the auction order, this would presumably give an additional reason to merge. However, this will not affect consumer welfare given the merger which is independent of the auction order (by Proposition 1)

merged entity exploits its newfound market power by coordinating sales across the two platforms: in practice this means selling one ad rather than two. This in turn induces the incumbent to choose a monopoly strategy and keep out the entrant. The overall platform revenue shoots up from case (b) to case (a). The proof of the proposition shows that the reduction in quantity is more than made up for by the increase in price.

This is the major positive effect of a merger for the new merged entity. There is also a minor positive effect: a segment in (c), where revenue is zero, could move to case (b), where revenue is small but strictly positive. If instead, the segment is already in (a), the merger has no effect because the merging entity will not benefit from reducing the number of ads it sells.

Proposition 1 also tells us when a merger will affect consumers. That happens when it tips the segment over from the entry case (b) to the monopoly case (a), which will obviously hurt consumers. In all other cases, consumers are unaffected.²³

Now that we know what the effect of a merger will be in each segment, we can compute the overall industry effect now by aggregating across segments, combining Propositions 4 and 5.

Proposition 6 *The aggregate loss in consumer surplus due to a merger between platform i and platform j is given by:*

$$-\Delta\bar{U} = (u_2 - u_1) \sum_{J \in M_{ij}} m_J \left(F \left(\frac{\pi_1 - \pi_2}{n_J - 1} \right) - F \left(\frac{\pi_1 - \pi_2}{n_J} \right) \right).$$

where M_{ij} is the set of segments where both platforms are present.

In the special case of a uniform distribution, the proposition simplifies to:

Corollary 2 *If F is a uniform distribution, consumer surplus loss is given by*

$$-\Delta\bar{U} = b \sum_{J \in M_{ij}} m_J \frac{1}{n_J (n_J - 1)},$$

where b is a constant.

The corollary helps understand the key factors that make a merger between platforms damaging to consumers. If n_J is constant in all segments where both platforms are present, then the effect is simply proportional to the share of consumers that are currently using both platforms

$$\sum_{J \in M_{ij}} m_J.$$

²³Notice a non-monotonic effect of mergers on welfare. Mergers are neutral either when the platform market is very fragmented (in which case there is always entry in the retail industry, as ads are mostly available at a zero price) or when it is very concentrated (in which case there is a foreclosure retail equilibrium, both pre- and post-merger). It is in the intermediate region where we identify welfare effects resulting from a merger.

If n_J varies across segments in M_{ij} , then the effect is stronger in segments that are already more concentrated.

Proposition 6, and the corollary, identify some possible issues for competition authorities that are dealing with social media platforms mergers.

First, a merger between attention oligopolists can hurt consumers through an anticompetitive effect on product markets. It can make it easier for incumbents to keep out entrants and thus lower consumer surplus.

Second, this negative effect depends on the extent of usage overlap between the merging platforms. The effect is nil if the two platforms have no common users before the merger, and it increases with the common usage share.

Third, the potential negative effect we have just identified is in principle measurable through platform usage. However, it is not enough to know the usage rates of the various platforms – or some other aggregate form of market share. One must also know the overlap between these platforms across consumers. This point is important because, to the best of our knowledge, previous merger cases in this area did not focus on this metric.²⁴

As highlighted in the discussion of modelling choices (Section 2.1), the optimal strategy for a merging pair arguably requires detailed information on the profits of the incumbent and of the entrant in the product market, and on the number of platforms which the consumer is using. However, we conjecture that the validity of our main insight is robust to coarser information for the platform operators as the underlying intuition is sufficiently simple. In essence, a firm with market power wants to restrict “output” – here, advertisers’ access to consumers – and its ability to do so depends on consumer-level market share – i.e., overlaps – rather than on its share of a broader market. We show next a theoretical result on the size of the error margin if one disregards the extent of consumer overlap. We also provide a numerical example to illustrate that the margin can be high.

5.1 Merger Assessment Based on Usage Shares

We illustrate the size of the measurement inaccuracy faced by a regulator who is trying to assess the effect of a platform merger on the basis of usage shares only, with no overlap information.

There are two platforms, 1 and 2. We consider the effect of a merger between the two platforms.

²⁴While it is well-known in Industrial Organization that market shares per se are not an indicator of market power, they still play a very relevant role in practice. The EC Horizontal Merger Guidelines state that market shares and concentration levels provide useful indications of the market structure and of the competitive importance of the merging parties (see Horizontal Merger Guidelines, OJ C 31, 5.2.2004, p. 5, ¶14). Specifically in the acquisition of WhatsApp by Facebook, the EC notes that “Facebook’s market shares are equal to [20-30]% in a number of Member States in a potential market for *overall online advertising*” (¶171 of the 2014 decision, emphasis added, http://ec.europa.eu/competition/mergers/cases/decisions/m7217_20141003_20310_3962132_EN.pdf).

There are four types of consumers: those who do not use either platform ($J = \emptyset$), those who use one platform only ($J = \{1\}$ and $J = \{2\}$), and those who use both ($J = \{1, 2\}$). If the share of consumers who use both platforms is $m_{\{1,2\}}$, the total consumer loss from a merger in the uniform case follows from Corollary 2 and is given by

$$-\Delta\bar{U} = \frac{b}{2}m_{\{1,2\}}.$$

Suppose we only know the usage rates of the two platforms, σ_1 and σ_2 , but not the extent to which usage overlaps. Given σ_1 and σ_2 , we can compute an upper and lower bound on $m_{\{1,2\}}$. The lowest value is when we minimize overlap given σ_1 and σ_2 , and that yields $\max(0, \sigma_1 + \sigma_2 - 1)$. The highest value is obviously $\min(\sigma_1, \sigma_2)$.

These bounds lead to the following:

Proposition 7 *Given σ_1 and σ_2 , the welfare effect of a merger can take any value:*

$$-\Delta\bar{U} \in \frac{b}{2} [m_{\min}, m_{\max}]$$

where

$$m_{\min} = \max(0, \sigma_1 + \sigma_2 - 1)$$

and


$$m_{\max} = \min(\sigma_1, \sigma_2).$$

If the two usage rates are the same and equal to σ , the expression becomes

$$-\Delta\bar{U} \in \frac{b}{2} [\max(0, 2\sigma - 1), \sigma]$$

The plot below depicts the gap between the two bounds for every possible value of the usage rate σ (and setting $b = 2$). If the usage rate is below 50%, the lower bound is zero. The maximal difference is when the usage rate is one half, in which case the difference between the upper and the lower bound is 0.5. Given that, in the example, we normalize to 1 the maximal consumer loss, this difference is also

large in magnitude.



