Attention Oligopoly*

Andrea Prat† Tommaso Valletti‡

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

Digital platforms have brought innovative benefits for users and contributed to the economy by opening new business opportunities and facilitating trading. They cover a wide range of daily activities including online marketplaces, social networking services, search engines, or application stores. A few large platforms increasingly act as gatekeepers between business users and end users and enjoy an entrenched and durable position.

In this paper we concentrate on those digital platforms whose business model is funded by advertising. Facebook and Google come immediately to mind: they provide free services to users (e-mails, searches, videos, social networking) in return for information on users' interests and behavior, to create more focused advertising to those very same users. Online advertising is a large and fairly concentrated market. Global spending on internet advertising is predicted to rise above $400 billion in

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

†Columbia Business School, Columbia University, andrea.prat@columbia.edu

‡Imperial College Business School, Imperial College London, t.valletti@imperial.ac.uk

2021 and will constitute over half of total ad spending. Two firms – Google and Facebook – command at least 60% of the market in the US (and over 80% in the UK).[1]

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 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.[2]

New and detailed evidence is emerging on the behavior of platform firms such as Google and Facebook, through an unprecedented number of antitrust lawsuits filed in the second half of 2020 by the DoJ, the FTC, State Attorney Generals, the US Congress, as well as enforcers in Europe and Australia. These high-profile probes have just begun and are expected to be fiercely debated, so it is too early to make an assessment. We just note two features. First, there seems to be a growing agreement that these companies are indeed attention brokers capable of deploying hyper-targeted ads, in contrast with other market players.[3] Second, past acquisitions that went unchallenged at the time, such as the $1 billion acquisition of Instagram by Facebook in 2012 and the $19 billion acquisition of WhatsApp also by Facebook in 2014, are being reconsidered also in light of the impact they had in the market for hyper-targeted ads.[4]

[1] See the 2020 Market study on online platforms and digital advertising conducted by the UK's CMA, available at https://assets.publishing.service.gov.uk/media/5fa557668fa8f5788db46efc/Final_report_Digital_ALT_TEXT.pdf

[2] Our model also allows for the presence of traditional mass media that sell non-targeted ads alongside attention brokers; see Section 6.

[3] See, e.g., the 2020 complaint from the State AGs (https://ag.ny.gov/sites/default/files/facebook_complaint_12.9.2020.pdf) or the 2020 FTC lawsuit (https://www.ftc.gov/system/files/documents/cases/1910134fbcomplaint.pdf).

[4] Quoting from the 2020 FTC lawsuit:

We report a related example that, though not perfect, shows that our framework can find an application also beyond these large and complex cases involving Big Tech companies: online real-estate portals. These are very specialized markets, and are potential "attention brokers" in a specific market when it comes to people searching for a property. They do not sell homes, or even real-estate agent services — rather they sell their search platform to users, and users' attention to advertisers. These portals know specific information about individual preferences over characteristics of the property (e.g., rooms, location, amenities, transport, and so forth) and they further learn from the clicks and views of online users.

In 2014, Zillow and Trulia announced plans to combine into Zillow Group (ZG).[5] At the time, the two firms were the largest and second-largest online real-estate portals, respectively. The FTC cleared the deal, without condition, after a review. According to Newman (2017), the combined firm then began to act like an "attention bottleneck". Prior to the merger, the search operators delivered neutral lists of realtors along with property listings. After the merger, ZG instituted the "Premier Agent" program that allowed realtors to show first and prominently in the search results, for a premium. Per the company's executives, ZG wanted to pursue a strategy of steering users to "Premier" real-estate agents. But by steering real-world buyers to certain favored suppliers, the platform also increased power in a separate — though related — market: that of local realtors, which has progressively become more concentrated too. The link between concentration among digital brokers and a corresponding concentration in a downstream market, is a central finding our paper.

Against this backdrop, how does market power among attention brokers affect consumers, even in zero-price platform markets? This paper explores one possible mechanism: increased concentration among attention brokers, achieved via a platform merger, can lead to reduced entry – and hence higher prices and less product variety – in retail product industries.

