

The longest second: Header bidding and the material politics of online advertising

Donald MacKenzie, Koray Caliskan & Charlotte Rommerskirchen

To cite this article: Donald MacKenzie, Koray Caliskan & Charlotte Rommerskirchen (2023) The longest second: Header bidding and the material politics of online advertising, *Economy and Society*, 52:3, 554-578, DOI: [10.1080/03085147.2023.2238463](https://doi.org/10.1080/03085147.2023.2238463)

To link to this article: <https://doi.org/10.1080/03085147.2023.2238463>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 23 Aug 2023.



Submit your article to this journal [↗](#)



Article views: 428



View related articles [↗](#)



View Crossmark data [↗](#)



The longest second: Header bidding and the material politics of online advertising

Donald MacKenzie , Koray Caliskan  and Charlotte Rommerskirchen 

Abstract

A user's online action is often followed, around a second later, by ads being shown to her/him. Much happens in that second, including near-instantaneous auctions (sometimes coordinated by the user's own phone or other device) in which algorithms bid to show particular advertisers' ads. Contributing to the burgeoning social-science literature on online advertising, we examine contending material forms these auctions take in 'open display' advertising. We trace the emergence of Google's centralized auctions, and how they have been challenged by decentralized 'header bidding'. We argue: first, that ad platforms should be seen as 'stack economization' processes, which layer different forms of economization in

Donald MacKenzie (corresponding author), School of Social and Political Science, University of Edinburgh, 15a George Square, Edinburgh, EH8 9LD, United Kingdom. E-mail: d.mackenzie@ed.ac.uk

Koray Caliskan, Parsons School of Design, The New School, 66 5th Ave, New York, NY, 10011, United States. E-mail: caliskak@newschool.edu

Charlotte Rommerskirchen, School of Social and Political Science, University of Edinburgh, 15a George Square, Edinburgh, EH8 9LD, United Kingdom. E-mail: charlotte.rommerskirchen@ed.ac.uk

Copyright © 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.



complex ways; second, that those processes are sometimes fiercely contested, and can be the site of intricate – and currently changing – material politics.

Keywords: online advertising; economic platforms; stack economization; material politics; header bidding; auctions.

Revenue from digital advertising – totalling around \$567 billion in 2022, and projected to reach \$836 billion by 2026 (Lebow, 2023) – funds much of today’s everyday digital world. The online advertising market helps shape economies, societies and politics globally as it interacts with issues of privacy, market evolution and platform power. It generates around 400 billion ad impressions daily, each of them seeking the user’s money, vote or attention, aiming to turn her/him into a customer, voter or spectator.¹ These ads compete with each other: in the second or so between a user’s online action and an ad being shown in response, there is usually a rapid, automated auction of the advertising opportunity. Billions of these auctions take place daily: there is no reliable estimate of their exact number. Their material organization is the focus of this paper.

Practitioners and their organizations (e.g. IAB, 2022) often divide online advertising into three main segments: ‘search’ (the ads that accompany the results of a query entered into Google or another search engine), ‘video’ (the ads on a video or streaming platform such as YouTube), and ‘display’, the ads shown on a social media platform, within a mobile app or on a website.² In this paper we examine what practitioners call ‘open display’: display ads shown not within the confines of a single platform such as Facebook, but via a multiplatform configuration that brings together multiple app or website publishers, multiple advertisers, and a host of intermediary organizations.³

The issue on which we focus is a conflict between two different ways of configuring display advertising, which we refer to as ‘centralized’ and ‘decentralized’. While initially decentralized, open-display advertising had become largely centralized by around 2013. At its heart were, and are, two interlinked Google platforms. One is Google’s ‘ad server’, DFP or DoubleClick for Publishers, which takes the final decision of which ad to show to which user. DFP operates on behalf of publishers, and can be programmed by them, but within constraints set by Google.⁴ The other is Google’s ad exchange, AdX, the main global platform for trading display-advertising opportunities.⁵

In this centralized approach, Google’s DFP is alerted electronically as soon as, e.g. a user accesses a webpage. DFP then seeks bids via AdX from algorithms, often Google’s own, acting on advertisers’ behalf, and within less than a second decides which ad(s) to show the user. From around 2014, however, that centralized configuration has been challenged, partially successfully, by a system of decentralized auctions known as ‘header bidding’, the chief topic of our paper. In header bidding, much of the auctioning process takes place outside of Google’s systems, with (in the technique’s canonical form) the

user's own phone, tablet or laptop gathering bids for the right to show her/him an ad or ads in around a second's time.