In this simple example, a regulator who knows usage shares but not usage overlaps risks making highly inaccurate decisions. This is particularly evident for usage share values just below 50%: they are consistent with a situation where the merger is completely harmless (the zero lower bound, when platforms do not overlap) and one in which it can destroy a large amount of consumer welfare (the upper bound, with full platform overlap).

5.2 Application to Mergers in Practice: The Case of US Social Media Platforms

How would a regulator have to proceed in a merger case in practice? Our framework suggests a two-step approach. First, one would have to identify the relevant platforms that can be considered as attention brokers, according to the definition used in this paper. Second, one would need to collect demand-side information about the users of such platforms - crucially accounting for the overlaps, so that a merger simulation could be run.

As for the first step, while time spent on different media platforms is a good starting point when thinking about users' attention, just looking at daily consumption across media as reported, for instance, by the Nielsen Total Audience Report would still not be sufficient to run a meaningful analysis. Data would have to be collected about those particular platforms that attract users' attention with the specific purpose of conducting hyper-targeted advertising. This implies also to form a view as to which product markets could be targeted, as well as to the relevant time span for a decision in a product market (e.g., the purchase of a car may take several weeks to complete, while the service of a local plumber may not be delayed for more than a few hours).

As for the second step, we report below a brief empirical section that applies our findings to publicly available platform usage data. In the example that follows, we skip the first step and simply

assume that the social media platforms identified are all capable of conducting hyper-targeted advertising. Both the underlying data and the empirical implementation are subject to obvious limitations. The objective is simply to illustrate how our theoretical analysis could be applied to actual social media.

The data come from the Pew Institute Survey on Social Media Trends.²⁵ The survey ran from July 12 to August 8, 2016 (wave 19). There were 4,579 participants, the majority of whom participate online (4,165) while the rest completed the survey through mail (414). Respondents were drawn from the American Trends Panel, a national sample of adults in the US. Panelists were obtained through two large RDD (random digital dial) surveys from the Pew Research Institute. 899 people refused to do the online survey out of the 5,064 people initially selected.

The survey asked participants whether they used any subset of these three social media platforms: Facebook, Instagram, and Twitter. Usage rates at the time were 70.7% for Facebook, 19.3% for Instagram, and 17.3% for Twitter. We now use these usage rates to describe the effects of pairwise mergers between these platforms.²⁶

As argued in the previous section, usage rates alone provide an imprecise measure of the potential welfare effect of a platform merger. We can see this point in columns (1) and (2) of Table 1, which report the lower bound m_{\min} (no overlap) and upper bound m_{\max} (maximal overlap) on the consumer welfare effect of a pairwise merger. The range between the two bounds represents all the possible welfare effect values assessed by a regulator that observes aggregate platform usage but not overlap, as derived in Proposition 7. The bounds are wide. They go from a zero effect when the merging platforms have no overlapping consumers to an effect equal to the less used of the two platforms when overlap is maximal, further illustrating the usefulness of overlap information.

The remainder of Table 1 derives the actual (negative) welfare effect of the three pairwise mergers on the basis of Corollary 1. Column (3) reports pairwise overlaps m_{ij} , while Column (4) reports the three-way overlap m_{123} . The latter is by definition the same for all pairs, while the former is greatest for the Facebook-Instagram pair. The total actual consumer welfare loss $-\Delta\bar{U}$, computed as in Corollary 2, is therefore largest for the Facebook-Instagram pair.²⁷

²⁵ Available from: <http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/>

²⁶ This an hypothetical exercise. As mentioned in the introduction, the merger between Facebook and Instagram was in fact already consummated in 2014.

²⁷ In the table, results are scaled by the factor $b/6$.

	(1)	(2)	(3)	(4)	(5)
<i>Merging Pair</i>	<i>Lower</i>	<i>Upper</i>	<i>Pairwise</i>	<i>Three-way</i>	<i>Welfare</i>
	<i>Bound</i>	<i>Bound</i>	<i>Overlap</i>	<i>Overlap</i>	<i>Effect</i>
Facebook-Instagram	0%	19.3%	9.4%	8.4%	6.1%
Facebook-Twitter	0%	17.3%	7.0%	8.4%	4.9%
Instagram-Twitter	0%	17.3%	0.5%	8.4%	1.7%

Table 1. Consumer welfare effects of platform mergers

6 Targeted and Non-Targeted Advertising

The model introduced in Section 3 included two advertising channels: non-targeted mass media ads and targeted platform ads. So far, we have abstracted from the first channel by assuming that the advertising fee a is sufficiently large that no entrant would buy non-targeted ads. We now explore the two advertising channels jointly. The section will show that the baseline results are essentially robust to the presence of this extension but the preemption effect is stronger in niche industries.

In the first stage of the game, the entrant in industry k chooses whether it pays a flat fee a to buy a non-targeted ad. If it does, then its product is known to consumers. If it does not buy the non-targeted ad, then we proceed to the auction subgame analyzed in Proposition 1. The only difference is that all firm payoffs are pre-multiplied by the size of the relevant industry: the monopolist incumbent's payoff is now $\gamma_k \pi_1$, the duopolist incumbent's payoff is $\gamma_k \pi_2$, and the entrant's payoff is $\gamma_k \pi_E$.

Note that, if the entrant buys the non-targeted ad, he receives a payoff $\gamma_k \pi_E - a$, where a does not depend on the size of the industry the advertiser is in. This means that non-targeted advertising is particularly useful for entrants in *mainstream* industries – the ones with a large consumer share γ_k – because the advertising fee is spread over a larger consumer base. The main result is that the effect of a merger between digital platforms is felt more intensely in *niche* industries.

To keep the analysis simple, we assume that all consumers have the same digital platform consumption, namely, they all use n platforms. If that is not the case, the results below are qualitatively similar but the expressions more involved because of the interaction between the size of the industry as a whole and the size of the segment in that industry.