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

¶165. Competing personal social networking providers would have been close competitors of Facebook Blue in the supply of advertising. This is because they would have been able to offer the distinctive advertising features described above that distinguish social advertising from other forms of display advertising, search advertising, and "offline" advertising. Instagram and WhatsApp, in particular, were each well-situated to develop into meaningful competitive constraints on Facebook Blue in the sale of advertising. Instagram's founders explicitly planned to develop advertising offerings to monetize the Instagram personal social network. And an independent WhatsApp that developed a personal social networking offering would have had incentives to monetize it by offering advertising. Competing social networks may also have explored and developed alternative advertising models that consumers and advertisers preferred.

¶166. Facebook's anticompetitive conduct to maintain its personal social networking monopoly therefore also has neutralized, suppressed, and deterred competition for the sale of targeted advertising, and deprived advertisers of the benefits of additional competition.

[5] We thank John Newman for bringing this case to our attention.

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.

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 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.

Preview of the results. The paper is structured as follows. Section 2 introduces a parsimonious model, which allows us to illustrate the intuition behind the results. An online appendix contains a more general setting. There are 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.[6]

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 2). Consumer welfare is shown to be a decreasing function of a concentration index that aggregates information about the number of platforms individual consumers use. Welfare is 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 4). 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. Finally, Section 7 concludes.[7]

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

[6] We 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).

[7] In the online appendix we conduct robustness checks, considering alternative ad selling specifications as well as a model with multiple incumbents, multiple entrants, multiple ads per platform, and imperfectly informative ads.

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.

Finally, our paper provides an explanation for a somewhat puzzling phenomenon observed by the empirical literature on online advertising. Companies spend large sums for brand keyword ads: a famous brand (e.g., Coca Cola) advertises for keywords containing their brand, which is of course well known, despite organic links would be likely to appear on top of the page (e.g., www.coca-cola.com) without having to pay for it. 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 as a defensive strategy: 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]

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

[8] AdWords (run by Google) is also an environment where bids are hyper-targeted that would fit our setting. Levin and Milgrom (2010) discuss conflation and the move toward increasingly finer and finer online ad targeting. See also Athey and Nekipelov (2014) and Decarolis and Rovigatti (2019).

[9] In 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).

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, 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). 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.[10] 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).[11] 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

[10] In the two-sided markets literature, because a platform's profits from advertising are roughly proportional to "eyeballs", 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.

[11] Loertscher and Marx (2019) analyze mergers based on auction models, typically applied to procurement settings.

multiple consumers and multiple industries.[12] The present model has also a connection to gatekeeping 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.

There is a connection between the ad selling problem we consider and the shelf space allocation problem (Shaffer, 2005, and Marx and Shaffer, 2010). In those models a retailer sells shelf space to producers who want to get access to consumers. Shaffer (2005) shows that in equilibrium the dominant firm will sometimes exclude a competitive fringe even when welfare would be higher if the fringe obtained distribution. Marx and Shaffer (2010) show that the choice of how much shelf space to allocate is a strategic decision by the retailer. By limiting the amount of shelf space – thus excluding some producers – the retailer can obtain a higher overall profit. Both these phenomena are present in our analysis. However, while that literature looks at a monopoly retailer we consider multiple platforms that compete in selling ads – a necessary extension to be able to discuss the effect of mergers among platforms.

Eliaz and Spiegler (2020) 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.

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

[13] Strictly 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.

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). We are not aware of empirical studies on mergers involving dominant online platforms. Due to the relevance of ongoing cases, a theory literature is just emerging on data driven mergers (e.g. de Corniere and Taylor, 2020; Chen et al., 2020).

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: the Online Appendix removes all of them and considers 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. 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 know what product k each of their users wants and they can run targeted ads that differ from user to user. Targeting distinguishes digital platforms from generalist mass media ads, which we will consider in Section 6.

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.

[14] See 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).

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]

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.

The model presented here is highly stylized. It is an attempt to capture in the most parsimonious way the pre-emption phenomenon we are trying to describe. The advantage is its analytical convenience and a simple closed-form characterization of equilibrium.