In examining Google's centralized system and the decentralized challenge to it, we make two arguments. First, an ad platform such as AdX/DFP cannot be understood as a mere two-sided or even multi-sided market. Such a platform is not a neutral 'market-plumbing' architecture that brings together an already existent buyer and seller, each with their proposed price. As Caliskan (2020) argues in the case of cryptocurrency exchanges, a platform is best conceived of as 'stacking' – layering together in an interdependent configuration – a variety of processes of 'economization' in the sense of Caliskan and Callon (2009).

The processes of economization stacked in platforms do not always take the form of marketization. Consider, for example, platforms' pervasive 'gifting' of 'free' services. Drawing on the classic anthropological literature on gifts, Caliskan and Callon (2009) note that, unlike a commodity as classically conceived, a gift 'circulates while preserving the presence of its giver embedded within it' (p. 387). As is well known, the always-present giver, the platform, extracts personal data (see, e.g. Fourcade & Klutzz, 2020). Platforms often reassure users that they will never 'sell' that data, and it does indeed usually remain 'untraded' (Beauvisage & Mellet, 2020, p. 92): attempts to commodify and marketize it have mostly failed. But that does not stop the 'economization' of personal data: the turning of it into an economic resource, an 'asset' in the sense of Birch and Muniesa (2020).

Second, we argue that innovation in online advertising is not 'linear' improvement driven by, e.g. cost-efficiency; rather, it is contestable and often contested. Advertisers *are* concerned with cost-effectiveness (Goldman, 2006), and digital technologies often reduce costs (Goldfarb & Tucker, 2019), but innovation's effects are often ambiguous: each individual innovation may offer increased efficiency, but aggregate fees paid by advertisers have been slow to fall, and advertising's carbon emissions remain alarmingly high (MacKenzie, 2023). Contestation within online advertising is rife. Controversy notoriously swirls around platforms' use of personal data. So too for the process of economization on which we focus: the structuring of how ad slots, in other words opportunities to show ads, are bought and sold. Conflict about this has been much more private, but different ways of materially organizing ad auctions structure the economic space of open-display advertising differently, advantaging some players and disadvantaging others. This makes the auctioning process a site of 'material politics' in the sense of Law and Mol (2008).

Jockeying for economic advantage is a crucial aspect of that material politics. If that seems too narrow a concern to merit the term 'politics', let us note, first, that income from open-display advertising is often crucial to news journalism. Second, also at stake in the material politics of header bidding is market power, and conceivably even the partial breakup of Google, the defining corporation of the information age. Third, ad auctioning's material politics has recently taken on a new aspect: environmental. Whether a single advertising opportunity (one

user, one ‘slot’ in which to show a display ad) gives rise to a small number of auctions, as in Google’s centralized system, or several dozen and perhaps as many as a thousand auctions, as in header bidding, affects digital advertising’s already substantial carbon footprint. An automated auction is not an abstract economic operation, but a material computational process that involves carbon emissions even before an ad is shown (MacKenzie, 2023).

Our two arguments – that platforms are stacks of economization processes, and that those processes are the site of material politics – are tightly inter-related. Material politics influences how processes are stacked, and stacking shapes material politics. When it concerns platforms, a ‘horizontal’ struggle over a particular economization process (such as our focus here, the structuring of ad trading) is unlikely to be self-contained, but will be subject to the ‘vertical’ influence of stacking on the process in question. For example, Google’s integrated system is, in practitioners’ terminology, ‘full-stack’: it incorporates almost the full range of display advertising’s economization processes. That has given Google greater capacity than other market participants to shape ad trading. And, as we shall show, the decentralizing challenge to Google’s integrated system from header bidding depended upon it turning out to be possible, albeit labour-intensive, to deploy the mechanisms of one economization process in the stack, targeting, to reshape another: the structuring of ad trading.

Our study of header bidding is a contribution to what is becoming a lively literature in economic sociology and related fields on online advertising. World-count limits prevent a full review, but let us briefly position our contribution with respect to it. By far the most influential single work discussing online advertising is Zuboff (2019), which however abstracts away almost completely from the material operations of what she calls ‘markets in future behaviour’ such as digital advertising, and, relatedly, tends to over-simplify and implicitly exaggerate Google’s role and power.

