Vertical Integration and Demand Steering with Information Frictions: Evidence From the Online Advertising Industry

Callum Shaw*

January 11, 2023

Abstract

I develop and estimate a structural model of the online display advertising intermediary market. I find substantial evidence that vertically integrated entities benefit from information frictions which restrict information transfer between non-integrated firms, giving integrated downstream firms a competitive advantage. A counterfactual simulation reveals that the market share of the downstream firm which is vertically integrated with the most dominant upstream firm (Google) would be 1.66 percentage points lower if information frictions were eliminated. My findings imply that privacy protection regulation like GDPR which inhibits the ease of data transfer between firms can have anti-competitive effects.

*Department of Economics, London School of Economics and Political Science, London, UK

1 Introduction

While the negative welfare effects of horizontal mergers are well documented in the literature on industry structure, there has historically been less agreement about the possible anti-competitive effects of vertical mergers.

A strand of the literature has argued in favour of allowing vertical integration on the basis that it eliminates the problem of double marginalisation (see Bork (1978) and Spengler (1950)). This school of thought is challenged by a literature on demand steering and market foreclosure (see Hart et al. (1990)), which proffers that vertical integration between an upstream and a downstream firm can result in an incentive for the upstream firm to favour its downstream partner in order to give that partner an advantage over its rivals in the downstream market (see, for example, Salop and Scheffman (1983)). Because of this, the welfare effects of vertical integration are more difficult to pin down than those of horizontal integration. Recently, progress has been made with empirical analysis tailored to the idiosyncrasies of specific industries, such as healthcare (Cuesta et al. (2019)) and television (Crawford et al. (2018)).

Vertical integration in the market for online advertising intermediation is quickly gaining more regulatory attention as the industry grows[1]. This industry comprises the tech firms which act as intermediaries (in several different layers) between advertisers purchasing online advertising space and publishers selling that advertising space.

Reports by the CMA (CMA (2019)) and the European commission (Jeon (2021)) have accused Google (a dominant firm in online search and in online ad intermediation) of anti competitive behaviour in this sector, and direct antitrust investigations into this conduct have been launched by authorities in the UK (Lomas (2022)), the EU (Lomas (2021)), and

[1]In 2021 the portion of advertising spend devoted to online channels was around 65% in both the US and the UK (Statista (2021); Ahuja (2020)). This is a significant item on balance sheets; digital advertising cost UK companies £14 billion in 2019 - around £500 per UK household, each of which sees these costs reflected in the prices of goods and services in the economy (CMA (2019)). Two of the largest ten companies in the world (by market capitalisation) make the majority of their revenue from online advertising: Alphabet (the parent company of Google, with a market cap of $1.46T in 2021) and Meta (the parent company of Facebook, with a market cap of $0.55T in 2021) (8marketcap (2022))

the US (O'Toole (2020)). One of the most significant policy concerns is the possibility that a firm like Google can use its dominance in one layer of the vertical market to favour its own partnered intermediaries in a different layer of the market.

I develop a model of the online display advertising industry and estimate it to empirically assess whether upstream intermediaries (ad exchanges) give their downstream integrated partners (demand side platforms) a competitive advantage by providing them an information advantage which they do not provide to their rivals. While previous work on vertical integration focuses mainly on how upstream prices are used to steer demand in a downstream market by raising rivals' costs, my model highlights a mechanism through which information frictions between upstream firms and downstream firms can work to steer demand in the downstream market by lowering the quality of service of downstream competitors. A unique proprietary data set is used to estimate the model, which has the advantage of containing information on the entire vertical market, from publisher, through intermediaries (exchanges and demand side platforms), to advertiser. Data on individual advertiser choice of demand side platform allows me to estimate my model by maximum likelihood estimation, in a manner similar to a "micro-BLP" approach (see Berry et al. (2004)).

I find substantial evidence that downstream firms which are integrated with dominant upstream firms benefit from the demand steering effect described above. The fact that this demand steering is facilitated by information frictions has implications for the consequences of policy which affects the ease of information transfer between firms, such as the General Data Protection Regulation in the EU, or the California Consumer Privacy Act in the US.

Related Literature

My paper is primarily a contribution to the literature on demand steering effects that can result from vertical integration. Relevant recent work includes Crawford et al. (2018), who investigate the effects of foreclosure and raising rival's costs in multichannel television markets, and Cuesta et al. (2019), who look at similar effects in the market for hospitals and insurers.

Scott Morton and Athey (2021) define the notion of “platform annexation”, whereby a platform firm has an incentive to “annex” an adjacent downstream industry in order to gain control of the access to the platform industry, thereby increasing its market dominance. The focus of this paper is related to the contribution by Scott Morton and Athey, as it studies a similar phenomenon in the same industry, but in a different layer of the vertical market.

This research also contributes to the literature on the value of data and its means of transaction online. See, for example, Alcobendas et al. (2021). My model departs from their conceptualisation in that I treat the transfer of a cookie[2] primarily as a transfer of information, rather than seeing it as something which directly raises the valuation of an advertising impression to an advertiser.

The remaining sections of the paper are organised as follows. Section 2 gives a more detailed description of the specific industry I study. In section 3 I describe my data and present motivating evidence for my modelling approach. Section 4 lays out my structural model of the display advertising industry, and analyses it to determine some relationships between its key elements. Section 5 describes my approach to identifying and estimating the parameters of my model. In section 6 I present my results and their implications, and in section 7 I conclude.

2 Industrial Setting

The display advertising industry is the subset of digital advertising comprising the buying and selling of all display advertising inventory[3]. Total spend on display advertising in the

[2]A piece of data used to track internet users

[3]Traditionally, the term “display” was used to refer to banner ads - the rectangular graphic adverts that appear at the top, bottom or along the side of web pages. As the market has developed, the term has expanded to encompass video advertising (the video adverts that appear before and after video content on sites like Youtube) and some elements of social media advertising. With the first display ads appearing as early as 1994 (LaFrance (2017)), it is normally considered the earliest form of internet advertising, and was

UK was around £5 billion in 2019, making up 36% of total digital advertising spend that year (CMA (2019)); the majority of the remainder is on PPC (Search engine[4]) advertising, and non-display social media advertising (including influencer advertising, etc.).

On the demand side of the display advertising market are advertisers who want to purchase advertising space in order to inform consumers about their products. On the supply side there are publishers and other content producers who have consumers viewing their content, and who therefore have a supply of "advertising inventory" (advertising space) on their sites which they wish to sell to advertisers to generate revenue. The unit of exchange is an "impression", which is an instance of one ad appearing on one person's screen when they load a new web page. When an advertiser purchases one impression, it will be his ad being shown in that particular advertising space that one time. There are various issues that arise with treating an impression as the unit for exchange. An advert shown to a web user for a second time is not necessarily as valuable to the advertiser as when it was viewed the first time, for instance. I will move on to discuss these concerns and how I integrate them into my model, but for now we can approximately consider the market for impressions to be a market for web-user attention, with publishers generating impressions (by drawing people to view their content) on one side, and advertisers buying these impressions on the other side.

A network of intermediaries sits between advertisers and publishers, known as the "ad tech" industry [5]. Technology has developed to the point where the majority of transactions of impressions are handled algorithmically by computer programs. The process I describe below is known as Real Time Bidding (RTB) (CMA (2019)).

The structure of the market is illustrated in figure 1. Advertisers delegate the buying

traditionally the main mechanism non-paid content producers used to generate revenue.

[4] PPC stands for "Pay per Click", and is the term normally used to refer to paid search engine advertising

[5] The buying and selling of impressions was traditionally carried out via direct transactions between advertisers and publishers, whereby an advertiser would pay a publisher some fixed fee to purchase a certain number of impressions to be shown over a certain time period. The network of intermediation that had developed means advertisers can show their ads across a range of websites without having to handle all of the individual interactions manually, and publishers benefit in a similar fashion by the fact that there they don't have to manually curate relationships with multiple advertisers to encourage competition for the advertising space on their sites.

of impressions to a demand-side platform (a DSP), and publishers delegate the selling of impressions to an exchange [6] [7]. We can consider DSPs as competing against each other for advertisers, and exchanges as competing against each other for publishers (the business model of these firms is normally a percentage fee charged on the spend that goes through them). Suppose a user loads a webpage which has several potential advertising spots on it. Take one of these advertising spots. What follows is an algorithmically controlled transaction which normally occurs in under a second. The publisher sends a message to the exchange to notify it that there is an impression up for sale, along with information about the impression, including the site it is on, and personal data about the user loading the webpage (often referred to as "cookie" data[8]). The exchange then auctions off the impression (via either a first price or second price auction), inviting DSPs to bid on it. Each DSP observes the information about the impression in order to learn its valuation of it, and then bids. The winner of the auction purchases the impression, thereby winning the right to serve its ad in the advertising spot in question. Several of these interactions happen each time any user loads a webpage, and as such billions happen every day.