In the Online Appendix we probe the robustness of our results to alternative assumptions. Section A1 lists and discusses all our modeling choices. We also prove that our results hold in more general settings. Section A2 shows the main results are robust to alternative specifications of the mechanism space (instead of second-pricee auctions, we consider take-it-or-leave it offers or a class of abstract mechanisms).[17] Section A3 in particular shows that most key results survive if one relaxes all of the following assumptions: one incumbent, one platform, one ad per platform, perfectly informative ads.

3 Consumer Segment Analysis

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).

[15] Under 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 $u_2 + \pi_2 + \pi_E$.

[16] All payoffs are expressed per consumer. If there are multiple consumers interested in the same product, aggregate values are obtained by integrating over those consumers. For instance, the total profit created by a mass s of consumers that generate individual profit π is $s\pi$, etc.

[17] One of the implicit assumptions we have made is that platforms do not bundle ads across consumer segments. This is not an innocent restriction in certain cases, and we discuss it in Section A1.

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.
- 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*

$$\bar{n} = \frac{\pi_1 - \pi_2}{\pi_E}. \tag{1}$$

*(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.*[18]

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

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

[18] To 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.

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.[19]

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 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

[19] The 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.

Figure 1

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

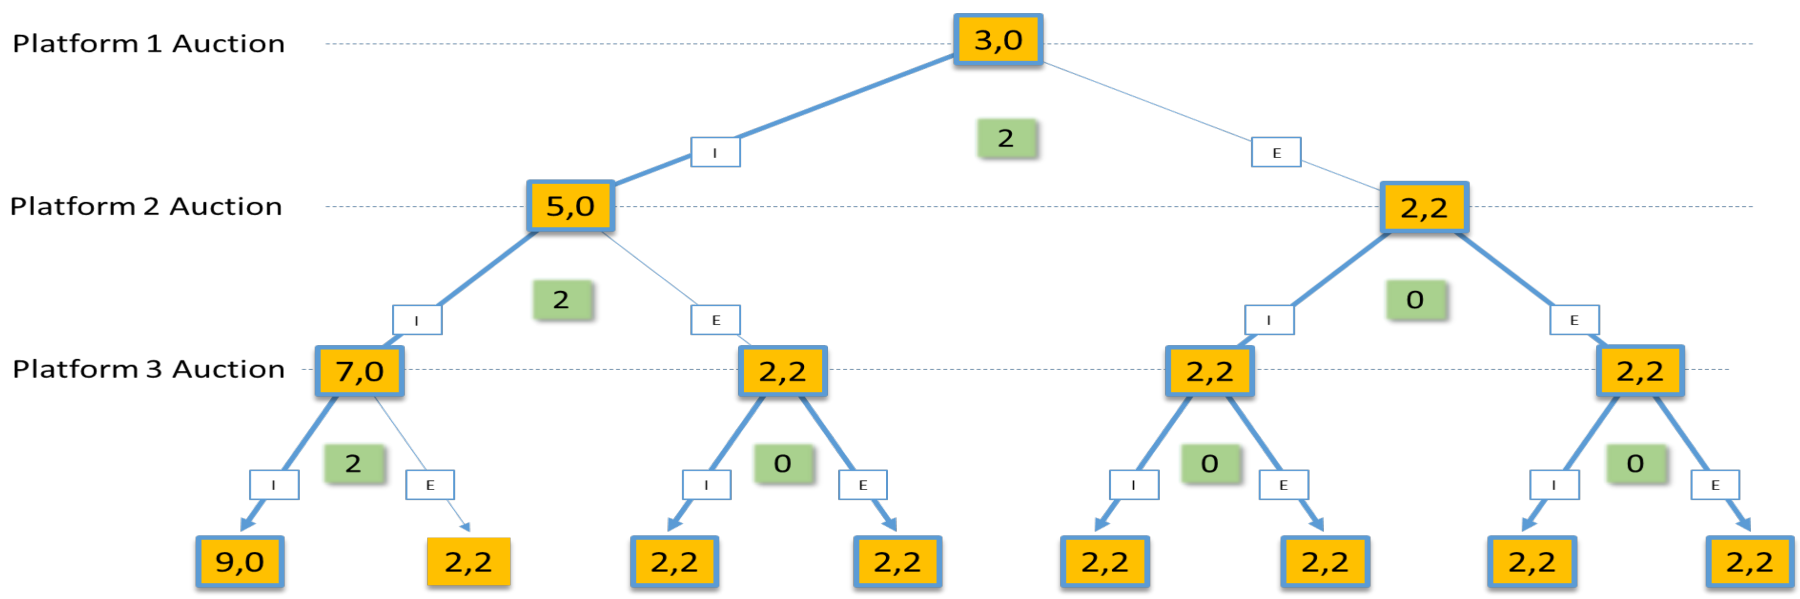

Example 2: $\pi_1 = 7$, $\pi_2 = \pi_E = 2 \rightarrow$ Entry with Positive Platform Profit