Two particularly important clusters of literature take us closer to advertising’s contested materiality. One is the research of Robert Cluley (e.g. Cluley, 2018, 2020; Cluley & Brown, 2015), which focuses, *inter alia*, on crucial controversies such as over privacy initiatives and how to define whether an ad is materially ‘viewable’ by the user. The other is the work of Jean-Samuel Beuscart, Kevin Mellet and Thomas Beauvisage (e.g. Beauvisage & Mellet, 2020; Beuscart & Mellet, 2013). In particular, Mellet and Beauvisage (2020) is by far the best social-science study of ‘cookies’, the single most important user-tracking tool employed in digital advertising.

Our central concern – the mechanisms via which ads are bought and sold – is, nevertheless, not a primary focus of either of these bodies of literature, and certainly not of Zuboff (2019). Our historical narrative does, however, draw upon Crain’s (2021) fine history of the early years of digital advertising. Alaimo and Kallinikos (2018) helpfully examine the standard that defines the digital objects that are traded in open-display ad markets. However, Viljoen *et al.* (2021) is to our knowledge the only sociological, as distinct from e.g.

economic, investigation of how advertising auctions are structured. They focus on the ‘mechanism design’ tradition in academic economics and the legitimization work it performs, although their insightful article perhaps does not press far enough into the materiality of auctions. However, the literature of competition law, notably Srinivasan (2020), contains useful discussions of header bidding and especially of Google’s reaction to it.⁶

Header bidding is, as Srinivasan (2020) suggests, a form of resistance to Google’s centralized system, making relevant other work on resistance to platforms, such as Cotter’s (2019) and O’Meara’s (2019) analyses of online influencers’ practices or Ziewitz’s (2019) discussion of the contested legitimacy of the ways search-engine ‘optimizers’ improve their clients’ positions in search results. Header bidding, however, goes beyond such practices of ‘external’ resistance to encompass a *recoding* of processes crucial to platform capitalism. Header bidding is indeed quite literally recoding: it requires the production and use of sophisticated new software, and this makes relevant the wider literature on software, digital life and platforms, such as Mackenzie (2006), Langley and Leyshon (2017, 2021), Amoore (2020), Stark and Pais (2020), Burrell and Fourcade (2021), and van der Vlist *et al.* (2022).⁷ For example, the culture of software engineering influences both how those involved talk about header bidding – in particular how, counterintuitively, its *proponents* often call it ‘a hack’, although an entirely legal one – and how it evolved into an open-source software project, a form of entangling gift familiar to, and well-regarded by, software engineers.

Like Mackenzie (2006) and Dourish (2017), though, we also emphasize coding’s materiality. Crucial to header bidding’s decentralizing challenge to Google’s centralized system are two down-to-earth material issues. First, *when* should the electronic call to Google’s ad server DFP happen: immediately, or only after a few tenths of a second? Second, *where* should bids be collected: by Google’s systems, or on the user’s phone/laptop or perhaps a server designated by the publisher? On these two apparently mundane aspects of material practices hinges the techno-economic structuring of open-display advertising, and thus, e.g. publishers’ revenues and even questions of corporate power. Although the classic studies of material politics (such as Barry, 2013, and Gabrys *et al.*, 2013) are of non-digital forms of materiality, digital systems are equally material and can therefore, as with header bidding, also be domains of material politics.

All forms of politics have material aspects and potentially influence the material world, but that does not mean that all politics should be considered to be material politics. Just as Barry (2002) rejects over-extended notions of ‘politics’, so we restrict the term ‘material politics’ to situations in which the (re)organization of the material world and/or material practices is primary, either as a means of achieving other goals or as an end in itself. For example, we touch briefly on lawsuits, largely triggered by Google’s reaction to header bidding, which at root involve how narrowly or broadly to define the behaviour or market power that violates competition law. We think of activism about that issue simply as ‘legal politics’: if successful, it may have material consequences,

but that does not make it material politics. On the other hand, environmental politics, which as noted is starting to impinge upon header bidding, surely *is* material politics, because it seeks consciously and directly to reshape material practices.

After this introduction and a section on data sources, our third section examines display advertising's initial decentralization and how, in contrast, Google constructed its largely centralized, integrated system. Section four examines header bidding's rise and its challenge to that system. The fifth section discusses Google's reaction to header bidding, the limitations on header bidding's success, and the challenge header bidding is facing from environmental politics. The sixth section is the conclusion.