If we assume all bidders receive the same information about the impression, the economic appeal of the structure of the market should be clear. The DSP's valuation of an impression is based on the characteristics of the advert it wishes to display (on behalf of the advertiser) and how well this matches with the characteristics of the impression that is being auctioned. That is, DSPs are incentivised to show ads to the people most like to respond well to those ads. Both first and second price auctions give rise to Nash equilibria which allocate each impression in an efficient manner, to the bidder with the highest valuation[9]. As such,

[6]Historically, there were multiple layers in the supply-side intermediary ecosystem, such as supply-side platforms (SSPs) which themselves intermediated between publishers and exchanges. Today, however, the delineation between SSPs and exchanges has largely disappeared (Srinivasan (2020)). As such, I abstract away from the distinction and describe exchanges only

[7]For the remainder of the paper, I will sometimes refer to DSPs as "downstream intermediaries" and exchanges as "upstream intermediaries"

[8]There are several different technologies used to transfer data (one of which is cookies), and the distinction between them is important for various issues in the industry, but for the scope of this paper it suffices to treat them as one unified "means of data transfer"

[9]As is shown in Milgrom and Weber (1982)

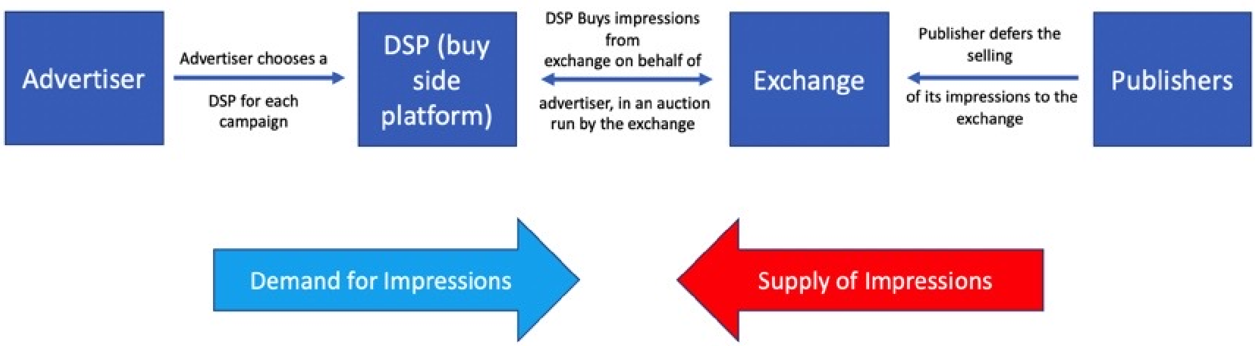

Figure 1: The structure of the display ad intermediary market

advertisers are matched up with the consumers most likely to respond well to those adverts - that is, the consumers who are most likely to want to purchase the item or service being advertised in response to being shown the advert. This should lead to advertising budgets being spent efficiently, minimising the amount wasted on showing adverts to people who will not benefit from seeing them. This is good for both consumers, who are shown fewer irrelevant ads (which, we can posit, causes more disutility than seeing relevant ads) and for advertisers, who generate more value from their advertising for less money. Publishers also benefit, as selling an impression to the DSP with the highest valuation also maximises revenue.

The other important thing to note is that the auction that occurs between exchanges and DSPs should ensure a level playing field across DSPs, and precludes the possibility of demand steering observed in vertical interactions within traditional markets. Even if an exchange is vertically integrated with a DSP, as long as there is no collusion in the auction, there is apparently no way for an exchange to try to favour that DSP over others, in order to attempt to give them an unfair advantage. The auction keeps competition in the DSP market "separated" from competition in the exchange market in the sense that an exchange with market share cannot leverage that dominance to try to increase the market share of a partnered DSP. This result is lost, however, if we consider the possibility of information frictions.

Information transfer is of vital importance to the impression transaction process. In order for the DSP to correctly identify its own valuation for the impression, the transfer of information about the impression from the web-user's computer, to the publisher, to the exchange, to the bidding DSPs all needs to be comprehensive and smooth. A hypothesis with growing support (see CMA (2019), Srinivasan (2020), O'Toole (2020), Jeon (2021), and Scott Morton and Athey (2021)) is that, in the last stage of information transfer between exchange and DSP, some DSPs receive imperfect information about the impression. The transfer of information from exchange to DSP involves encryption and is technically involved but, in short, it is possible for the transfer of information from an exchange to a vertically integrated DSP to be more efficient than it is to other, non integrated DSPs. This can be, for example, because a vertically integrated entity faces fewer legal and technical barriers to transferring information between its different components (exchange and DSP) than do exchanges and DSPs owned by separate entities [10]. Such information frictions can benefit vertically integrated entities at the expense of welfare in the market.

If such information frictions exist, it is possible that the market does not work efficiently. If some DSPs do not learn about their actual valuation for an impression, they will not bid appropriately (or they will drop out of the auction altogether (Srinivasan (2020))), and the impression will not be allocated efficiently to the DSP with the highest valuation. This results in harm to consumers, advertisers and publishers by decreasing the efficiency of the real time bidding process described above, as well as lowering the expected revenue from auctions.

A merged entity can benefit from these information frictions. Advertisers value DSPs which can reliably provide reach (access to advertising inventory from a broad range of websites) and pacing (a reliable, consistent flow of impressions). A DSP which faces information frictions with exchanges can therefore be expected to lose market share (and revenue earned through fees) as these frictions make it less effective at finding and buying impressions on

[10] Regulation intended to protect consumer privacy, like GDPR or the CCA, can restrict the transfer of information between counterparties in this way

behalf of advertisers and lower the quality of service of the DSP. This lowers market share and profits for the DSP. In this way, information frictions can have anti-competitive effects in the market for DSPs.

A vertically integrated entity can thereby benefit from information frictions provided that the increase in profits of the integrated DSP is higher any loss in profit the exchange may incur from these frictions. Further, we would expect that the benefit is higher for an entity which owns a dominant exchange, as having poor information transfer with an exchange is worse for a DSP if that exchange controls a large proportion of the supply of impressions.

Higher market concentration and a lack of competition can lead to increased prices which in turn harm consumers. On the publisher side, higher fees to intermediaries means less money is spent on content. On the advertiser side, higher fees can be passed onto consumers in the form of higher prices for goods and services. Figure 2 is from a CMA report published in 2019, and describes the various ways consumers can be harmed by increased market share in the ad intermediary industry. Through information frictions and the demand steering they can facilitate, a market which exhibits market concentration on the publisher side (and thus incurs the damage caused by this) is more likely to incur further damage on the advertiser side as this market concentration is transferred from the upstream intermediary side to the downstream intermediary side.

There is documented speculation that firms owned by Google do benefit from the information frictions described above. Google owns both exchanges and DSPs, and has historically been amongst the dominant players in both of these markets (CMA (2019)). Industry publications and legal literature (Srinivasan (2020)) have complained not only that Google benefits from information frictions, but also that privacy regulation like GDPR aggravates these information frictions by worsening interoperability between separately owned exchanges and DSPs, thus forcing more demand through Google's vertically integrated structure. As this literature points out, such asymmetric information practices have been banned in financial

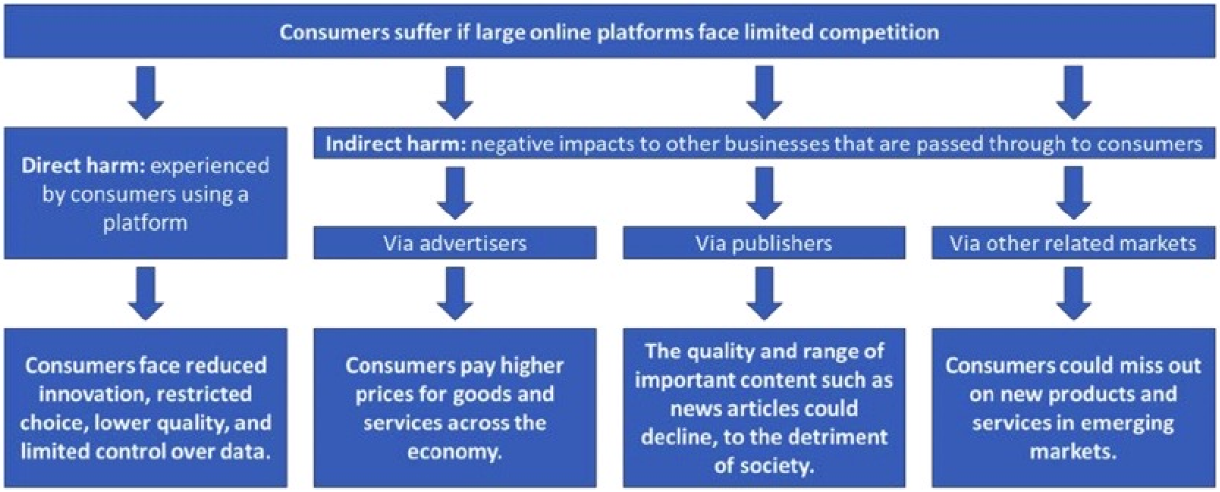

Figure 2: How consumer harm can arise through limited competition in the ad intermediary industry (source: CMA (2019))

markets with similar electronic trading elements, owing to analogous regulatory concerns[11]. The aim of this paper is to quantitatively analyse the extent to which a firm like Google benefits from these information asymmetries.