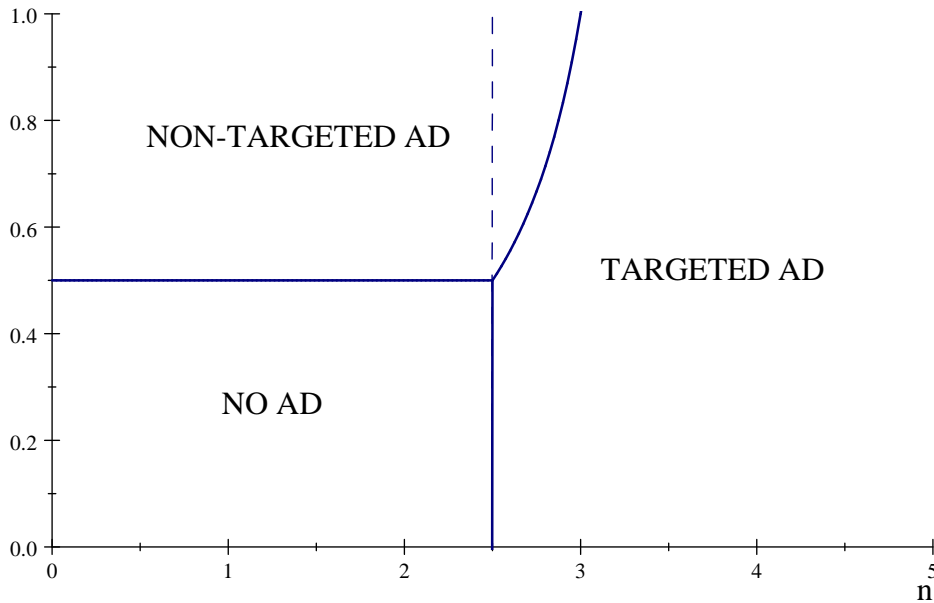
Proposition 8 *Equilibrium in industry k depends on industry size γ_k and on the number of independent platforms n :*

(i) *If k is sufficiently niche ($\gamma_k < a/\pi_E$) and there are few platforms ($n < \bar{n}$), then the entrant buys no ads and the industry is monopolized.*

- (ii) If k is sufficiently mainstream ($\gamma_k > a/\pi_E$) and there are an intermediate number of platforms ($n \in [\bar{n}, \bar{n} + 1 - \frac{a}{\gamma_k \pi_E})$), then the entrant buys a non-targeted ad and there is entry;
- (iii) In the remaining cases, the entrant buys a targeted ad and there is entry.

Proposition 8 is illustrated in the figure below. The x -axis represents the number of independent platforms (n) and the y -axis is the importance of the industry (γ_k). Entry can be achieved in two ways: if the industry is sufficiently mainstream (large γ_k and low to intermediate n), the entrant will be willing to buy a non-targeted ad; if the number n of platforms is sufficiently large, the entrant will manage to win at least one targeted ad. The latter condition is the same as in the baseline result: $n \geq \bar{n}$. Monopoly occurs in the bottom-left quadrant of the figure, that is, in niche industries if the number of platforms is sufficiently low.

A secondary effect of the presence of non-targeted advertising is some erosion of the profits of platforms. Absent non-targeted advertising, when $n \in (\bar{n}, \bar{n} + 1)$ platforms always receive a profit selling the targeted ad to the entrant. With non-targeted advertising, this is no longer always the case, as the entrant may prefer to buy a non-targeted ad if the industry is sufficiently mainstream. In the plot, that corresponds to the v-shaped region between the dashed line and the curved solid line.



Presence of an Entrant Ad in Equilibrium: x -axis: number of platforms (n);
 y -axis: size of industry (γ_k).

Once we have characterized equilibrium behavior, we can derive consumer welfare and determine

the effect of a merger in the presence of non-targeted ads. As before, we assume that π_E is distributed according to F . As before, the merger occurs before the game starts and before any ad – targeted or not – is purchased.

Proposition 9 *Expected consumer surplus is given by*

$$\bar{U} = u_2 - (u_2 - u_1) \left(M + \nu F \left(\frac{\pi_1 - \pi_2}{n} \right) \right),$$

where

$$\nu = \sum_{k: \gamma_k < \frac{an}{\pi_1 - \pi_2}} \gamma_k \text{ and } M = \sum_{k: \gamma_k \geq \frac{an}{\pi_1 - \pi_2}} \gamma_k F \left(\frac{a}{\gamma_k} \right).$$

Proposition 9 shows that the welfare costs are concentrated in niche industries. Mainstream industries are shielded because they find it economical to use non-targeted advertising.

A merger has an inframarginal effect on niche industries through $F \left(\frac{\pi_1 - \pi_2}{n} \right)$. Within the set of industries for which $\gamma_k < \frac{an}{\pi_1 - \pi_2}$, the probability of product monopolization increases in the same way as it does in the baseline case (Proposition 6). The merger also has a partially offsetting marginal effect: the set of firms that uses non-targeted advertising increases as the right-hand side of condition $\gamma_k < \frac{an}{\pi_1 - \pi_2}$ increases. However, as the figure above illustrates, this offset is partial because it can only affect mainstream industries ($\gamma_k > a/\pi_E$).

7 Robustness Checks

7.1 Alternative Ad Selling Mechanism

The model was based on the assumption that platforms use a simple selling mechanism. The selling order is randomized and each platform uses a second-price auction. This section asks whether our key equilibrium characterization (Proposition 1) is robust to alternative specifications of the mechanism space.

We approach the question in two ways: a more concrete one – we suggest a different mechanism and show that it yields the same result – and a more abstract one – we consider a corresponding cooperative game and show that the same condition determines the presence of a stable coalitional partition.

Starting with the more concrete approach, the auction mechanism assumed in the baseline model may be criticized because it assigns a purely passive role to platforms: they always run the same kind of auction, thus forgoing potential gains they may make by exploiting their market power. Here we examine what can be seen as the polar opposite case: suppose every platform can make a take-it-or-leave-it offer to one of the two producers, with the understanding that, if that producer rejects the offer, the ad will go to the other producer. For every platform, the mechanism is therefore $(z_i, t_i) \in \{I, E\} \times \mathfrak{R}^+$, where

z_i is the identity of the producer the offer is made to and t_i is the amount requested. Platforms have commitment power and the commitment is observed by everyone: they each announce (z_i, t_i) .²⁸

As before, we focus on a segment with n platforms. The game is composed of three stages:

- The order of platforms is randomized and observed by all players: 1, 2, ... n .
- In the announcement stage, platforms announce their mechanisms simultaneously.
- In the acceptance stage, the mechanisms selected by the platforms are played publicly in order 1, 2, 3, ..., n .²⁹

We focus on the set of pure-strategy subgame perfect equilibria. As in the baseline case, we have a monopoly equilibrium when the incumbent buys all ads, while an entry equilibrium is the set of complementary cases where at least one ad is purchased by the entrant.