Data sources

This paper reports early results from multi-year research on digital advertising, which has taken two main forms. First, we have participated in four face-to-face and 12 online sector meetings, along with 29 more specific online events such as webinars. Public presentations to these meetings and events are often thinly-disguised marketing pitches for particular advertising technologies or platforms. Nevertheless, informal interactions in the sidelines of meetings, and even some formal presentations, sensitized us to divides in the sector, for example between independent advertising-technology or AdTech firms and the big-tech platforms (especially Google/Alphabet, Facebook/Meta and Amazon), and between those big platforms and other publishers, especially news publishers.

Second, we have conducted 91 interviews with 69 advertising practitioners: 34 working in the United States, 25 in the United Kingdom, nine in continental Europe, and one in India. Their roles range across digital advertising, from advertiser and advertising agency to publisher and platform, including the demand-side and supply-side platforms discussed below. Interviews are typically 45–60 min long; all but three were audio recorded and transcribed. We are not primarily seeking to elicit an interviewee's general 'attitudes' but to gain insight into issues such as: the systems, data and material practices s/he employs; what s/he is trying to achieve with them; the challenges and obstacles faced; and the internal divides in the sector that s/he perceives. Interviews are only semi-structured: we focus the questions for each interviewee on the systems and issues with which s/he has most experience.

Our research is iterative, with interviews leading us to analytically interesting topics, which we then focus on in greater depth. Header bidding is one of those topics: our early interviews revealed it to be a central site of both the above tensions. We therefore particularly tried to locate, largely by snowballing, people involved with header bidding, finding 21 such interviewees; 27 of our 91 interviews are with them. In what follows, we anonymize interviewees by using two-letter labels, with the exception of AdTech entrepreneur Brian O'Kelley, whose specific roles render him identifiable and who has kindly allowed us to name

him. We supplemented these interviews with the use of documentary sources such as the trade press, especially *Ad Exchanger*, which has reported in greatest depth on header bidding, and regulatory reports, the most detailed of which is by the UK's Competition and Markets Authority (2020).

Our analysis of these data sources focused on constructing a robust narrative history of header bidding's context and emergence, its taking the form of an open-source software project, its implementation by publishers, resistance to it, etc. That involved extensive triangulation of what individual interviewees told us against what others say, and, where possible, also against documentary sources. While we have generally found consistency, we need to be clear about methodological limitations. Almost all the 27 header-bidding interviews were with people who could broadly be described as supporters of it. Indeed, their willingness to be interviewed often seems related to their self-perception of having been involved in a crucial challenge to the techno-economic *status quo*. Similarly, *Ad Exchanger's* reporting on header bidding, while always fully professional, has typically been sympathetic.

It has been much harder to obtain interviews with staff of the big platforms, which often forbid employees speaking to outsiders about controversial topics. Google's reaction to header bidding is one such topic, severely limiting our interview-based data on Google's decision-making processes. We are therefore deliberately agnostic about those processes and thus about why Google took the actions it did regarding header bidding. As parties to what they frequently saw as conflict with Google, other interviewees, especially from news publishers, often suggested scathingly antipathetical interpretations of Google's actions, but we do not adopt those interpretations as ours.

We are also careful to avoid a methodological danger evident, e.g. in competition-law research on digital advertising: the tendency, in the absence of contrary evidence, to reason 'backwards' from a decision to an imputed unitary corporate intent. Any corporation of any size is divided into departments or teams, which may well have different priorities and preferences, and our fieldwork provides tentative evidence that this is indeed the case for Google. A decision that looks as if it is 'by Google' may therefore perfectly conceivably be the result of private internal conflict or compromise rather than of any unitary motivation, and that is another reason for our agnosticism on the causes of Google's decisions.

A heterogeneous stack versus 'a single and complete advertising system'

Online display advertising began in the mid-1990s, and for at least its first decade it was largely decentralized and heterogeneous. Initially, indeed, much of it was simply a patchwork of direct bilateral deals between publishers and advertisers. If a high-prestige outlet such as a prominent news publisher set up a website, it would aim to negotiate deals with advertising agencies acting for

large advertisers at an agreed-upon price, typically several dollars for every thousand impressions. Interviewee BZ's characterization of the process as 'schmooze and booze', or BW's 'going out to lunch, drinking a martini', are stereotypes, but deals were indeed negotiated person-to-person.