3 Data and Stylised Facts

I use proprietary data from a large advertising agency. The dataset contains a sample of 200 advertisers represented by this agency or one of its subsidiaries in the US from July 2017 to July 2020 and the UK from January 2018 to July 2020 [12]. The sample of advertisers was selected to be representative of the wider population, and contains around 110 billion impressions. I define a "market" as a country-month.

The data set is highly granular. For each impression purchased, I observe the DSP it was purchased by, the advertiser on who's behalf the DSP purchased the impression, the exchange it was purchased from, and the publisher who owns the website on which the advert is finally shown. Along with this, there is additional information on the type of transaction through

[11]As Srinivasan (2020) points out, one of the guiding principles in regulation of financial markets is that "exchanges must provide traders with fair access to the marketplace, including access to the data transmitted by exchanges"

[12]The UK and the US are the largest display advertising markets the company operates in

which the impression was purchased (this allows me to identify those impressions purchased via the Real Time Bidding (RTB) auction process).

I observe the price paid for each impression (the "media cost"), which is passed through to the advertiser, as well as the "tech cost", which gives the fee charged by the DSP (for its services) to the advertiser over and above the media cost. I do not observe losing bids.

As I observe individual advertiser characteristics and choices of DSP, I can estimate the discrete choice model described in section 4 using a "micro BLP" approach (see Berry et al. (2004)).

In my data I observe two stylised facts which are consistent with the hypothesis described in section 2 and which motivate my modelling approach. First, as figure 3 shows, there is a positive correlation between the market share of Google's exchange (the largest exchange in my data set, by number of impressions) and the market share of Google's DSP.

If information frictions do allow a vertically integrated exchange to steer demand to its integrated DSP, we would expect the demand steering effect to be larger when that exchange has a higher market share. This is because an information friction between a non-integrated DSP and an exchange will have a more detrimental effect on that DSP's service providing if that exchange controls a larger proportion of the supply of the impressions. Figure 3 shows that we do in fact observe this relationship [13].

Second, figure 4 shows that the Google-owned DSP (DBM) purchases a higher proportion of its impressions from Google's exchange than do the other two DSPs in my data set.

As showing an ad impression can have a negative impact on an advertiser (owing to the possibility that the advert is shown next to harmful content, or is shown repeatedly to the same internet user), a DSP will often drop out of an auction where it does not receive information about the impression (Srinivasan (2020)). If there is better information transfer between Google's exchange and its own DSP, we would therefore expect to see the

[13] It should be noted that there are clearly other candidate explanations for this correlation, besides the possibility of information frictions. A vertically integrated entity could, for instance, undergo investment in its brand image which benefits the market share of its DSP and its exchange simultaneously. This possibility is controlled for in my structural estimation

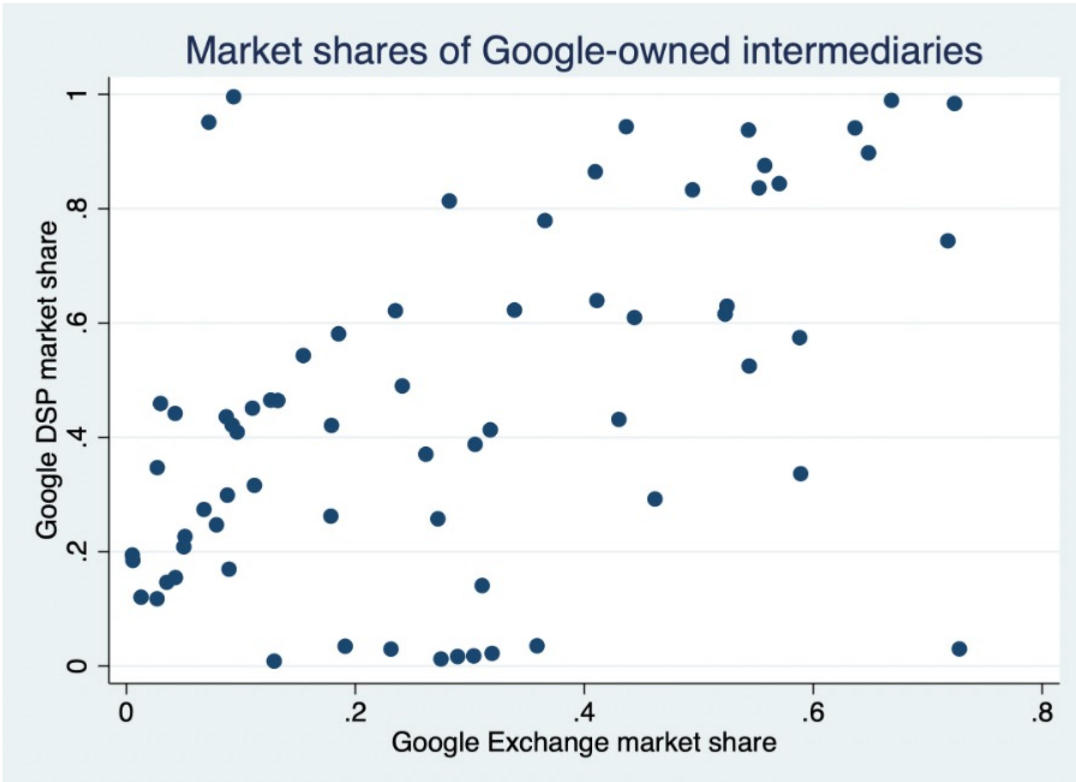

Figure 3: The relationship between Google DSP's market share and Google exchange's market share in the UK and the US from mid 2017 to mid 2020, where market share is measured by number of impressions. Each point is a market, defined as a country-month. Regression Coefficient: 0.70, Standard error: 0.15. Source: Advertising agency data

trend observed in figure 4, that DBM participates in and wins more auctions with its own exchange than do other DSPs.

4 Structural Model

This section presents a model of the display ad intermediary market.

There are N DSPs labelled $d = 1, 2, ..., N$, E exchanges labelled $s = 1, 2, ..., E$, and T markets labelled $t = 1, 2, ..., T$ [14]. In each market t, each advertiser makes a discrete choice of DSP to maximise its utility. Advertiser utility depends on the level of service the DSP can provide. The utility that advertiser a obtains by choosing DSP d in market t is