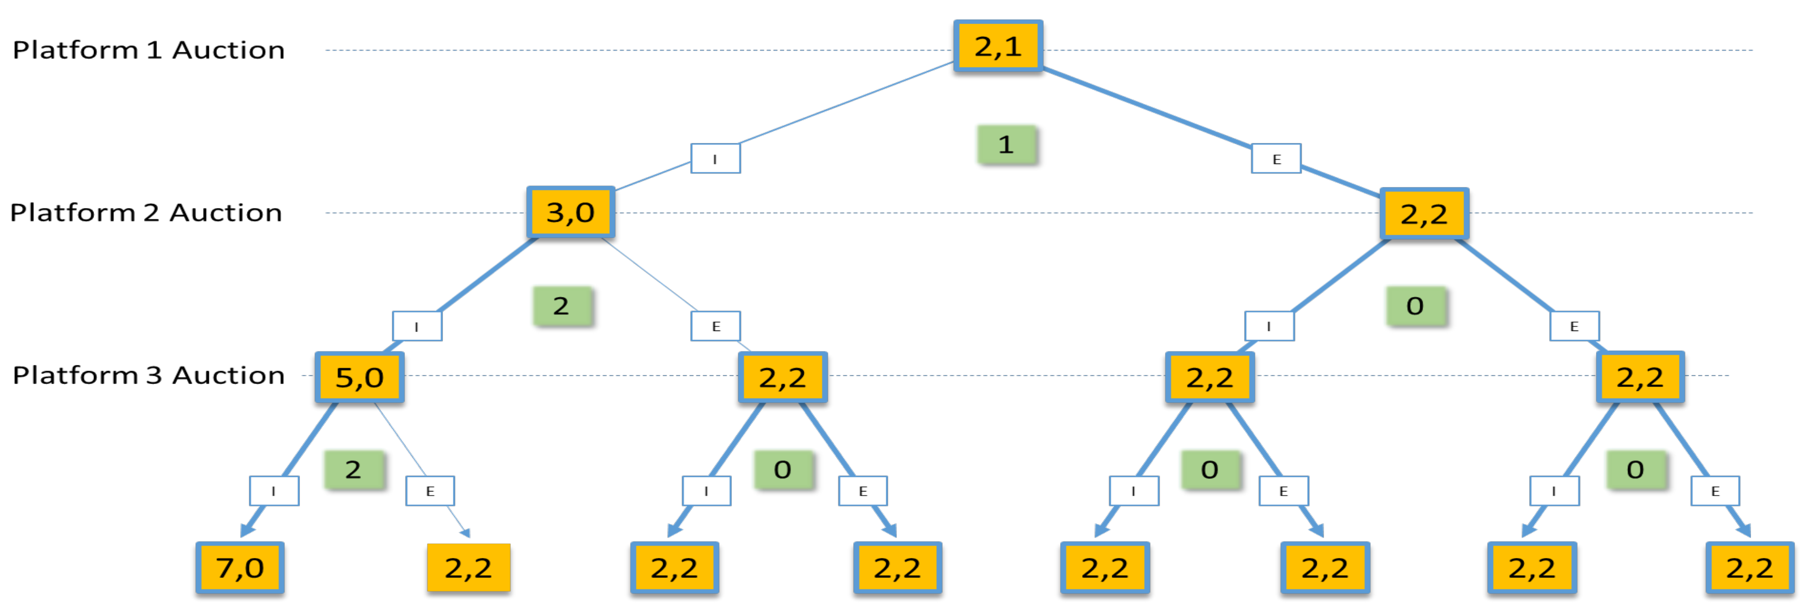

Example 3: $\pi_1 = 5$, $\pi_2 = \pi_E = 2 \rightarrow$ Entry with Zero Platform Profit