Proposition 10 *A monopoly equilibrium exists if and only if $n \leq \bar{n}$.*

Let us illustrate the Proposition by re-visiting the examples in Figure 1. In Example 1, there is a symmetric equilibrium where each platform demands $t_i = 7/3$ from the incumbent. If any of the platforms tried to demand more, the incumbent would rather not buy any ad. If a platform sold an ad to the entrant instead of the incumbent, the most it would get is 2.³⁰

In Example 2, there is no equilibrium where all the ads go to the incumbent. If such an equilibrium existed, at least one of the platforms would have to charge a price below 2, but then that platform would rather switch to selling its ad to the entrant. Instead, there is an equilibrium where all ads are sold to the entrant for zero.

The more abstract approach shows that the $n < \bar{n}$ condition captures a strategic feature of the environment we are considering. We move from non-cooperative to cooperative game theory and we

²⁸The analysis of the no-commitment case is very simple.

If platforms cannot commit to a price and $n \geq 2$, the monopoly equilibrium never exists.

To see this, note that in a monopoly equilibrium, the incumbent must start the last auction holding $n - 1$ ads. In that case, the last platform would charge her the whole surplus: $\pi_1 - \pi_2$. Predicting this, the incumbent would be willing to pay at most zero in any of the previous auctions, implying that previous platforms would prefer to sell to the entrant, who is willing to pay any price up to π_E to avoid the monopoly outcome.

Obviously, if $n = 1$, the monopoly equilibrium will arise.

Thus, the consumer welfare analysis of the no-commitment case is exactly the same as in the baseline with $\bar{n} \in (0, 1)$. The effect on consumer welfare of a merger is also immediately derived: a merger matters if and only if it leads to $n = 1$.

²⁹The assumption that the acceptance stage is sequential guarantees the existence of a pure-strategy equilibrium. If the acceptance stage was simultaneous, we conjecture that the proposition could be re-written as: A pure-strategy equilibrium where all ads are sold to the incumbent exists if and only if $n \leq \bar{n}$.

³⁰Obviously, this is not the only equilibrium: there is a continuum of asymmetric equilibria with $t_1 + t_2 + t_3 = 7$ and $t_i \geq 2$ for all platforms.

show that the $n < \bar{n}$ condition determines whether the game has a stable outcome where only the incumbent advertises or not.

Assume that instead of playing the auction described in Section 2, the incumbent, the entrant, and the n platforms are engaged in a transferable utility game (TUG). The set of players is $\{I, E, 1, \dots, n\}$. The characteristic function v has the following properties:

- The coalition that includes the incumbent and all platforms gets $v_{I,1,\dots,N} = \pi_1$ for itself, while the coalition that includes only the entrant gets $v_E = 0$.
- A coalition that includes the incumbent and a set of platforms P that does not include all platforms gets $v_{I,P} = \pi_2$.
- A coalition that includes the entrant and a non-empty set of platform P gets $v_{E,P} = \pi_E$.
- Coalitions including both E and I are not feasible: the two producers cannot directly cooperate because they cannot collude (if they could, the solution would be simple: the incumbent would pay the entrant an amount between π_E and $\pi_1 - \pi_2$ to stay out).
- Any other coalition gets zero.

This is not a superadditive transferable utility game because the grand coalition is not feasible, hence the concept of core is not applicable. However, one can ask whether the game has a stable partition of the players into coalitions. A partition is stable if there exists an imputation x – a vector of payoffs to players that is compatible with the characteristic function – such that no coalitional deviation can guarantee a strictly higher payoff for each of its members.

Proposition 11 *If $n < \bar{n}$, the only stable partition of the TUG is $\{\{I, 1, \dots, n\}, \{E\}\}$. If $n > \bar{n}$, the set of stable partitions is empty.*

The intuition for this result generalizes the intuition for the two non-cooperative games we considered: the sequential auctions and the take-it-or-leave-it offers. Namely, if $n < \bar{n}$ the incumbent's monopoly profit π_1 is large enough for the incumbent to pay at least π_E to every platform and leave at least π_2 for herself. If instead the condition fails, the incumbent is not willing to pay at least π_E to every platform and that leaves at least one of platform open to switching to the entrant.

7.2 General Model

The goal of this section is to relax four assumptions that were made in the baseline case: (i) There is one incumbent only; (ii) There is one entrant only; (iii) Each platform sells one ad only; (iv) Being exposed to one ad provides the consumer with perfect information. In a stylized extension, we will allow

for any number of incumbents, a large number of potential entrants, any number of ads per platform, and imperfectly informative ads.

Consider a consumer segment with n platforms. Each platform shows k ads. We focus on one consumer who is interested in a product made by a retail industry with q incumbents and a large number of potential entrants.

Every platform runs a $k + 1$ th price auction (similar to the one used by Google). Each firm submits a bid and can buy at most one ad on that platform. The highest k bidders receive an ad and they all pay the bid of the $k + 1$ th bidder. All auctions are run simultaneously. Thus, each firm submits a n -vector of non-negative bids to all platforms.

The probability that a consumer sees any specific ad is given by $p \in [0, 1]$. We assume that this probability is independent across ads and platforms.

There are a large number of entrants. The payoff of each entrant from a specific consumer is π_E if the consumer learns about the entrant's product and zero otherwise. The consumer learns about the product if she sees an ad by the entrant.

There are q incumbents. To make the problem interesting, assume that $q \geq k$. The incumbent's utility depends on the number of entrants the consumer is aware of. With no entrants, it is π_H , with at least one entrant, it is π_L , with $\pi_H > \pi_L$.³¹

The probability of entry – namely the probability that the consumer becomes aware of at least one entrant – depends on the number of ads bought by entrants. By the independence assumption above, it does not depend on which platforms the ads appear on, or which set of entrants get ads (or whether multiple ads are bought by the same entrant).³² If m ads end up with entrants, the probability that the consumer becomes aware of at least one entrant is:

$$P_m = 1 - (1 - p)^m.$$

We say we are in a *pre-emption equilibrium* if all the ads from all the platforms are bought by incumbents. We restrict attention to pure-strategy equilibria where every ad receives at least one full-value bid from some entrant, where a full-value bid is one that equals the additional payoff the entrant would receive if he won the ad.³³

The following result is a partial extension of Proposition 1 to this more general environment:

Proposition 12 *A pre-emption equilibrium exists if and only if the following condition is satisfied*

$$p\pi_E \leq \frac{P_K (\pi_H - \pi_L)}{K}.$$

³¹The analysis can be extended to a payoff that depends on the number of successful entrants as long as the marginal effect of entry is decreasing in the number of entrants.