[14] A market is defined as a geography-time period

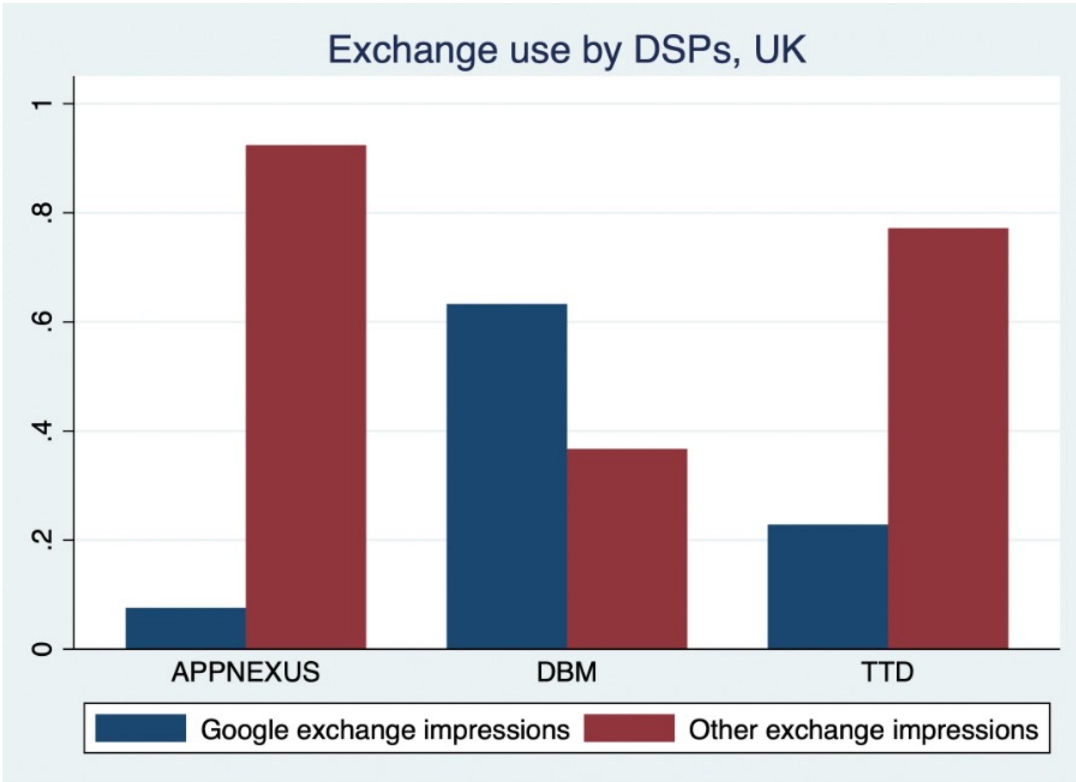

Figure 4: Proportion of impressions purchased from Google's exchange for each DSP in the UK, 2018-2020. Source: Advertising agency data. A similar trend is observed for the US.

$$u_{adt} = \delta_{dt} + \lambda_{adt} + \epsilon_{adt}, \tag{1}$$

$$\delta_{dt} = \beta' x_{dt} + \gamma m_{dt} + \xi_{dt}, \tag{2}$$

$$\lambda_{adt} = \mu_d' z_{at}, \tag{3}$$

where x_{dt} is a vector of observed characteristics of DSP d in market t, ξ_{dt} is the unobserved characteristics of DSP d in market t, and z_{at} is a vector of observed characteristics of advertiser a in market t. β is a vector of coefficients multiplying the elements of x_{dt}, and μ_d is a vector of coefficients specific to DSP d multiplying the elements of z_{at}. ϵ_{adt} is an idiosyncratic shock term, which I assume is i.i.d. extreme value distributed with a variance of $\pi^2/6$, following the standard practice in the literature on discrete choice (see McFadden (1974) and Train (2009)).

m_{dt} is the search friction measure of DSP d in market t, and is a function of the Ex1 vector of information frictions DSP d faces with each exchange i. If DSP d faces large information frictions, in particular with large exchanges, it will be more difficult for DSP d to buy impressions in a timely manner when an advertiser requests them [15]. m_{dt} is calculated by assuming that each DSP d takes part in the following buy process each time it is directed to buy an impression:

Step 1: Randomly draw an exchange. Each exchange i has a probability q_{it} of being drawn, where q_{it} is constant across DSPs

Step 2: Randomly draw a binary "information draw" for that exchange (i), where DSP d has a probability p_{dit} of a good information draw and $1 - p_{dit}$ of a bad information draw.

Step 3(a): If the DSP receives a bad information draw, the DSP does not purchase an impression in this buy process and the buy process finishes

Step 3(b): If the DSP receives a good information draw, it then draws a private valuation v from the distribution with cumulative distribution function $F(.)$, where $F(.)$ is common to all DSPs at all exchanges. The DSP then bids in an auction run by the exchange. If it wins, it receives an impression and the buy process finishes. If it does not win, the DSP does not purchase an impression in this buy process and the buy process finishes. Let the ex-ante probability that the DSP wins the auction (i.e. the probability it wins an impression after it has received its information draw, but before it has received its valuation draw) be r_{dit}

q_{it} is the market share of exchange i in market t, determined exogenously by publisher decisions of which exchange to sell its impressions through. p_{dit} is a measure of the ease of

[15] Advertisers value DSPs which can reliably provide reach and which can quickly source impressions on command to respond to the pacing needs of the advertiser. We can see a higher m_{dt} as something which makes it more difficult for a DSP to search for impressions and provide this service.

information transfer DSP d experiences with exchange i in market t (and so $(1 - p_{dit})$ is a measure of the size of the information friction). A good information draw can be interpreted as a DSP receiving the cookie information associated with the impression being sold. The fact that a DSP withdraws from auctions where it has a bad information draw reflects the fact that impressions can be damaging to an advertiser (the above specification implies $\int v \, dF(v) < 0$).

The probability that a DSP finishes an iteration of the buy process with a bad information draw is $\sum_{i=1}^{E} q_{it}(1 - p_{dit})$. This is a measure of how bad a DSP is at sourcing impressions, and this is how I define the search friction measure m_{dt} for DSP d in market t:

$$m_{dt} \equiv \sum_{i=1}^{E} q_{it}(1 - p_{dit}). \tag{4}$$

Note that m_{dt} is a weighted average of the information friction DSP d faces with all exchanges, where the weight is the market share of each exchange.

Assuming advertiser characteristics are jointly distributed with a cumulative distribution function $G(.)$, the market share of DSP d in market t is

$$P_{dt} = \int \frac{e^{\delta_{dt} + \lambda_{adt}}}{\sum_j e^{\delta_{jt} + \lambda_{ajt}}} \, dG(z_t). \tag{5}$$

I assume that publisher decisions of which exchange to select to sell impressions through are exogenous. An exchange always accepts the highest bid in an auction, there are no reserve prices, and the information frictions (that is, the p_{dit}'s) are determined exogenously. My assumptions on the supply side highlight that mine is a partial equilibrium analysis, a decision motivated by the absence of information on decisions made by exchanges in my data.

Proposition 1 *Provided γ is negative, there is a positive causal link between the market share of an exchange i and that of a DSP d if*

$$\int P_{adt}(1-P_{adt})(1-p_{dit})\,dG(z_t) < \sum_{k\neq d}^{N} \int P_{adt}P_{akt}(1-p_{kit})\,dG(z_t). \tag{6}$$

(Proof in Appendix)

Proposition 1 provides the condition under which an increase in the market share of exchange i causes an increase in the market share of DSP d. This condition is more likely to hold if p_{dit} is high and p_{kit} is low for $k \neq d$ - that is, if there is good information transfer between exchange i and DSP d, but poor information transfer between exchange i and other exchanges k. This result shows that, if γ is negative (a condition which is addressed by the empirical content of my paper), my model generates the positive correlation we observe in figure 3 when an exchange offers stronger information transfer to an integrated DSP and weaker information transfer to non integrated DSPs (which will occur when exchange i is vertically integrated with DSP d, as the Google-owned exchange and DSP are, and there are regulations in place restricting information flow between non-integrated firms).

Proposition 2 *If* $p_{jit} = 1$ *for all* j, i, t *then* $\frac{\partial P_{dt}}{\partial q_{it}} = 0$ *for all* d, i, t.

(Proof in Appendix)

In words, proposition 2 states that, if there were no information frictions, there would be no causal links between the market shares of exchanges and the market shares of DSPs[16].

Vertical Integration

I have not modelled the decisions of publishers nor of exchanges in my analysis, and I will proceed with estimation of demand on this basis. However, it will aid the interpretation of

[16]Note that I have only discussed here the possible *causal* link between exchange market share and DSP market share, and have so far omitted mention of the possibility that there are forces determining market share in both markets which potentially cause a correlation between P_{dt} and q_{it} which is not the result of this causal link. I discuss this in section 5.

my results to briefly analyse the impact of information frictions on exchanges and how they are affected when an exchange vertically integrates with a DSP.

An exchange's profits π_{it} and the vector of information frictions that exchange has with all DSPs ($\boldsymbol{p_{it}}$) are related in the following way[17]:

$$\pi_{it} = \pi_{it}(\boldsymbol{p_{it}}),$$
$$\frac{\partial \pi_{it}}{\partial p_{dit}} > 0 \qquad \forall d.$$

From the fact that p_{dit} is a probability, it follows that

$$\underset{p_{dit}}{\operatorname{argmax}}(\pi_{it}(\boldsymbol{p_{it}})) = 1 \qquad \forall d.$$

That is, the optimal situation for a non vertically integrated exchange is for there to be no information frictions between it and any DSP.