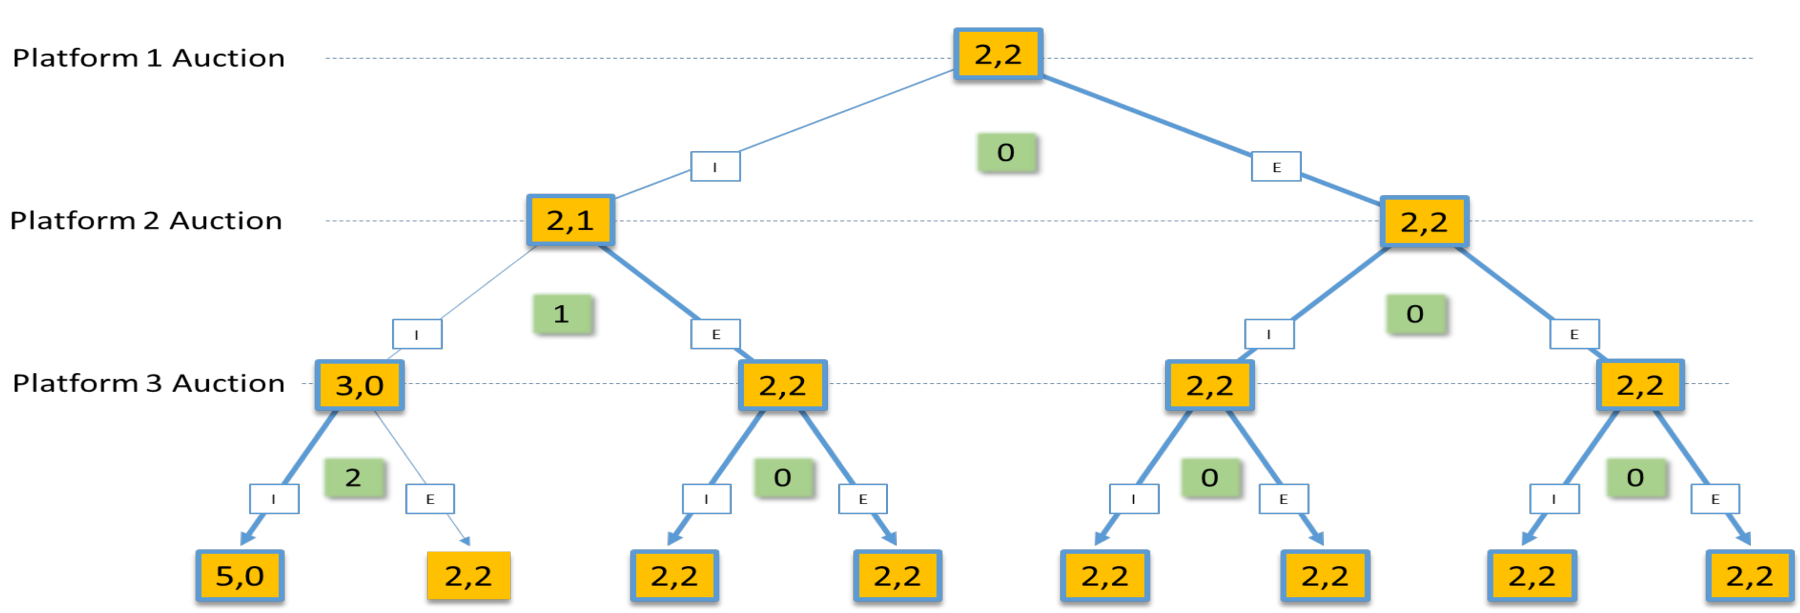

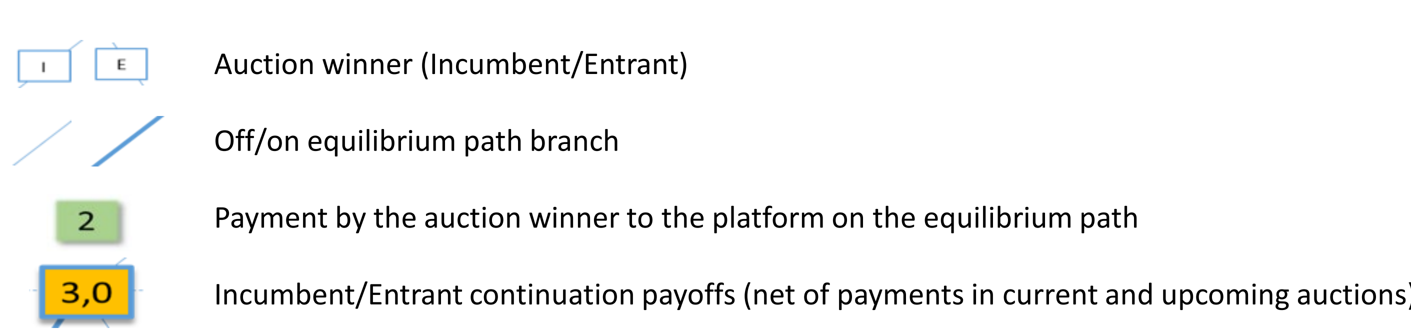

worth more than duopoly to the two firms in aggregate, and the incumbent must pay a substantial part of her rent to the platforms in order to defend her monopoly.

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.

There are two remarks on Proposition 1:

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

Remark 2 *If condition (1a) fails, total platform profit is bounded above by*

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

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).[20]

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 2 *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

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

[20] The objective of Step 1 is to add an element of variability so that the pre-emption condition in Proposition 1 is sometimes satisfied and sometimes not. We achieve this goal by randomizing one of the parameters (π_E); we could also have randomized other parameters or all parameters.

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.

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.[21]

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.

[21] As we further discuss in the Online Appendix, note that the results so far 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. This knowledge instead matters when it comes to mergers. We observe that Facebook acquired in 2013 the data analytics firm Onavo and then used its analytics platform to monitor competitors (including individual usage of several apps). This influenced Facebook to make various business decisions and acquisitions, including its acquisition of WhatsApp.

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 3 *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 3 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 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 2 and 3.