³²Alternative formulations can be accommodated at a higher notation cost.

³³This restriction eliminates equilibria where the auction winner pays less than the losers' valuation. In auctions with one object, these equilibria are usually ruled out by invoking weak dominance. However, with multiple simultaneous auctions, weak dominance is not applicable.

where $K = \lceil nk/q \rceil$ (the smallest integer that is at least as large as nk/q).

To understand the proposition, we consider two special cases. In what follows we assume that $\frac{nk}{q}$ is an integer.

1. One incumbent, perfectly informative ads, one ad per platform. When $q = 1$, $k = 1$ and $p = 1$, the no-deviation condition in Proposition 12 boils down to the monopolization condition in our baseline Proposition

$$n \leq \frac{\pi_H - \pi_L}{\pi_E}.$$

2. Perfectly informative ads. When $p \rightarrow 1$, Proposition 12 boils down to the monopolization condition in our baseline Proposition

$$\frac{nk}{q} \leq \frac{\pi_H - \pi_L}{\pi_E}.$$

This condition is similar to the baseline case, but the number of independent platform is multiplied by the ratio between k and q . A higher number of ads per platform makes monopolization harder. A higher number of incumbents makes it easier (but it is important to keep in mind that a higher number of oligopolists is also likely to modify $\pi_H - \pi_L$ and consumer welfare).

As mentioned above, Proposition 12 is only a partial extension of Proposition 1. Both propositions provide a necessary and sufficient conditions for pre-emption. However, only the latter characterizes what happens when the pre-emption equilibrium does not exist. The problem is that, if there are multiple ads, when the pre-emption does not exist, all (reasonable) equilibria involve mixed strategies: entry occurs with positive probability but we do not know how to characterize that probability. The general result is therefore qualitative: entry increases consumer welfare but we cannot say by how much in a general way. One could of course compute the mixed-strategy equilibrium numerically for specific instances.

This difficulty is not specific to our framework. It is a manifestation of a general issue when there multiple auctions. One can show that an equilibrium exists (e.g., Simon and Zame, 1990). However, a general analytical characterization does not exist. Szentes and Rosenthal (2003) provide a solution with three objects and two bidders for first- and second-price auctions: the complexity of solving that case explains a general characterization has yet to be found.

8 Conclusions

- Competition among online platforms for users lacks one of the hallmarks of competition: prices are typically zero and do not determine competitive outcomes. Competition shifts to the other side, that is, producers who reach out for users through advertising offered by the platforms.

- Time/attention of users is a scarce resource. This is controlled by platforms who can become bottlenecks for producers who want to reach users that utilize a small set of platforms. This in turn hurts users as they end up having less product choice and paying higher prices.
- Platforms may be tempted to exploit the attention bottleneck to their advantage, to the possible detriment of consumers. This paper analyzed a highly stylized model of attention bottleneck and found conditions under which platforms increase their revenues by shrinking the bottleneck and directing all consumers' attention to incumbents.
- We argue that this situation typically happens when online platforms are concentrated, and they manage to set prices for ads that select incumbent producers to the detriment of entrants. In this sense, online platform market power is not good for innovative entrants (who need advertisements to be known by users), nor is it for users (who end up buying at a high price from incumbent producers). Market power begets market power.
- One contribution of our paper is therefore to see the effects of online platform competition through the impact it has on product markets. This assessment does not require a radical departure from existing competition policy, rather to correctly apply first-order economic principles to these markets.
- In fact we show that, with many competing platforms, it becomes very expensive for incumbents to bid out entrants. Entrants get to be known, consumers are typically well off and marginal mergers would not matter.
- As online platforms become more concentrated, competition for ad scarcity works to the advantage of incumbent producers. A merger between online platforms allows them to better control this scarcity, tilting the game even more in favor of incumbents.
- The crucial element in an online merger assessment (as in any other merger) is to look at the overlaps of users across platforms. If consumers multi-home, scarcity is more likely to disappear, making entry also more likely.
- A problematic merger is therefore one between concentrated online platforms with overlapping users. We show that standard metrics that ignore these overlaps and just concentrate on usage can lead to large biases. Even more so, metrics that only focus on the supply-side (market shares) are inappropriate in these markets. In contrast, we show that existing individual-level platform usage data can be used to account for overlaps and obtain more meaningful estimates.
- While we concentrated the policy discussion on mergers, we note that our analysis could also be used to inform the debate on other regulatory questions. For instance, in our model a dominant

platform should not be allowed to sell too many ad slots to incumbent retail firms when this results in the exclusion of entrants.

- We also draw a line between advertising on traditional media and targeted ads on online platforms. We find that potential competition problems are more likely to be related to niche product markets for whom targeted ads are the ideal channel to reach consumers.

References

- [1] Agrawal, A, J. Gans and A. Goldfarb (2018), *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press, Boston.
- [2] Ambrus, A., E. Calvano and M. Reisinger (2016), “Either or Both Competition: A Two-sided Theory of Advertising with Overlapping Viewership”, *American Economic Journal: Microeconomics*, 8, 189-222.
- [3] Armstrong M. and J. Vickers (2019), “Discriminating Against Captive Customers”, *AER: Insights*, 1, 3, 257-272.
- [4] Anderson, S. and S. Coate (2005), “Market Provision of Broadcasting: A Welfare Analysis”, *Review of Economic Studies*, 72, 947-972.
- [5] Anderson, S. and J. McLaren (2012), “Media Mergers and Media Bias with Rational Consumers”, *Journal of the European Economic Association*, 10, 831-859.
- [6] Anderson, S., O. Foros and H. Kind (2018), “Competition for Advertisers and for Viewers in Media Markets”, *Economic Journal*, 128, 34-54.
- [7] Armstrong, M. (2006), “Competition in Two-sided Markets”, *RAND Journal of Economics*, 37, 668-91.
- [8] Asker, J. and H. Bar-Isaac (2014), “Raising Retailers’ Profits: On Vertical Practices and the Exclusion of Rivals”, *American Economic Review*, 104, 2, 672-686.
- [9] Athey, S. and D. Nekipelov (2014), “A Structural Model of Sponsored Search Advertising Auctions”, mimeo.
- [10] Athey, S., E. Calvano and J. Gans (2018), “The Impact of Consumer Multi-homing on Advertising Markets and Media Competition”, *Management Science*, 64, 1574-1590.
- [11] Bagwell, K. (2007), “The Economics Analysis of Advertising”, in *Handbook of Industrial Organization*, vol. 3, Elsevier (Amsterdam), 1701-1844.