Even prestigious publishers, though, often could not fill every online ad slot with direct deals, creating the need to find ads to fill 'what we call[ed] "remnant" ... the leftovers' (O'Kelley interview). 'Schmooze and booze' played a much smaller role in that effort. Unfilled 'remnant' ad slots were sold to 'ad networks', the first main category of intermediary between publishers and advertisers. Ad networks also handled advertising for newly established/low-prestige websites and apps that found it difficult to attract direct deals. Networks made their money by buying ad slots cheaply, paying publishers mere 'pennies' (interviewee BI) per thousand slots, and charging advertisers a substantial margin on top of that to have the network show their ads in those slots.

Ad networks' business model was thus simple enough, but crucial underlying platforms – ad servers – often were not: indeed, all but the simplest ad servers clearly exemplify Caliskan's (2020) 'stacking' of economization processes. An ad server is the system via which ads that advertisers or agencies have created are delivered to users' screens, and it controls and keeps track of that process, ensuring that the publisher's or ad network's commitments to advertisers are met and that the advertiser can be billed. A publisher will often input its classification of its content into its ad server and program it with its rules, for example about the content ads should appear alongside and any advertisers, such as tobacco companies, from which it does not accept ads. In recent years, ad servers have also needed, e.g., to store and implement 'consent strings': the digital evidence of a user having consented to tracking.

Even in the early years of display advertising, for a publisher or ad network to build its own ad server was a substantial software-development project. Some large publishers did this, but ad serving was also available as a commercial service, most importantly from the New York-based ad network DoubleClick, founded in 1996 (Crain, 2021, p. 63). DoubleClick was the best-known pioneer of the first of the two most crucial economization processes that are 'stacked' in ad serving: the turning of the gigantic, diverse, poorly structured, and poorly recorded body of online actions into a set of knowable users, who thus become economically valuable targets for advertising.

DoubleClick was famous within advertising above all for its cookies. When a user visited a website that employed DoubleClick's ad server, that server would, as well as rendering ads on her/his screen, also deposit a cookie – essentially, a short string of letters and numbers unique to each user – in the user's browser. When the user then visited a different website that also used DoubleClick's ad server, that server could retrieve the cookie and recognize that it was most likely the same user, enabling DoubleClick gradually to build a picture of her/his behaviour across websites. Its ad server could then use that knowledge to choose which ads to show that user.

The second crucial economization process stacked in ad servers such as DoubleClick's is structuring the buying and selling of online ads. If a publisher and advertiser strike a direct deal, junior staff then have to 'traffic' the deal, using keyboard and mouse to set up the corresponding 'line-items' (the basic components of an advertising campaign) in the ad server. As users visited the publisher's website, early ad servers would often show them ads in essentially a set order: 'It was deliver A [an ad for direct-deal A], then B, then C, then rotate through those' (interviewee BW). If, however, there was no direct deal to fill an ad slot, the server would trigger a lower-priority line-item that offered the 'remnant' slot to a pre-selected ad network. If it chose not to buy it, 'they would send it to the next one [again in a fixed "daisy chain" or "waterfall"]', and the next one, and the next one ...' (O'Kelley interview).

These two elements of the AdTech stack – sophisticated targeting and relatively inflexible ways of selling ad slots – were implicitly in tension. Among those who recognized this was computer scientist Brian O'Kelley of the ad network Right Media, who developed an ad server that employed machine learning to estimate the economic value of showing a specific ad to a particular user, and in effect ran an internal auction, in which these predicted values were treated as equivalent to bids. O'Kelley and his colleagues then turned that ad server into what they thought of as an ad exchange (RMX Direct, launched in August 2006), by allowing other ad networks also to 'input their supply and demand' (interviewee CB), in other words the ad slots they wanted to sell or buy. The ad server 'would do the rest' (interviewee CB), allocating each individual slot – in real time (interviewee BW) – to the network that would pay the most.

Ad exchanges were soon joined by other new layers in the evolving AdTech 'stack'. These include 'demand-side platforms' or DSPs, which bid on exchanges on behalf of advertisers or advertising agencies, and 'supply-side platforms' or SSPs, discussed below. By 2010, independent advertising-technology firms were sufficiently numerous and heterogeneous that what was to become a famous graphical portrayal of the sector, the Lumascape, was already visually daunting.