The profits of DSP d in market t are given by

$$\pi_{dt} = f_{dt} P_{dt}(m_{dt}, x_{dt}) M,$$

where M is total DSP market size (measured as spend on media through DSPs) and f_{dt} is the fee charged by DSP d.

Assume an exchange i and a DSP d merge. Now, the manager of the merged entity

[17] An exchange makes money by charging a % fee on the price of each ad impression that is sold through it. Therefore, having either higher traffic going through it, or higher maximum bids in each auction, improves the revenue of an exchange. For my model, I have assumed publisher decisions of which exchange to sell their advertising space through are exogenous, and so, given this assumption, exchange incentives should align with those of publishers. That is, an exchange's only goal should be to maximise the payment being made for each impression sold on its exchange. Whether we assume a second price auction or a first price auction, the expected payment is increasing in the number of bidders, and so an exchange maximises profits by maximising the number of bidders in each auction. For an exchange i, an increase in p_{dit} for any DSP d lowers the expected number of bidders in each auction, and therefore lowers i's profits. This means we can model exchange i's profits π_{it} as being a monotonically increasing function of p_{dit} for all d

maximises

$$\begin{aligned}\Pi_t &= \pi_{it} + \pi_{dt} \\ &= \pi_{it}(\boldsymbol{p_{it}}) + f_{dt} P_{dt}(m_{dt}, x_{dt}) M.\end{aligned}$$

Differentiating with respect to p_{kit} for some non-integrated DSP k yields

$$\begin{aligned}\frac{\partial \Pi_t}{\partial p_{kit}} &= \frac{\partial \pi_{it}}{\partial p_{kit}} + f_{dt} M \frac{\partial P_{dt}}{\partial p_{kit}} \\ &= \frac{\partial \pi_{it}}{\partial p_{kit}} + \gamma f_{dt} M q_{it} \int P_{adt} P_{akt} \, dG(z_t).\end{aligned} \tag{7}$$

The first term of this derivative is positive, but, provided that γ is negative, the second term of the derivative is negative. This means it is now possible that the profit maximising value of p_{kit} for the merged entity is lower than 1. That is, information frictions (decreasing p_{kit} for non integrated exchanges k) may increase the profits of a merged entity, whereas they decrease the profits of any non vertically integrated exchange. Further, the second (negative) term is larger for higher q_{it}, meaning that information frictions are of greater benefit to merged entities where the integrated exchange already has a high market share[18].

If we were to extend the model to allow exchanges to choose the level of p_{dit}, what we have shown here is that large vertically integrated exchanges plausibly have an incentive to introduce information frictions themselves.

We can consider this mechanism a variation of the notion of raising rivals' costs, common in the literature on vertical integration. In Salop and Scheffman (1983), the authors describe how, "if the upstream merger partner has some market power, input price increases to downstream rivals (perhaps to a level above the monopoly price) will raise their costs, allowing the dominant firm to increase price or output" in such a way that "upstream profits are sacrificed but downstream profits rise disproportionately".

[18]Note the assumption here that $\frac{\partial \pi_{it}}{\partial p_{kit} \partial q_{it}} = 0$

Equation 7 illustrates a similar effect, but where the relevant instrument is the information friction p_{kit} that exists between the upstream firm i and downstream rivals k, rather than the price charged for intermediate goods and services. We thus observe here a novel form of this antitrust concern that should be considered by regulators when assessing vertical integration and privacy regulation which restricts information flows in this market.

5 Estimation Approach

In this section I discuss identification of the parameters of my model.

For my demand estimation, I consider only the market for single-homing on a DSP. That is, in each market I consider only advertisers who use only one DSP in that market. When an advertiser multi-homes, it can direct a particular DSP to only buy impressions from exchanges with which it has low information frictions [19], and have another DSP buy impressions from the remaining exchanges with which the first DSP had relatively higher information frictions. Because of this, the proportions in which a DSP uses exchanges when it is purchasing impressions on behalf of a multi-homing advertiser do not in fact reflect the information frictions that DSP experiences with each exchange. For this reason, I restrict my model to only include the discrete choice of single-homing advertisers.

Given the assumptions about the process through which DSPs buy impressions in section 4, the fraction of DSP d's impressions that are bought from exchange i in market t can be expressed using Bayes' rule as:

$$s_{dit} = \frac{q_{it} p_{dit} r_{dit}}{\sum_{i=1}^{E} q_{it} p_{dit} r_{dit}}. \tag{8}$$

The fraction of DSP's d's impressions that are bought from exchange i in market t is estimated as:

[19]DSP's allow advertisers to select which exchanges the advertiser wants that DSP to connect to - any exchanges which are "shut off" will be sampled with zero probability by that DSP in step 1 of the buy process described above.

$$\hat{s}_{dit} = \frac{(single - homing\ impressions)_{dit}}{\sum_{j=1}^{E}(single - homing\ impressions)_{djt}}. \tag{9}$$

Due to the assumption of a common valuation distribution, all DSPs have the same ex-ante probability of winning the auction conditional on a good information draw, and so the r_{dit} term drops out of equation 8:

$$\begin{aligned} s_{dit} &= \frac{q_{it}p_{dit}\hat{r}}{\sum_{i=1}^{E} q_{it}p_{dit}\hat{r}} \\ &= \frac{q_{it}p_{dit}}{\sum_{i=1}^{E} q_{it}p_{dit}}. \end{aligned}$$

To estimate the exchange sampling probabilities q_{it}, I use the market shares according to multi-homing impressions[20]:

$$\hat{q}_{dit} = \frac{(multi - homing\ impressions)_{it}}{\sum_{i=1}^{E}(multi - homing\ impressions)_{it}}.$$

With these estimates for q_{dit}, r_{dit} and s_{dit}, and by rearranging equation 8, we can derive the following linear system of $E-1$ equations for each DSP d in each market t[21]:

[20] q_{dit} cannot be estimated as the aggregate market shares of the exchanges using single-homing impressions, as these market shares are the endogenous outcome of the impression buying process outlined in my model, and hence are themselves determined by p_{dit}. In order to identify q_{dit} separately from p_{dit}, multi-homing impressions are used. As mentioned above, when purchased on behalf of multi-homing advertisers, the proportions in which DSPs use different exchanges do not reflect information frictions, and so only reflect sampling probabilities.

[21] For each DSP d in each market t, there are E unknowns (p_{d1t}, p_{d2t}, ..., p_{dEt}) but only $E-1$ distinct equations. This is because the DSP-specific market share for the last exchange (exchange E) is determined once we know the DSP-specific market share for all the other exchanges, i.e.:

$$\hat{s}_{dEt} = 1 - \sum_{j=1}^{E-1} \hat{s}_{djt}$$

$$p_{dit} = \frac{\hat{s}_{dit}}{\hat{q}_{it}} \sum_{j=1}^{E} \hat{q}_{jt} p_{djt} \qquad \forall i. \tag{10}$$

This system identifies the E-1 ratios (which I label ρ):

$$\rho_{dit} \equiv \frac{p_{dit}}{p_{dEt}}.$$

A normalisation is required to recover the values p_{dit}, reflecting the fact that my model only identifies information frictions relative to a "baseline" exchange, which I define as an "other" category into which I place all exchanges with which there is no a priori reason to expect heterogeneity in information frictions across DSPs.

Dividing equation 10 through by p_{dEt}, the linear system becomes:

$$\rho_{dit} = \frac{\hat{s}_{dit}}{\hat{q}_{it}} \sum_{j=1}^{E} \hat{q}_{jt} \rho_{djt} \qquad \forall i \neq E. \tag{11}$$

Rearranging yields

$$\rho_{dit} = \frac{\hat{s}_{dit}}{\hat{q}_{it}} \sum_{j=1}^{E-1} \hat{q}_{jt} \rho_{djt} + \frac{\hat{s}_{dit}}{\hat{q}_{it}} \hat{q}_{Et} \qquad \forall i \neq E, \tag{12}$$

because $\rho_{dEt} = 1$. Defining $(E-1)$x$(E-1)$ matrix $\hat{\boldsymbol{A}}_{\boldsymbol{dt}}$ and $(E-1)$x1 vectors $\hat{\boldsymbol{S}}_{\boldsymbol{dt}}$ and $\boldsymbol{\rho_{dt}}$ appropriately, we can write out the linear system for each DSP d in each market t as

$$\begin{aligned} \boldsymbol{\rho_{dt}} &= \hat{\boldsymbol{A}}_{\boldsymbol{dt}} \boldsymbol{\rho_{dt}} + \hat{\boldsymbol{S}}_{\boldsymbol{dt}} \hat{q}_{Et} \\ &= (I - \hat{\boldsymbol{A}}_{\boldsymbol{dt}})^{-1} \hat{\boldsymbol{S}}_{\boldsymbol{dt}} \hat{q}_{Et}. \end{aligned} \tag{13}$$

Equation 13 allows estimation of the full vector $\boldsymbol{\rho_{dt}}$ for each DSP d in each market t by using a matrix inversion to solve the linear system[22].