- [12] Blake, T., C. Nosko and S. Tadelis (2015), “Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment”, *Econometrica*, 83, 155-174.
- [13] Bernheim, D. and M. Whinston (1986), “Common Agency”, *Econometrica*, 54, 923-942.
- [14] Caillaud, B. and B. Jullien (2013), “Chicken & Egg: Competition among Intermediation Service Providers”, *RAND Journal of Economics*, 34, 309-328.
- [15] Chandra A. and A. Collard-Wexler (2009), “Mergers in Two-sided Markets: An Application to the Canadian Newspaper Industry”, *Journal of Economics & Management Strategy*, 18, 1045-70.
- [16] De Carolis, F. and G. Rovigatti (2019), “From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising”, mimeo.
- [17] de Corniere, A. and R. De Nijs (2016), “Online Advertising and Privacy”, *RAND Journal of Economics*, 47, 48-72.
- [18] Diermeier, D. and R. B. Myerson (1999) “Bicameralism and Its Consequences for the Internal Organization of Legislatures”, *American Economic Review*, 89(5).
- [19] Eaton, B.C. and R. Lipsey (1979), “The Theory of Market Pre-emption: The Persistence of Excess Capacity and Monopoly in Growing Spatial Markets”, *Economica*, 46, 149-158.
- [20] Eliaz, K and R. Spiegler (2019), “Incentive-Compatible Advertising on Non-Retail Platforms”, *Rand Journal of Economics*, forthcoming.
- [21] Evans, D. (2017), “The Economics of Attention Markets,” mimeo.
- [22] Gilbert, R. and D. Newbery (1982), “Preemptive Patenting and the Persistence of Monopoly”, *American Economic Review*, 72, 514-526.
- [23] Kaldor, N. (1950), “The Economic Aspects of Advertising”, *Review of Economic Studies*, 18, 1-27.
- [24] Khan, L. and S. Vaheesan (2017), “Market Power and Inequality: The Antitrust Counterrevolution and its Discontents”, *Harvard Law & Policy Review*, 11, 235-294.
- [25] Jeziorski, P. (2014), “Effects of Mergers in Two-sided Markets: The U.S. Radio Industry”, *American Economic Journal: Microeconomics*, 6, 35-73.
- [26] Lynn, B. (2013), “Antitrust: A Missing Key to Prosperity, Opportunity, and Democracy”, *Demos*, policy paper, October 2013.
- [27] Milgrom, P. and S. Tadelis (2019), “How Artificial Intelligence and Machine Learning Can Impact Market Design”, in *The Economics of Artificial Intelligence: An Agenda* (eds. A. Agrawal, J. Gans and A. Goldfarb), National Bureau of Economic Research Conference Report.

- [28] Prat, A. (2018), “Media Power”, *Journal of Political Economy*, 126, 4, 1747-1783.
- [29] Prat, A. and A. Rustichini (2003), “Games Played Through Agents”, *Econometrica*, 71, 989-1026.
- [30] Rey, P. and J. Tirole (2007). “A Primer on Foreclosure,” in *Handbook of Industrial Organization* Vol. 3, Elsevier (Amsterdam).
- [31] Rochet, J.-C. and J. Tirole (2003), “Platform Competition in Two-sided Markets”, *Journal of the European Economic Association*, 1, 990-1029.
- [32] Rysman, M. (2004), “Competition between Networks: A Study of the Market for Yellow Pages”, *Review of Economic Studies*, 71, 483-512.
- [33] Schmalensee, R. (1978), “Entry Deterrence in the Ready-to-Eat Breakfast Cereal Industry”, *Bell Journal of Economics*, 9, 305-327.
- [34] Simon, L. and W. Zame (1990). “Discontinuous Games and Endogenous Sharing Rules”, *Econometrica*, 58, 861-872.
- [35] Simonov, A., C. Nosko and J. Rao (2018). “Competition and Crowd-Out for Brand Keywords in Sponsored Search”, *Marketing Science*, 37, 2, 200-215.
- [36] Szentes, B. and R. Rosenthal (2013). “Beyond chopsticks: Symmetric equilibria in majority auction games”, *Games and Economic Behavior*, 45, 278-295.
- [37] Telser, L. (1964), “Advertising and Competition”, *Journal of Political Economy*, 72, 537-562.
- [38] Whinston, M. (2006). *Lectures on Antitrust Economics*, MIT Press (Boston, MA).
- [39] Wu, T. (2019), “Blind Spot: The Attention Economy and the Law”, *Antitrust Law Journal*, 82, 3.

Appendix: Proofs

Proposition 1

Focus on a segment J . To simplify notation, replace n_J with n .

Solve the sequential auction game by backward induction. After all n auctions are run, the gross payoffs of the incumbent and the entrant are respectively $(\pi_1, 0)$ if the incumbent wins all auctions and (π_2, π_E) if the entrant wins at least one auction. After $n - 1$ auctions, both the incumbent and the entrant are willing to bid at most zero if the entrant has already won at least one auction, and they are willing to bid $(\pi_1 - \pi_2, \pi_E)$ respectively if the incumbent has won all previous auctions, in which case the incumbent wins the n -th auction if $\pi_1 - \pi_E > \pi_2$.

