

“From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising” by Francesco Decarolis and Gabriele Rovigatti

Summary by Gina Markov

This paper studies the effect of intermediary concentration on search engine revenue in digital advertising. Digital advertising is mostly divided between display advertising (ads on a website like The New York Times) and sponsored search. This study focuses on the latter. In sponsored search advertising, advertisers bid on one of a limited number of “slots” available on a search engine’s result page for a given search term, or keyword. For example, if a user googles ‘best blender,’ the first few results may be ads for Vitamix blenders and KitchenAid blenders. Oftentimes, advertisers do not bid on slots themselves. Rather, they pay marketing agencies (MAs) to run their marketing campaigns and bid on their behalf. Furthermore, many MAs are owned, at least in part, by large agency networks. The authors classify an MA as part of a network if the network holds operating control (at least 51%) of the MA. Agency networks are similar to brokers in that they facilitate transactions between advertisers and publishers, both of which may be members of the network; the networks manage publishers’ unsold inventory as well as advertisers’ marketing campaigns (providing information like target audience and budget), matching an advertiser’s demand with a publisher’s supply. However, in the context of this paper – sponsored search – the network is only acting on behalf of the advertiser, since the relevant publisher is Google.


There are only around 7 networks in the industry. The combined market share of the four largest networks – WPP, Omnicom, Publicis Groupe, and Interpublic Group of Companies – reached 74% in 2017. Due to these networks’ consolidated market power and prominence, they comprise a large buyer for a search engine. On the supply side, Google dominates the sponsored search market, earning 75-80% of total search advertising revenue in the US between 2016 and 2018; Amazon and smaller players in sponsored search have also emerged.

This paper addresses the idea of “countervailing power.” Are advertising intermediaries like networks able to counterbalance the dominance of online ad suppliers – namely Google? This study quantifies the effectiveness of increasing buyer concentration in offsetting the high bargaining power of dominant platforms. Where previous works emphasize the positive roles of intermediaries in improving use of information, effectively managing advertisers’ budgets, and fostering healthy competition in ad auctions, this empirical study expands on the perils of coordinated bidding for Google’s sponsored search revenue.

The authors gather data from Redbooks, a comprehensive database about marketing agencies covering the years 2014 to 2017. There is information about a few thousand MAs, the seven networks to which most of them belong, and around 6,000 advertisers. The study combines this data in a novel way with data on Google’s sponsored search auctions from the software and marketing analytics company SEMrush. For each keyword and year, the authors know the search volume and the average price advertisers paid for a user’s click on the ad, known as the Cost-Per-Click (CPC). The dataset also highlights the hierarchical structure of advertising intermediation: which advertisers belong to which

MAs, and which MAs belong to which networks. This structure allows the authors to extract information about coalitions – networks that participate in auctions on behalf of two or more different competing advertisers. In the sample, some advertisers belong to multiple MAs, but these MAs almost always belong to the same network. The authors make several important observations from the dataset based on keyword statistics. They

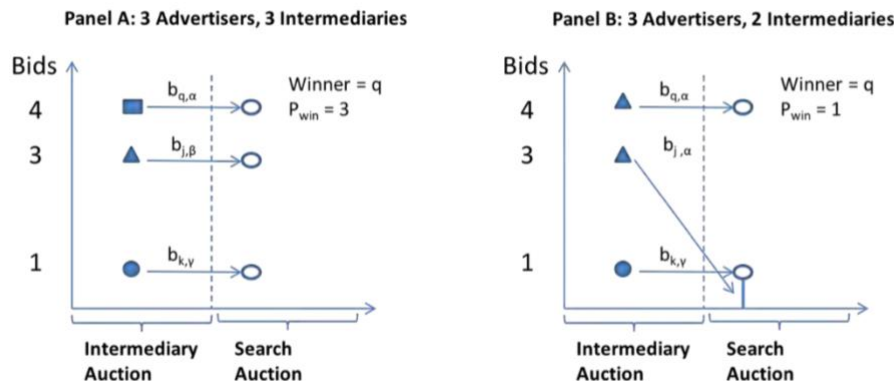
Figure 2: Redbooks-SEMrush Data Structure



compare statistics for keywords with at least one network advertiser to those with at least one independent advertiser (an advertiser who bids on its own without an MA or with an MA not affiliated to a network). These categories contain significant overlap. Since keywords with at least one network advertiser receive less traffic (traffic is

defined as the share of the advertiser’s traffic associated with that keyword, aggregated across all relevant advertisers) on average than independent advertisers, the authors infer that network advertisers tend to place ads over more keywords than independent advertisers. The authors also find that for each keyword, there is almost always exactly one coalition and it has an average size of 2.38 (the vast majority of coalitions have size 2). This suggests that different networks specialize in different segments of the keyword market.