[22] If 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)

[23] Notice 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.

Proposition 4 *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

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

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

Proposition 4, 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]

[24] While 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 acqui-

As we highlight in the discussion of modelling choices (Section A1 of the online appendix), 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. There are four types of consumers: those who do not use either platform ($J = \varnothing$), 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 5 *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

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

sion 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).

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

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

Figure 2 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.

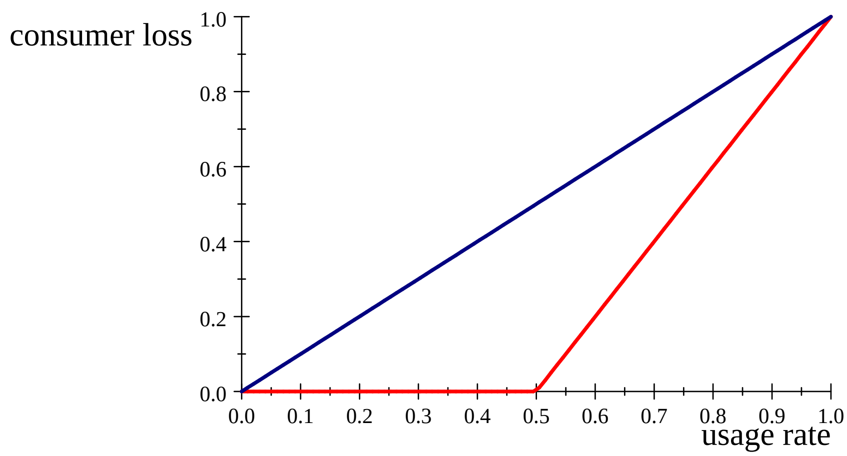

Figure 2. Bounds to inaccuracy when employing usage shares instead of actual overlaps

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 5. 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.

[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.

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]

	(1)	(2)	(3)	(4)	(5)
Merging Pair	*Lower Bound*	*Upper Bound*	*Pairwise Overlap*	*Three-way Overlap*	*Welfare Effect*
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 2 included only one advertising channel: targeted platform ads. We now allow also for the realistic possibility of non-targeted mass media ads and 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 what we call "niche" industries.