Alongside this loosely interconnected sector, though, a quite different approach to display advertising was being built: a less heterogeneous, more tightly integrated stack, owned by a single company, Google. Founded in 1998, Google initially focused on search advertising, but in 2003 entered display advertising, launching a highly-automated ad network, AdSense, which enabled even small, technically unsophisticated publishers to sell ad slots.

AdSense was a stepping-stone. In their 2006 annual letter to investors, Google's co-founders wrote: 'Our goal is to create a single and complete advertising system' (Page & Brin, 2006).⁸ Statements of grand corporate ambition need a pinch of salt, but Google did indeed go on to buy both an SSP and DSP, and, most importantly, the best-known independent AdTech firm, DoubleClick. That \$3.1 billion purchase was announced in April 2007, although

competition-law scrutiny by the US Federal Trade Commission and European Commission delayed the acquisition to March 2008.

Buying DoubleClick brought Google a nascent ad exchange, AdX, and, crucially, DoubleClick's ad server. In approving the acquisition, European policy-makers reasoned that ad-serving fees were a small proportion of advertising's costs, and other ad servers were available (Brockhoff *et al.*, 2008, p. 58). That economic reasoning, however, hugely underestimated the significance of ad servers' materiality. The material forms taken by display advertising were changing fast (with, e.g. the rise of video ads and ads on mobiles), so ad-server developers had 'to be constantly innovating ... to support the latest ... thing' (interviewee BR). Ad servers, though, also needed to be scalable: to serve billions of ads daily, they must run on networks of hundreds, thousands, or tens of thousands of physical machines. Almost all developers of ad servers struggled with the twin challenge.

The exception was the ad server that Google acquired from DoubleClick, to which it devoted two years of intensive development, 'rebuil[ding] ... the entire code base from scratch' (Nolet, 2010). Over the previous decade, Google had learned – not without struggles of its own – how to make systems run at giant scale, even itself designing and assembling the million or so machines needed by the massive data centres it had started to build. Google kept DoubleClick's name, calling its ad server DoubleClick for Publishers or DFP, but it was actually thoroughly re-engineered, as was AdX, the ad exchange Google inherited from DoubleClick.

DFP's most important capacity, which has become crucial to many publishers' ad revenues, was 'dynamic allocation': automated real-time choice between the highest real-time bids for a publisher's ad slots and the demands of the publisher's direct deals with advertisers. A direct deal involves the publisher agreeing to show a certain number of the advertiser's ads, but if DFP is confident that this contractual obligation will be met, it will choose a real-time bid if it is higher than the direct-deal price.

Google's new ad server, DFP, thus fused together the previously largely distinct economic worlds of direct deals and 'remnant'. But the real-time market that DFP integrated with direct deals was not any market: it was Google's exchange, AdX. Bids to show ads flowed into AdX in huge volume through Google's 'stack', via the algorithmic ad-buying systems Google made available to both small and large advertisers. AdX rapidly became 'by far the largest source of exchange inventory' (McDermott, 2014), making Google's integrated system, the core of which is shown in [Figure 1](#), a powerful, indeed potentially overwhelming, rival to the more heterogeneous independent sector. The US Federal Trade Commission had approved Google's merger with DoubleClick, but it began to have competition-law concerns, with one of its lawyers privately contacting independent AdTech firms (McDermott, 2014). Her informal enquiries, though, did not generate a lawsuit. Serious legal challenges to Google's centralized, integrated display-advertising system were to come, but only after that system had been challenged materially by header bidding.