The authors identify a pattern called bid coordination. This is when networks split the keyword market and selectively allocate slots that multiple advertisers want. Bid coordination, the authors explain, is the one of the mechanisms by which intermediaries can lower ad auction revenue. A network analyzes the large amount of information it has regarding how much its advertisers can bid and their ad quality. With this knowledge, the network determines how much to bid on behalf of its advertisers at a given auction. Figure 1 provides an example of bidding through intermediaries. In this second price auction, there are three advertisers (q, j, k) bidding, and the shape of the advertiser represents its intermediary. There is no bid distortion in Panel A. In Panel B, the intermediary representing q and j depresses j’s bid from 3 to 1. While advertiser q wins in both panels, it pays much less in Panel B because of the intermediary’s bid distortion – the cost drops from 3 to 1.



An alternative coordination strategy involves market splitting. That is, the network can assess which keywords might provide its advertisers with the most revenue by matching the keyword traffic with the ad's target audience and quality. In this way, the network selectively chooses which ads to bid with for a certain keyword. For example, an intermediary may have its advertisers bid only on their own branded keywords. If Ford and Toyota belonged to the same intermediary, the intermediary can ensure that only Ford bids on the keyword 'Ford truck,' and only Toyota bids on 'Toyota truck.' Keyword targeting can also occur on geographical, demographic, device, and temporal levels. An intermediary may also rotate bids so that its advertisers never meet in the same auction for a certain person in a certain time and area. This behavior reduces competition among the network's clients, while still exposing the advertisers to the same audience. The keyword market can be segmented in nearly an infinite number of ways, and it is difficult to detect this kind of coordination. Furthermore, it is worth noting that advertisers' explicit coordination to not bid on each other's brands is a type of behavior recently sanctioned by the Federal Trade Commission, but not if this coordination is implemented by their intermediary. However, such coordination may also raise efficiency in the sense that a consumer searching for Adidas sneakers is probably going to click on an Adidas ad at a higher rate than a Nike ad.

The authors first use keyword-level data to analyze if intermediaries' behavior is indicative of strategies to alter search engine revenue. They analyze two types of increases in intermediary concentration: individual advertisers joining MAs and networks expanding via MA acquisitions. For the first case, they find that, compared to when an advertiser bids on its own, MA bidding is associated with an increase in the number of keywords bid on and a decrease in the average CPC. Particularly, the number of long tail keywords bid on increases. Long tail keywords are longer, more specific keywords (at least four terms, like "charity donations furniture pickup") that see less competition (by being rare) and are more likely to be searched by users close to finalizing their purchase decision, leading to more clicks. While more clicks increase search engine revenue, they also likely lead to greater revenue for the company who is advertising. However, it is risky to analyze the impact of formerly-independent advertisers joining MAs on the CPC, because some advertisers might join an MA due to an inability to optimize their bidding strategies; this would mean that a lower CPC can be explained by excessively high bids before joining the MA, not bid coordination. Therefore, the authors also analyze network expansions via MA acquisitions (they assume the MAs have already optimized bidding from the individual bidder's perspective). In order to detect bid coordination and market splitting, they look at changes in the overlap of clients' keywords before and after the acquisition. The clients belong to either the network or the acquired MA. If the overlap declines, the network might be engaging in market splitting. If the overlap stays the same or grows, the network might affect search engine revenue within the auctions themselves. There is evidence in the data that both strategies are adopted, so the authors move to an empirical strategy with statistical regression.

In order to perform the empirical analysis, the authors offer a new way to define the keyword market. They use natural language processing techniques to create non-overlapping clusters of keywords to represent new individual markets. The clusters are created in two layers: first thematic, then competitive. Thematic clustering is based on semantic similarities between words (similarities related to meaning). For example, car

and automobile would belong to the same thematic cluster. This cluster reflects consumer perception, since it resembles how consumers learn about products on the internet. The second layer, competitive clustering, further partitions the thematic clusters in terms of the set of advertisers who bid for them. Thematic clusters are further grouped together if they are frequently bid on by the same advertisers. In this way, the competitive clusters reflect advertiser perception. Under this market definition, the average market has 37 keywords and 4 competing advertisers, with the number of competing advertisers for a single keyword being 1.62 on average.