Traditional media is modeled in a simple way. Alongside platforms, there are also 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. Let also denote as γ_k the share of consumers that want product k. Industry size will play a role in determining whether industry k advertises on social media or mass media. Without loss of generality assume that k is ordered so that γ_k is nonincreasing in k.

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

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

– 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 6 *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 the number of platforms is not too high ($n < \bar{n} + 1 - \frac{a}{\gamma_k \pi_E}$), 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.

In Proposition 6, entry can be achieved in two way. First, if the industry is sufficiently mainstream and platform concentration is relatively high (large γ_k and low n), the entrant can by-pass the platform bottleneck by buying a non-targeted ad. Second, if the number n of platforms is instead sufficiently large, there is competition among them and targeted ads are cheap; as a consequence 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 rather 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.

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 7 *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 7 shows that the welfare costs are concentrated in niche industries. Mainstream industries are shielded because they find it economical to use non-targeted advertising.[28]

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 4). 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, this offset is partial because it can only affect mainstream industries ($\gamma_k > a/\pi_E$).

We have modelled traditional media in a rather coarse way, via the exogenous price a of ads, which still captures the idea that traditional media can potentially give access to larger audiences but is also more expensive. A natural follow up question, which we leave for further research, would be to endogenize the price of non-targeted ads, which may indeed change as a consequence of a merger among digital platforms.

7 Conclusions

Online digital platforms have created markets, letting demand meet supply in ways that were not available before. Network externalities have also brought concentration and market power, allowing the possibility of rents to be generated and extracted by very few players. This tension is reflected in a flurry of antitrust cases and regulatory proposals that are ongoing in the U.S. and worldwide.

In this paper, we propose a novel analysis of advertising-funded digital platforms. We formalize a theory of attention brokers. Attention brokers attract individual users offering free access to some services, learn about them and keep them engaged on their platform, and then monetize by selling access to the attention of those individuals, typically via hyper-targeted ads. In our model, ads are bought by producers in retail markets who ultimately sell products to the very same individuals. We study the link between market structure in attention markets, and market structure in product markets.

Platforms may be tempted to exploit the attention bottleneck to their advantage, to the possible detriment of consumers. We argue that this situation typically happens when online platforms are concentrated, and they manage to orchestrate prices for ads that select incumbent retail producers to the detriment of entrants. Despite the zero-price paid to the platform by users, they are hurt as they end up having less product choice and paying higher product prices.

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 (local)

[28] Procter & Gamble famously scaled back, in 2016, its targeted advertising on Facebook, redirecting its budget on more traditional media, suggesting that, for large companies, a wider marketing strategy can be more successful (see https://www.wsj.com/articles/p-g-to-scale-back-targeted-facebook-ads-1470760949).

product markets for whom targeted ads are the ideal channel to reach consumers. This is a theoretical prediction of our model that leads itself to an empirical test, with suitable disaggregated data.

One contribution of our paper is therefore to see the effects of online platform competition through the impact it has on product markets. Regulators, at least in the past, were not on the lookout for these types of harms. We believe, however, that the harm potentially highlighted by our model is cognizable under current antitrust principles. This assessment does not require a radical departure from existing competition policy, rather to correctly apply first-order economic principles to these markets. In an application to mergers, we discuss that 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 in retail markets also more likely. In fact we show that, with many competing platforms, it becomes very expensive for retail product incumbents to bid out entrants. Entrants get to be known, consumers are typically well off and marginal mergers would not matter.

A problematic merger is instead one between concentrated online platforms with overlapping users. We discuss 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.

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] Anderson, S. and S. Coate (2005), “Market Provision of Broadcasting: A Welfare Analysis”, Review of Economic Studies, 72, 947-972.

[4] Anderson, S. and J. McLaren (2012), “Media Mergers and Media Bias with Rational Consumers”, Journal of the European Economic Association, 10, 831-859.

[5] Anderson, S., O. Foros and H. Kind (2018), “Competition for Advertisers and for Viewers in Media Markets”, Economic Journal, 128, 34-54.

[6] Armstrong, M. (2006), “Competition in Two-sided Markets”, RAND Journal of Economics, 37, 668-91.

[7] 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.