[22] This can be done iteratively over all DSPs and time markets to yield ρ_{dit} values for every exchange i, for every DSP d in every market t. Note that a condition for identification is that the matrix $(I - \hat{\boldsymbol{A}}_{\boldsymbol{dt}})$ is non-singular. This will not be the case if $s_{dEt} = 0$ for a DSP d in some market t (i.e. if a DSP does not use

Using equation 4, along with the estimates for p_{dit} and q_{it}, I recover estimates for the level of search frictions at each DSP in each market, m_{dt}. Using these m_{dt} values, along with other characteristics x_{dt} of each DSP in each market, I carry out the demand estimation of the parameters of equation 1.

As I have data on individual advertiser choices of DSP, I am able to estimate the parameters of the utility function using a micro-BLP approach (Berry et al. (2004)). The parameters of the model can be estimated by maximising the likelihood function

$$l = \sum_t \sum_a \sum_d \mathbb{1}_{adt} ln(P_{adt}). \tag{14}$$

Where $\mathbb{1}_{adt}$ is an indicator equal to 1 if advertiser a chose DSP d in market t, and 0 otherwise. P_{adt} is the probability advertiser a selects DSP d in market t:

$$P_{adt} = \frac{e^{\delta_{dt}+\lambda_{adt}}}{\sum_j e^{\delta_{jt}+\lambda_{ajt}}}. \tag{15}$$

The main aim of the paper is to estimate γ. The identification issues that arise for the estimation procedure can be elucidated by examining equation 2.

I normalise utility by setting the utility from one of the DSPs to zero (the DSP which is most frequently selected amongst the choice set). In doing so, I redefine the choice situation as being the choice of DSP conditional on selecting a (single) DSP (see Train (2009)).

There may be a spurious correlation between the market share of an exchange and the market share of its vertically integrated DSP which is the result of factors *other* than the information frictions which are the focus of my model. For example, an entity-wide improvement in productivity or brand image (unobservable characteristics) could raise the service level of both a DSP and its vertically integrated exchange. As m_{dt} is by definition correlated with exchange market shares q_{it}, this could cause m_{dt} to be correlated with the unobservable ξ_{dt}, biasing the estimate for γ.

the "other exchange" option in a particular market). Hence, a condition for identification is that each DSP purchases at least one impression from the "other exchange" option.

I account for this with a control function approach. The argument above suggests there is some set of variables y which should be included on the right hand side of equation 2 (brand image, customer service, etc.) which are correlated with the market share of any exchange i with which DSP d is vertically integrated. The effect of this set of variables on the value of δ_{dt} for DSP d can be written as $g(q_{it})$, which I write as a function of q_{it} to reflect the correlation with the market share of exchange i. I control for this by including the market share of any exchange i with which d is vertically integrated as a variable on the right hand side, effectively assuming that $g(q_{it}) = \alpha_i q_{it}$. Note that α_i does not have a causal interpretation, but simply reflects the (what I have assumed to be linear) relationship that there is between the δ_{dt} of DSP d and the market share of its vertically integrated exchange i as a result of the type of factors described above (correlated shifts in brand image, productivity, etc.). With this addition, equation 2 becomes

$$\begin{aligned}\delta_{dt} &= \beta' x_{dt} + \gamma m_{dt} + \alpha_1 \mathbb{1}_{Vd}(1) q_{1t} + ... + \alpha_E \mathbb{1}_{Vd}(E) q_{Et} + \xi_{dt} \\ &= \beta' x_{dt} + \gamma m_{dt} + \left(\sum_{i=1}^{E} \alpha_i \mathbb{1}_{Vd}(i) q_{it} \right) + \xi_{dt},\end{aligned} \tag{16}$$

where V_d is the set of all exchanges with which DSP d is vertically integrated, and therefore $\mathbb{1}_{Vd}(i) = 1$ if exchange i is vertically integrated with DSP d, and 0 otherwise. I name the variable $\mathbb{1}_{Vd}(i) q_{it}$ the external factor correlation between exchange i and DSP d in market t.

The additional DSP characteristics I observe are the CPM (a measure of the media cost per impression advertisers are having to pay for the impressions that are being bought by the DSP on their behalf), CTR (the number of clicks on all ad impressions divided by the total number of impressions), as well as a proxy for the fee charged by the DSP to advertisers - note this is not included in the calculation of CPM and is measured as a percentage of the media cost.

As is standard in the literature on demand estimation (Train (2009)), the question arises over whether the price variable is endogenous, in which case an instrumental variable approach would be warranted. In the absence of an appropriate instrument, I assume that the search friction variable m_{dt} is uncorrelated with the fee variable. As a robustness check, I estimate my model both including and excluding the fee variable.

I estimate my model on data on the US from July 2017 to July 2020 and on the UK from January 2018 to July 2020, with market t defined as a country-month. In total there are 78 markets and 3 DSPs which have at least one advertiser single homing on them in at least one market. There are two DSPs which are vertically integrated with exchanges: Google and AppNexus. I give the exchanges with which they are vertically integrated the same name as the corresponding DSP.

6 Results

Figure 5 demonstrates the difference in information frictions between integrated and non integrated DSPs. This figure compares the Google exchange information draw (p_{dit}) of Google's DSP to Appnexus' DSP in the UK from July 2018 to July 2020. Note the different scales on the axes. It is clear that Appnexus (the second largest DSP in the data set, after Google's DSP) faces substantial information frictions while using Google's exchange, and that Google's DSP therefore has a significant advantage over Appnexus in gaining impressions from Google's exchange.

The results of my model estimation are displayed in table 1. The table contains the estimates for the parameters in δ_{dt} - these are γ and the parameters of β, shown in equation 2. The estimates for the parameters in μ_d for each DSP d are omitted for ease of exposition.

The most important thing to note is the negative and statistically significant coefficient on m_{dt} - i.e. that $\gamma < 0$. I discussed the implication of this finding in section 4. This indicates that search frictions do in fact have a negative impact on the utility gained from a particular DSP, and thus that the presence of information frictions between exchanges and non vertically integrated DSPs can work to steer advertiser demand away from those DSPs. Apart from its sign, the value of the coefficient does not allow for easy interpretation as it