The continuation payoffs for the two players after $n - 1$ auctions is

$$\left(\Pi_I^{(n-1)}, \Pi_E^{(n-1)}\right) = \begin{cases} (\pi_1 - \pi_E, 0) & \text{if the incumbent has won all previous auctions} \\ & \text{and } \pi_1 - \pi_E > \pi_2 \\ (\pi_2, \pi_E) & \text{otherwise} \end{cases}$$

By induction we see that the continuation payoff after j auctions is always (π_2, π_E) if the entrant has won at least one auction.

$$\left(\Pi_I^{(j)}, \Pi_E^{(j)}\right) = \begin{cases} \left(\Pi_I^{(j+1)} - \pi_E, 0\right) & \text{if the incumbent has won all previous auctions} \\ & \text{and } \pi_1 - (n - j)\pi_E > \pi_2 \\ (\pi_2, \pi_E) & \text{otherwise} \end{cases}$$

At the beginning of the game (when $j = 0$), the expression above becomes

$$\left(\Pi_I, \Pi_E\right) = \begin{cases} (\pi_1 - n\pi_E, 0) & \text{if } \pi_1 - \pi_2 > n\pi_E \\ (\pi_2, \pi_E) & \text{otherwise} \end{cases}$$

If $\pi_1 - \pi_2 > n\pi_E$, the incumbent wins all auctions and total platform revenue is $n\pi_E$. If $\pi_1 - \pi_2 < n\pi_E$, the entrant wins the first auction and pays

$$\max(\pi_1 - \pi_2 - (n - 1)\pi_E, 0)$$

for the first ad; all bids are zero thereafter and it does not matter who gets the ad.

Proposition 4

By Proposition 1, the entrant buys an ad in segment J if and only if $n_J > \frac{\pi_1 - \pi_2}{\pi_E}$, which happens with probability

$$1 - F\left(\frac{\pi_1 - \pi_2}{n_J}\right).$$

Aggregating across segments yields the proposition.

Corollary 1

Follows directly from Proposition 4, where $a = u_2$ and $b = \frac{(u_2 - u_1)(\pi_1 - \pi_2)}{M} > 0$.

Proposition 5

Bundling the two ads is equivalent to selling only one ad (using one platform and letting the other one idle). Call $R(n)$ the expected platform revenue in a segment with n platforms, as in Proposition 1. In the unbundled case, the expected platform revenue for the merged entity is $2R(n)$, while in the bundled case it is $R(n - 1)$. Compare now $2R(n)$ and $R(n - 1)$ in cases (a), (b), and (c).

If the segment was in (a) before the merger, it remains in (a) after the merger: $R(n) = R(n-1) = \pi_E$, and therefore $2R(n) > R(n-1)$.

If the segment was in (b), it must be that $n_J \in (\bar{n}, \bar{n} + 1)$. By bundling ad sale, the merged entity moves the equilibrium from (b) to (a) and increases its expected revenue from twice the expected revenue in (b), namely

$$2 \frac{\pi_1 - \pi_2 - (n_J - 1) \pi_E}{n_J}$$

to once the expected revenue in (a), namely π_E . Note that, as $n_J > \bar{n}$ and $n_J \geq 2$,

$$2 \frac{\pi_1 - \pi_2 - (n_J - 1) \pi_E}{n_J} < 2 \frac{\pi_1 - \pi_2 - (\bar{n} - 1) \pi_E}{n_J} = 2 \frac{\pi_E}{n_J} \leq \pi_E,$$

where the first inequality is due to $n_J > \bar{n}$ and the second one is due to $n_J \geq 2$ (because at least the two merging platforms are present in that segment). Thus, $R(n-1) > 2R(n)$.

If the segment was in (c), $R(n-1) \geq 2R(n)$ because $R(n) = 0$.

The second part of the proposition derives from the following observations combined with Proposition 1. If J was in (a), it remains in (a) after the merger. If J was in (b), it goes to (a). If J was in (c), it remains in (c) or goes to (b).

Proposition 6

In each segment J , consumer welfare changes whenever $n_J \in (\bar{n}, \bar{n} + 1)$. That is

$$\frac{\pi_1 - \pi_2}{\pi_E} < n_J < \frac{\pi_1 - \pi_2}{\pi_E} + 1.$$

The second inequality rewrites as

$$\pi_E < \frac{\pi_1 - \pi_2}{n_J - 1}.$$

The first inequality becomes

$$\pi_E > \frac{\pi_1 - \pi_2}{n_J}.$$

Corollary 2

We have

$$\left(F \left(\frac{\pi_1 - \pi_2}{n_J - 1} \right) - F \left(\frac{\pi_1 - \pi_2}{n_J} \right) \right) = \frac{1}{M} \left(\frac{\pi_1 - \pi_2}{n_J - 1} - \frac{\pi_1 - \pi_2}{n_J} \right) = -\frac{\pi_1 - \pi_2}{M} \frac{1}{n_J(n_J - 1)}.$$

The Corollary follows, with $b = \frac{(u_2 - u_1)(\pi_1 - \pi_2)}{M}$ as defined in Corollary 1.

Proposition 8

From Proposition 1, if $n \geq \bar{n}$, the entrant prefers to buy a targeted ad rather than buying no ad; however he prefers to buy a non-targeted ad to a targeted ad if the cost of the latter is smaller than

the cost of the former. The cost of a targeted ad is $(\pi_1 - \pi_2 - (n - 1)\pi_E)\gamma_k$ if $n_J \in [\bar{n}, \bar{n} + 1)$, and 0 if $n \geq \bar{n} + 1$. Thus, if $n \geq \bar{n} + 1$ the entrant will always prefer targeted ads that are available for free instead of mass media ads that involve the payment of the fixed fee a . Instead, if $n \in [\bar{n}, \bar{n} + 1)$ the entrant prefers a non-targeted ad if

$$a \leq (\pi_1 - \pi_2 - (n_J - 1)\pi_E)\gamma_k.$$

That is

$$n \leq \frac{\pi_1 - \pi_2 - \frac{a}{\gamma_k}}{\pi_E} + 1,$$

which can be rewritten as

$$n \leq \bar{n} + 1 - \frac{a}{\gamma_k \pi_E},$$

which is a number in $[\bar{n}, \bar{n} + 1)$ if $\gamma_k > a/\pi_E$.

Proposition 9

By Proposition 8, industry k is monopolized if and only if $\gamma_k < a/\pi_E$ and

$$n < \bar{n} = \frac{\pi_1 - \pi_2 - \frac{a}{\gamma_k}}{\pi_E}$$

Thus, the industry is monopolized if

$$\pi_E < \max\left(\frac{a}{\gamma_k}, \frac{\pi_1 - \pi_2}{n}\right),$$

and is in the entry equilibrium otherwise.