[8] Athey, S. and D. Nekipelov (2014), "A Structural Model of Sponsored Search Advertising Auctions", mimeo.

[9] 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.

[10] Bagwell, K. (2007), "The Economics Analysis of Advertising", in Handbook of Industrial Organization, vol. 3, Elsevier (Amsterdam), 1701-1844.

[11] Blake, T., C. Nosko and S. Tadelis (2015), "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment", Econometrica, 83, 155-174.

[12] Bernheim, D. and M. Whinston (1986), "Common Agency", Econometrica, 54, 923-942.

[13] 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.

[14] Chen, Z., C. Choe, J. Cong and N. Matsushima (2020), "Data-Driven Mergers and Personalization", mimeo.

[15] Decarolis, F. and G. Rovigatti (2019), "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising", American Economic Review, forthcoming.

[16] de Corniere, A. and R. De Nijs (2016), "Online Advertising and Privacy", RAND Journal of Economics, 47, 48-72.

[17] de Corniere, A. and G. Taylor (2020), "Data and Competition", mimeo.

[18] Diermeier, D. and R. B. Myerson (1999) "Bicameralism and Its Consequences for the Internal Organization of Legislatures", American Economic Review, 89, 5, 1182-1196.

[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 (2020), "Incentive-Compatible Advertising on Non-Retail Platforms", Rand Journal of Economics, 51, 2, 323-345.

[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] Jeziorski, P. (2014), “Effects of Mergers in Two-sided Markets: The U.S. Radio Industry”, American Economic Journal: Microeconomics, 6, 35-73.

[24] Kaldor, N. (1950), “The Economic Aspects of Advertising”, Review of Economic Studies, 18, 1-27.

[25] Levin, J. and P. Milgrom (2010), “Online Advertising: Heterogeneity and Conflation in Market Design”, American Economic Review P&P, 100, 603-607.

[26] Loertscher, S. and L. M. Marx (2019), “Merger Review for Markets with Buyer Power”, Journal of Political Economy, 127, 9, 2967-3017.

[27] Marx, L. M.. and G. Shaffer (2010), “Slotting Allowances and Scarce Shelf Space”, Journal of Economics, Management and Strategy, 19, 3, 575-603.

[28] 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.

[29] Prat, A. (2018), “Media Power”, Journal of Political Economy, 126, 4, 1747-1783.

[30] Prat, A. and A. Rustichini (2003), “Games Played Through Agents”, Econometrica, 71, 989-1026.

[31] Rey, P. and J. Tirole (2007), “A Primer on Foreclosure”, in Handbook of Industrial Organization Vol. 3, Elsevier (Amsterdam).

[32] Rochet, J.-C. and J. Tirole (2003), “Platform Competition in Two-sided Markets”, Journal of the European Economic Association, 1, 990-1029.

[33] Schmalensee, R. (1978), “Entry Deterrence in the Ready-to-Eat Breakfast Cereal Industry”, Bell Journal of Economics, 9, 305-327.

[34] Simonov, A., C. Nosko and J. Rao (2018), “Competition and Crowd-Out for Brand Keywords in Sponsored Search”, Marketing Science, 37, 2, 200-215.

[35] Shaffer, G. (2005), “Slotting Allowances and Optimal Product Variety”. Advances in Economic Analysis & Policy, 5, 1, Article 3.

[36] Telser, L. (1964), “Advertising and Competition”, Journal of Political Economy, 72, 537-562.

[37] Whinston, M. (2006). Lectures on Antitrust Economics, MIT Press (Boston, MA).

[38] 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

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

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.

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

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

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

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

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

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

Proposition 2

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 2, where $a = u_2$ and $b = \frac{(u_2-u_1)(\pi_1-\pi_2)}{M} > 0$.

Proposition 3

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 4

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

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

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

Proposition 6

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 $\left(\pi_1-\pi_2-(n-1)\,\pi_E\right)\gamma_k$ 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, the entrant prefers a non-targeted ad if

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

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 7

By Proposition 6, industry k is monopolized if and only if $\gamma_k<a/\pi_E$ and

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

Thus, the industry is monopolized if

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

and is in the entry equilibrium otherwise.

Adapting Proposition 2, we have

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