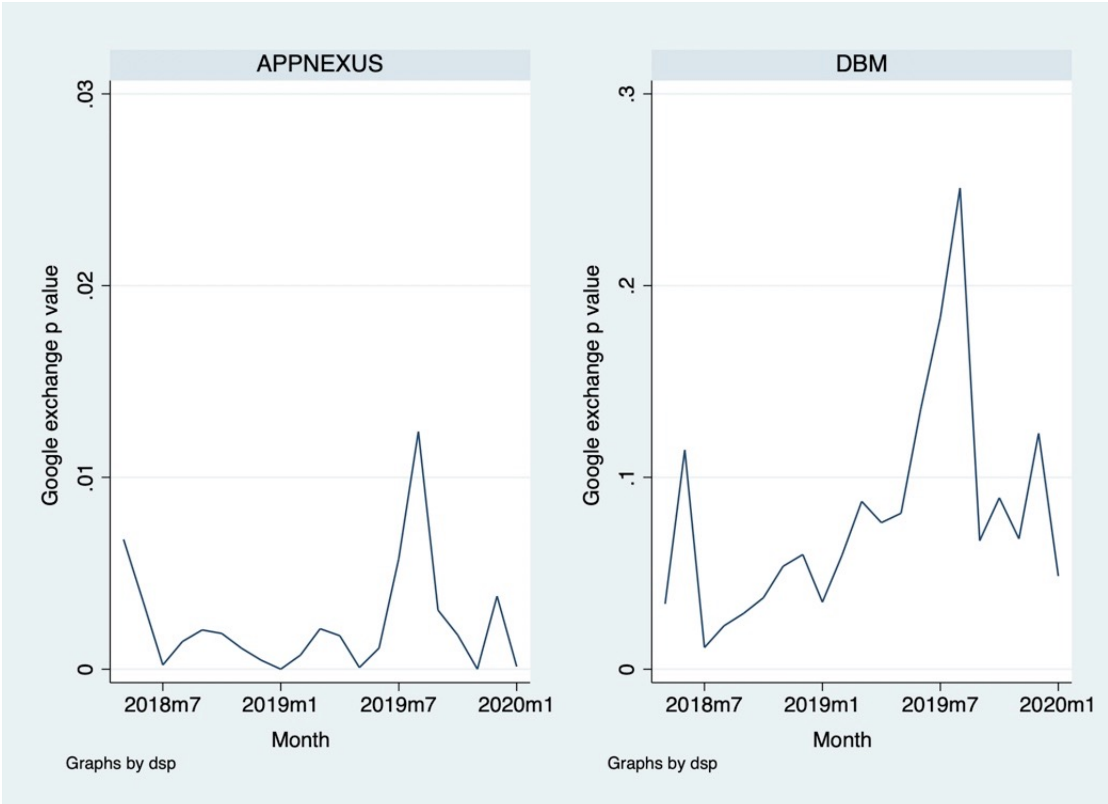

Figure 5: Comparison of the Google exchange information draw (p_{dit}) of Google's DSP to Appnexus' DSP in the UK. Note the different scales on the axes.

is. Below, I simulate a counterfactual which gives a clearer quantitative interpretation of my result.

As I discussed in section 4, the fact that the search friction measure does act as a demand driver means that information frictions between exchanges and non-integrated DSPs can act to the benefit of merged entities by steering demand toward integrated DSPs and boosting their market share.

It is also interesting to note that the coefficients on both of the external factor correlation variables (for Appnexus and Google, the two entities which own both a DSP and an exchange) are positive and significant. This indicates that there is some positive relationship between the market share of a DSP and of its vertically integrated exchange which is not attributable to the search friction mechanism, and that my model has successfully identified it separately.

All the results appear to be robust to the removal of fee from the estimation.

Table 1: The mixed logit estimates of the parameters in δ_{dt}

Variable	(1) Fee included	(2) Fee not included
Fee	-2.656	
	(1.695)	
CPM	71.09**	94.02***
	(30.65)	(27.16)
CTR	0.0485	-0.00702
	(0.486)	(0.486)
m	-6.427***	-5.605***
	(1.850)	(1.776)
Appnexus EFC	3.638***	3.550***
	(1.164)	(1.160)
Google EFC	2.123***	1.914***
	(0.423)	(0.401)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Robustness Check

My estimation is of a discrete choice model where each advertiser's choice of DSP is weighted equally. There may be a concern that this is not the most relevant means of defining market share, as different advertisers may in fact provide different amounts of business by requesting different numbers of impressions. That is, I have implicitly defined market share of a DSP as number of single homing advertisers it serves, where as one might instead define market share as share of impressions purchased on behalf of single homing advertisers.

To address this concern, as a robustness check I also estimate a simple logit model, where market share is defined as the share of impressions purchased on behalf of single homing advertisers. Using the familiar inversion of the market share equations (Berry (1994)), for this model, δ_{dt} is equal to $log(P_{dt}) - log(P_{Ot})$ (where P_{dt} is the market share of DSP d in

Table 2: Simple Logit Robustness Checks

Variable	(1) Fee included	(2) Fee not included
Fee	8.55	
	(7.23)	
CPM	−0.21	−0.29
	(0.62)	(0.58)
CTR	14.32^{***}	15.29^{***}
	(1.60)	(0.63)
m	-11.68^{***}	-14.15^{***}
	(1.99)	(4.27)
Google EFC	1.49	2.17
	(1.42)	(2.13)

Standard errors in parentheses

$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$

market t, and O denotes the outside option DSP) [23]. This means the estimation becomes an OLS estimation of the following linear model, which I estimate with DSP fixed effects:

$$\begin{aligned}\log(P_{dt}) - \log(P_{Ot}) &= \beta' x_{dt} + \gamma m_{dt} + \alpha_1 \mathbb{1}_{Vd}(1) q_{1t} + ... + \alpha_E \mathbb{1}_{Vd}(E) q_{Et} + \xi_{dt} \\ &= \beta' x_{dt} + \gamma m_{dt} + \left(\sum_{i=1}^{E} \alpha_i \mathbb{1}_{Vd}(i) q_{it} \right) + \xi_{dt}.\end{aligned} \tag{17}$$

The results of the robustness check is shown in table 2. We can see that the results are similar, with a significantly negative γ in both specifications of the model. One point of interest with these estimates is the significantly negative coefficient on CTR. This can be explained by the fact we are defining market share as share of impressions here. CTR (click through rate) gives the number of clicks per thousand impressions, and as such will be higher for "efficient" marketing campaigns which use fewer impressions but still generate the same number of clicks.

[23] As Appnexus is the DSP defined as the outside option, the external factor correlation for Appnexus is no longer a variable in this model.

Counterfactual simulations

Using the estimated model, I now simulate a counterfactual scenario where all values of m_{dt} are set to 0, thereby eliminating the heterogeneity in search frictions across DSPs. This has policy relevance as we can interpret this restriction as being a simulation of what might happen if an authority imposed regulation eliminating the existence of any information frictions between exchanges and DSPs. My estimated model allows me to investigate the effect of such a hypothetical policy on the market share of DSPs during the time period of my data.

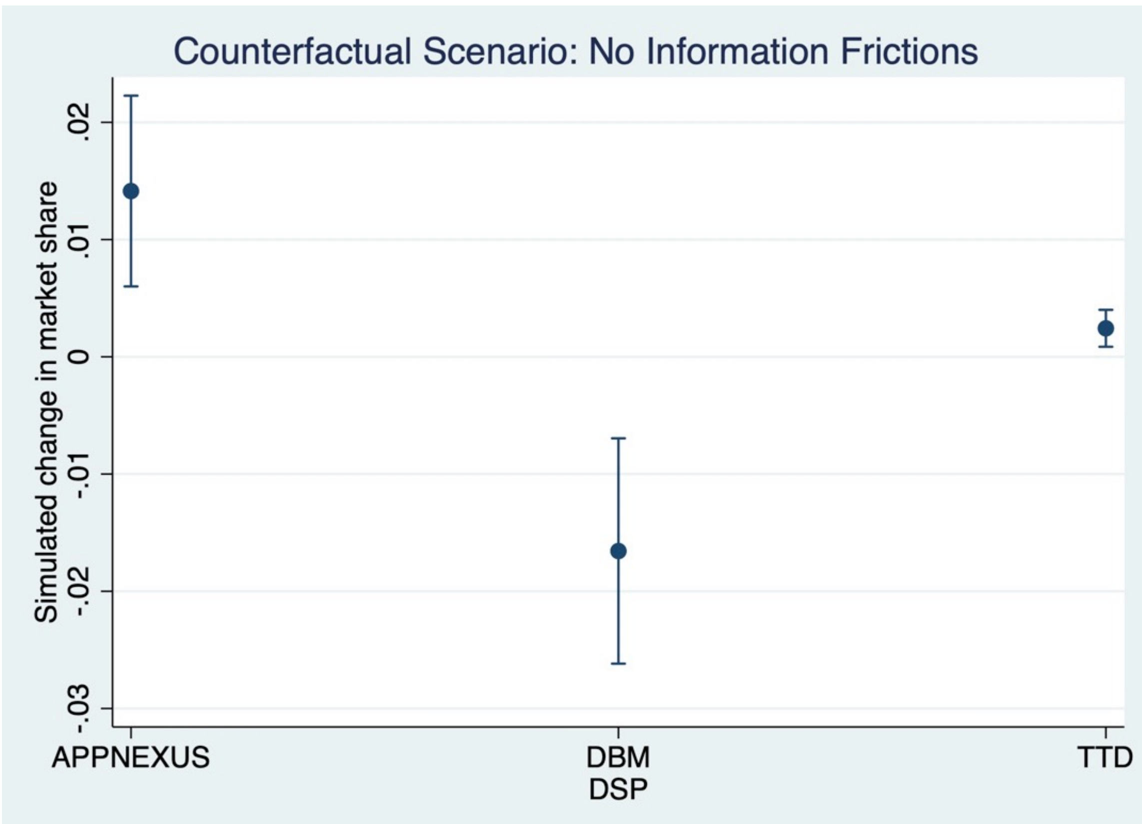

Figure 6: The effect of eliminating all information frictions on single-homing market share for each DSP

Figure 6 shows the results from the simulation using the discrete choice model estimated by maximum likelihood. The number displayed for each DSP is the change in that DSP's market share that would have occurred in my sample if all information frictions were eliminated. The number is averaged over all markets (country-months) for each DSP. The error bars are 95% confidence intervals for these estimates. Eliminating all information frictions only has a negative impact on the market share of DBM, showing clearly that it is this DSP which benefits from the demand steering effect of information frictions. The negative effect on DBM's market share is on average around 1.66 percentage points.