Adapting Proposition 4, we have

$$\begin{aligned} \bar{U} &= u_2 - (u_2 - u_1) \sum_k \gamma_k F\left(\min\left(\frac{a}{\gamma_k}, \frac{\pi_1 - \pi_2}{n}\right)\right) \\ &= u_2 - (u_2 - u_1) \left(\sum_{k:\gamma_k \geq \frac{an}{\pi_1 - \pi_2}} \gamma_k F\left(\frac{a}{\gamma_k}\right) + \sum_{k:\gamma_k < \frac{an}{\pi_1 - \pi_2}} \gamma_k F\left(\frac{\pi_1 - \pi_2}{n}\right) \right) \\ &= u_2 - (u_2 - u_1) \left(M + \nu F\left(\frac{\pi_1 - \pi_2}{n}\right) \right). \end{aligned}$$

Proposition 10

For the “if” part, consider the following candidate equilibrium: Each platform demands $(I, \frac{\pi_1 - \pi_2}{n})$. Given that $\frac{\pi_1 - \pi_2}{n} > \pi_E$, no platform gains by deviating and selling to the entrant. No platform gains by demanding a higher t_i , as that would induce the incumbent to reject all offers.

For the “only if” part, suppose for contradiction that $n > \bar{n}$ and there exists an equilibrium where the incumbent wins all the ads. Clearly, the entrant is paying zero to all platforms. Let (t_1, \dots, t_n)

be the transfers that the incumbent is paying in equilibrium to the n platforms. Note that platform i could switch to $(E, \pi_E - \varepsilon)$, where ε is an arbitrarily small positive number. To guarantee that such deviation is not profitable, it must be that $t_i \geq \pi_E$. This implies that the total amount of transfers paid by the incumbent is $n\pi_E$. But that would be greater than the incumbent's benefit $\pi_1 - \pi_2$ if $n > \bar{n} = \frac{\pi_1 - \pi_2}{\pi_E}$, yielding a contradiction.³⁴

Proposition 11

It is easy to see that if the condition for monopolization ($n < \bar{n}$) is satisfied the corresponding TUG has a unique stable partition: $\{\{I, 1, \dots, n\}, \{E\}\}$. For instance, this is sustained by the imputation $x_i = \pi_E$ for every platform $i = 1, \dots, n$, $x_I = \pi_1 - n\pi_E$ for the incumbent, and $x_E = 0$ for the entrant, which yields the same payoff vector as in the non-cooperative equilibrium. The incumbent forms a coalition with all platforms and offers each of them enough money to defend against a coalition of a subset of platforms with the entrant, which can achieve at most π_E .

One can also see that if the monopolization condition fails ($n > \bar{n}$), there exists no stable partition. The partition above, $\{\{I, 1, \dots, n\}, \{E\}\}$, is no longer stable because $v_{E,1,\dots,N}$ is not large enough to guarantee at least π_2 to the incumbent and π_E to each platform. No coalition of the form $\{E, X\}$ is stable because: (i) If the imputation to X is zero, there is a deviation to $\{I, 1, \dots, n\}$; (ii) If the imputation to X is positive, there is a deviation to $\{E, \bar{X}\}$ (because in turn any stable coalition $\{I, \bar{X}\}$ must have a zero imputation to \bar{X}).

Proposition 12

“Only If”. In a monopoly equilibrium nk ads are bought by q incumbents. There must be an incumbent who in equilibrium buys at least $K = \lceil nk/q \rceil$ ads. An entrant who buys exactly one ad receives a payoff $p\pi_E$. Therefore, the bid on every ad must be at least $p\pi_E$. In equilibrium, an incumbent who buys K ads receives payoff $\pi_H - Kp\pi_E$. If the incumbent deviates by offering zero on all ads, her payoff would become $(1 - P_K)\pi_H + P_K\pi_L$. The deviation is profitable if

$$p\pi_E > \frac{P_K(\pi_H - \pi_L)}{K}.$$

“If”. We construct the following equilibrium. Some incumbents bid $p\pi_E$ on $K = \lceil nk/q \rceil$ ads, while the remaining incumbents buy $\lceil nk/q \rceil - 1$ ads (the number of incumbents who buy K ads is $nk - q(K - 1)$). Every ad receives a bid by exactly one incumbent and at least one entrant.³⁵

³⁴Note the set of equilibria is non-empty because there is a pure-strategy subgame perfect equilibrium where every platform offers $(E, 0)$. No platform has a strict incentive to deviate.

³⁵For instance, with $n = 3$ platforms, $k = 4$ ad slots, and $q = 7$ incumbents, it is $K = \lceil 12/7 \rceil = 2$. Hence $nk - q(K - 1) = 5$ incumbents buy 2 ads each, and the remaining 2 incumbents buy 1 ad each.

No entrant has a profitable deviation because $p\pi_E$ is the additional payoff they get from buying one ad and buying more than one ad generates a lower payoff per ad.

We check that no incumbent has a profitable deviation. Clearly no incumbent benefits by increasing her bid or bidding on more ads. If an incumbent bids on $K' < \tilde{K}$ ads instead of $\tilde{K} \in \{K, K - 1\}$ (or reduces her bid on the ads she is buying in equilibrium), entrants would win those auctions, and the incumbent's payoff would become $(1 - P_{K'})\pi_H + P_{K'}\pi_L - K'p\pi_E$. Thus a deviation to K' is profitable if

$$(\tilde{K} - K')p\pi_E - P_{\tilde{K}-K'}(\pi_H - \pi_L) > 0,$$

namely

$$p\pi_E - \frac{P_{\tilde{K}-K'}(\pi_H - \pi_L)}{\tilde{K} - K'} > 0.$$

Note that because $P_{\tilde{K}-K'}$ exhibits decreasing differences it is then

$$\arg \min_{K'} \frac{P_{\tilde{K}-K'}}{\tilde{K} - K'} = 0.$$

Thus, if the inequality holds for some K' it also holds for $K' = 0$ and the necessary and sufficient condition for the deviation is

$$p\pi_E - \frac{P_{\tilde{K}}(\pi_H - \pi_L)}{\tilde{K}} > 0.$$

As $\tilde{K} \in \{K, K - 1\}$, by a similar argument, the necessary and sufficient condition for a deviation for some incumbent is

$$p\pi_E - \frac{P_K(\pi_H - \pi_L)}{K} > 0,$$

which yields the statement