The authors then measure network concentration at the cluster level, and assess the relationship between that concentration and Google's revenues. The empirical model assumes a linear relationship between the (log) of Google's revenues and demand concentration across each time period (the year between 2014 and 2017) and market (defined in the clusters described above). To measure the magnitude of this effect, the authors need to create proxies for Google's revenue and demand concentration. Google's revenue per market and time period is measured (loosely) by the product of search volume, CTR (cumulative click-through rate) and CPC. Demand concentration is defined in a more complicated manner. An intermediary's share is the sum of the clicks that the market keywords received in slots occupied by advertisers under that intermediary. Each intermediary's market share is then squared and summed to produce a concentration measure for each market and time. The authors find that around 40% of markets are highly concentrated.

The study highlights the imperfect nature of the proxy for demand concentration as a source of measurement error, and discusses the risk for omitted variable bias. Omitted variable bias might occur if a keyword becomes popular for exogenous reasons, like a trend that changes consumer taste, causing advertisers to hire the intermediary that specializes in that keyword market; this would then lead to an increase in intermediary concentration and in search engine revenue, although there is no causal relationship implied by this event. To eliminate bias, the authors control for differences across thematic clusters and time (using a fixed effects specification), as well as market-time varying observables, like the sudden change in the popularity of a keyword. They also use an instrumental variables (IV) strategy to refine the coefficient of the concentration variable. The IV approach is used in the context of mergers and acquisitions (M&As), which occur when intermediaries purchase MAs. In essence, the authors simulate the change in concentration of a certain market *induced by a merger*.

The authors find that M&As substantially increase the concentration of the markets affected by these mergers. They also find that greater network concentration in keyword markets leads to lower search engine revenues. The paper notes that the decrease in search engine revenue is caused primarily by a decrease in CPC. They do not find a decrease in output. As discussed previously, networks may be lowering CPC via bid coordination, relaxing competition among their subsidiary advertisers. Another way networks lower CPC is by bolstering competition among ad selling platforms. For example, they can split advertisers' marketing budgets across Google, Amazon, Bing, display advertisers, and social media advertising. Some intermediaries also bargain with Google due to their huge economic size. Google can create side deals with networks by

making small tweaks to its auction structure. For example, Google may bolster an ad's quality score, which reduces its CPC.

The authors also outline several ways that intermediary concentration can be beneficial to search engines. Intermediaries can foster entry to advertisers who would not otherwise enter certain auctions, increase the number of keywords on which advertisers bid, and enhance bidding speed. Nevertheless, they conclude that intermediary concentration transfers a significant portion of the value created by search advertising away from the platform and toward buyers. This is true in a static sense, not considering the many actions that search engines are taking to limit this phenomenon. Indeed, selling platforms are taking several steps to reduce their loss of market power. They are increasing auction reserve prices, reducing the number of ad slots offered, and promoting disintermediation services. The first two methods threaten market efficiency, however, because they likely hurt the "wrong" advertisers – those not sharing a common intermediary. Small independent advertisers bidding close to the reserve price may have to pay higher prices or are excluded from the auction altogether, undermining market efficiency.

The final question that the paper raises is the extent to which Google's reduced revenues can positively contribute to consumer welfare. Greater consumer welfare would suggest that bolstering buyer power may be a desirable policy to address supply-side monopoly prices and platform concentration. Decreasing a network-owned advertiser's costs may result in lower consumer prices. However, this can only occur if the network passes through the lower prices to advertisers. A network's revenue stream has two main sources: an initial buy-in cost for the advertiser to become a member of the network, and commissions on a per-ad basis. Lower ad prices would reduce commissions, passing the lower prices through to the advertiser. Greater competition among intermediaries could also reduce the initial buy-in price and commissions. Regarding intermediary competition, the authors claim that there is no conclusive evidence on the issue. Networks increasingly face competition not just from each other, but also independent agencies and consulting firms. Nevertheless, there are several hinderances to healthy competition among advertising intermediaries: the difficulties in measuring returns to advertising, lack of transparency about budget spending between intermediaries and advertisers (causing advertisers to fail to optimize their bidding campaigns), the lock in of advertisers with an intermediary due to bid coordination, common ownership of intermediaries causing collusion, and network coordination with Google. The largest, most economically powerful networks may engage in mutually-beneficial murky business practices with Google, like Google's "*cash rebate program*." Under this program, networks receive cash back for buying more ads on Google's search engine, which may not be in the best interests of the network's members. Such a program may help Google establish its long-run market dominance. However, it is unclear if, or how much of, the rebate money is actually passed down to the network's clients. There can also be harms to consumer welfare due to increasing buyer power. These include reduced incentives for suppliers like Google to innovate, as well as worse, fewer options for advertisers' intermediary needs – causing higher advertising costs and harming consumers.