Google benefits the most from these information frictions because its vertically integrated exchange has a strong market position, meaning its downstream DSP gains a large advantage when information frictions steer demand from exchanges to their integrated DSPs. In this way, market dominance is proliferated from upstream firms to downstream integrated partners.

This counterfactual simulation demonstrates how vertical integration, coupled with the existence of information frictions, has had anti-competitive effects in the downstream market.

7 Conclusion

I propose and estimate a structural model of the display advertising intermediary industry that highlights and quantifies the demand steering effect that information frictions between exchanges and DSPs can have on the market for DSPs when some firms are vertically integrated. I find that these information frictions cause search frictions which have a significant negative effect on demand for a DSP, in such a way that advertiser demand is steered to vertically integrated DSPs which have an information advantage with their partnered exchange. My estimated model allows me to carry out a counterfactual analysis which simulates the effect of eliminating information frictions, showing that such a policy would reduce the market share of the DSP in my data which has most benefited from these frictions (Google's DBM) by around 1.66%.

This paper builds on previous work studying the potential for demand steering that can result from vertical integration, looking at a novel context where an auction mechanism governing firm to firm interactions means that information provision plays a key role in this demand steering effect. As with other empirical work in this literature, my analysis supports the view that vertical integration can have anti-competitive effects.

My analysis carries policy implications. I have shown empirically that information frictions between upstream and downstream entities can have anti-competitive effects which allow a dominant upstream firm to proliferate its market power to reduce competition in the downstream market. Consumer privacy protection policies like GDPR can exacerbate these information frictions by making it more difficult for firms to transfer information on web users to other non-integrated firms. While such protection policies have clear benefits for consumer privacy, my analysis demonstrates that regulators must be careful that they understand the potential anti-competitive effects that such policies can have, extending the power of large "walled garden" data hoarders like Google.

A potential avenue for future research should be to more directly quantify the welfare effects of reduced competition amongst DSPs, and of a reduction in the allocative efficiency of online display ads, such that this can also be balanced against the positive effects of privacy protection regulation. Another direction for future research is to study the role that dynamics may have in this market, as it is possible that the market's two-sided nature could result in dynamic indirect network effects.

8 Appendix

8.1 Proof of Proposition 1

Proof.

The own-search-friction derivative of market share is

$$
\begin{aligned}
\frac{\partial P_{dt}}{\partial m_{dt}} &= \int \frac{e^{\delta_{dt}+\lambda_{adt}}}{\sum_j e^{\delta_{jt}+\lambda_{ajt}}} \frac{\partial(\delta_{dt}+\lambda_{adt})}{\partial m_{dt}} - \frac{e^{\delta_{dt}+\lambda_{adt}}}{(\sum_j e^{\delta_{jt}+\lambda_{ajt}})^2} e^{\delta_{dt}+\lambda_{adt}} \frac{\partial(\delta_{dt}+\lambda_{adt})}{\partial m_{dt}} \, dG(z_t). \\
&= \int \gamma P_{adt}(1-P_{adt}) \, dG(z_t).
\end{aligned}
$$

The cross-search-friction derivative of market share is

$$
\begin{aligned}
\frac{\partial P_{dt}}{\partial m_{nt}} &= -\frac{e^{\delta_{dt}+\lambda_{adt}}}{(\sum_j e^{\delta_{jt}+\lambda_{ajt}})^2} e^{\delta_{nt}+\lambda_{ant}} \frac{\partial(\delta_{nt}+\lambda_{ant})}{\partial m_{nt}} \\
&= -\int \gamma P_{adt} P_{ant} \, dG(z_t).
\end{aligned}
$$

Given that $\frac{\partial m_{dt}}{\partial q_{it}} = (1-p_{dit})$, applying the chain rule yields

$$
\begin{aligned}
\frac{\partial P_{dt}}{\partial q_{it}} &= \sum_k^N \frac{\partial P_{dt}}{\partial m_{kt}} \frac{\partial m_{kt}}{\partial q_{it}} \\
&= \gamma \left(\int P_{adt}(1-P_{adt})(1-p_{dit}) \, dG(z_t) - \sum_{k \neq d}^N \int P_{adt} P_{akt}(1-p_{kit}) \, dG(z_t) \right).
\end{aligned}
\tag{18}
$$

Provided γ is negative, the above implies that $\frac{\partial P_{dt}}{\partial q_{it}} > 0$ provided that

$$
\int P_{adt}(1-P_{adt})(1-p_{dit}) \, dG(z_t) < \sum_{k \neq d}^N \int P_{adt} P_{akt}(1-p_{kit}) \, dG(z_t).
$$

■

8.2 Proof of Proposition 2

Proof.

Follows by inspection of equation 18.

■

References

8marketcap, 2022, 'Largest companies by market cap,' .

Ahuja, S., 2020, 'Marketing budget allocation 2021 – breakdown, digital spend, and template,' .

Alcobendas, M., Kobayashi, S. and Shum, M., 2021, 'The impact of privacy measures on online advertising markets,' *Available at SSRN: https://ssrn.com/abstract=3782889 or http://dx.doi.org/10.2139/ssrn.3782889*.

Berry, S., Levinsohn, J. and Pakes, A., 2004, 'Differentiated products demand systems from a combination of micro and macro data: The new car market,' *Journal of political Economy*, 112(1), pp. 68–105.

Berry, S. T., 1994, 'Estimating discrete-choice models of product differentiation,' *The RAND Journal of Economics*, pp. 242–262.

Bork, R., 1978, *The Antitrust Paradox: A Policy at War with Itself* (Free Press).

CMA, 2019, 'Online platforms and digital advertising market study,' *https://www.gov.uk/cma-cases/online-platforms-and-digital-advertising-market-study*.

Crawford, G. S., Lee, R. S., Whinston, M. D. and Yurukoglu, A., 2018, 'The welfare effects of vertical integration in multichannel television markets,' *Econometrica*, 86(3), pp. 891–954, doi:https://doi.org/10.3982/ECTA14031.

Cuesta, J. I., Noton, C. and Vatter, B., 2019, 'Vertical integration between hospitals and insurers,' *Available at SSRN 3309218*.

Hart, O., Tirole, J., Carlton, D. W. and Williamson, O. E., 1990, 'Vertical integration and market foreclosure,' *Brookings papers on economic activity. Microeconomics*, 1990, pp. 205–286.

Jeon, D.-S., 2021, 'Market power and transparency in open display advertising – a case study,' *European Commission*.

LaFrance, A., 2017, 'The first-ever banner ad on the web,' *The Atlantic*.

Lomas, N., 2021, 'Eu is now investigating google's adtech over antitrust concerns,' *TechCrunch*.

Lomas, N., 2022, 'Uk finally opens antitrust probe of google's role in the adtech stack,' *TechCrunch*.

McFadden, D., 1974, 'Conditional logit analysis of qualitative choice behavior,' *in P. Zarembka, ed., Frontiers in Econometrics*.

Milgrom, P. R. and Weber, R. J., 1982, 'A theory of auctions and competitive bidding,' *Econometrica*, 50(5), pp. 1089–1122.

O'Toole, D., 2020, 'The state of texas, et al. v. google llc, civil no. 4:20-cv-957-sdj (e.d. tex. may. 20, 2021),' *Filed civil complaint*.

Salop, S. C. and Scheffman, D. T., 1983, 'Raising rivals' costs,' *The American Economic Review*, 73(2), pp. 267–271.

Scott Morton, F. M. and Athey, S., 2021, 'Platform annexation,' *Available at SSRN: https://ssrn.com/abstract=3786434*.

Spengler, J. J., 1950, 'Vertical integration and antitrust policy,' *Journal of Political Economy*, 58(4), pp. 347–352.

Srinivasan, D., 2020, 'Why google dominates advertising markets,' *Available at SSRN: https://ssrn.com/abstract=3500919*.

Statista, 2021, ‘Digital advertising in the uk 2021 - statistics and facts,’ .
Train, K. E., 2009, *Discrete choice methods with simulation* (Cambridge university press).