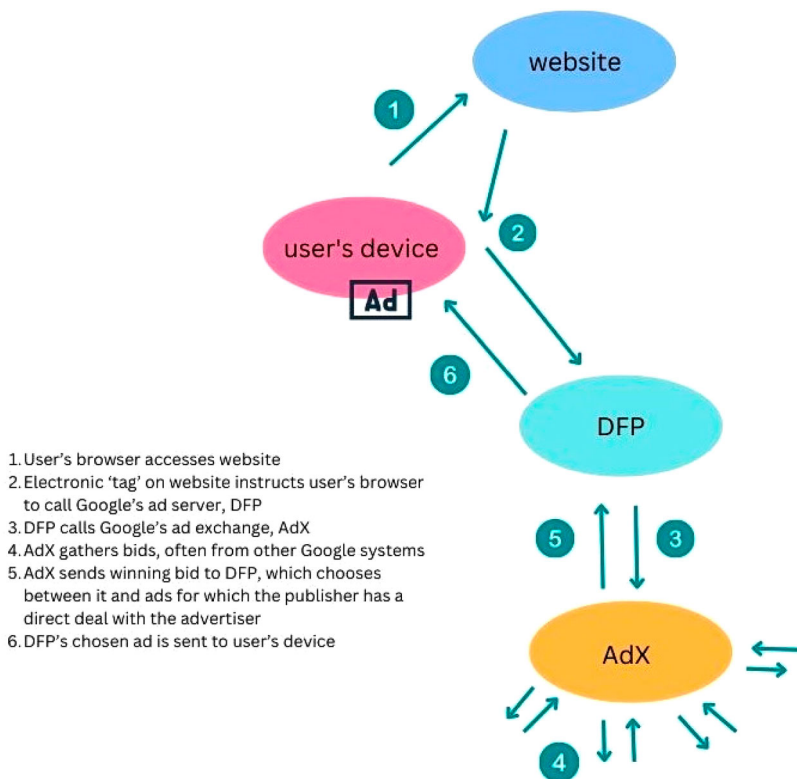


Figure 1 The core of Google's integrated display-advertising system (simplified). Note the central roles of Google's AdX, which gathers real-time bids for publishers' ad slots, and DFP, which weighs up these bids against publishers' direct deals with advertisers

Material politics I: The rise of header bidding

Header bidding gets its name from its use of webpages' normally invisible 'headers' to instruct users' browsers to gather bids from sources other than Google's AdX. What gave force to previously *ad hoc* ways of doing this was dissatisfaction with the techno-economic *status quo* among one particular category of firm in the independent AdTech sector: supply-side platforms. SSPs aggregate publishers' ad slots, selling them in what is in effect the SSP's own exchange. For example, a Los Angeles-based SSP, the Rubicon Project, helped the *Guardian* build its presence in online advertising, enabling the newspaper to keep its website accessible free-of-charge.

Unlike Google's AdX, though, SSPs could not bid directly in real time into Google's ad server, DFP. So individual SSPs started to try to persuade publishers to insert code into their webpage headers to instruct users' browsers

to request a bid from that SSP prior to electronically calling DFP. The publishers of those webpages, at first often suspicious, began to see header bidding as potentially a useful way of getting higher prices for their ad slots than they were receiving from Google's AdX.

As header bidding gained momentum, it faced, but largely overcame, three significant problems. The first was a 'general lack of trust ... across the industry', as interviewee BR puts it: 'there's a lot of things that can happen where money is potentially lost in the friction of the pipes or whatever'. Not only did publishers not always trust SSPs, but SSPs did not trust each other enough to coordinate their actions. Individual SSPs separately approached each publisher, and the code that each SSP asked the publisher to incorporate into its webpage headers was specific to that SSP and would cause the user's browser to request a bid only from that particular SSP.

The lack of collective action had a potentially disastrous effect on how long it took webpages to load. 'There's this terrible publisher's freak-out about this', says interviewee BR, 'a really bad user experience'. The 'scripts' containing the code that instructed the user's browser to gather a bid from each SSP could each take around a second to load, and, worse, they loaded one after the other, with 'every single script ... block[ing] the page [from loading]' (interviewee BQ), until it had finished executing. Nor was there an incentive for individual SSPs to make their scripts faster: 'it's to the SSP's advantage to make [load time] really long so that they can collect as many bids as possible [from demand-side platforms]', says interviewee BQ.

One way to control the process was to create a software 'wrapper' that would load SSPs' scripts simultaneously, timing out of the auction any that loaded too slowly. For publishers, though, creating and maintaining such a wrapper would be onerous (interviewee BV), so in practice it seemed as if an SSP had to do it. That, though, failed to solve the underlying material-politics problem, because it raised the question of whether the SSP that wrote the wrapper would give itself a surreptitious advantage, such as more time to gather bids (interviewee BX).

But perhaps trust in a header-bidding wrapper could be built by making its code public as open-source software, so that it could be scrutinized to see if it was doing anything underhand? 'Open-sourcing makes a lot of sense ... for political reasons', says interviewee BQ. That, however, meant in practice making the software a gift: freely available, free of charge. The engineers who initiated the open-source project (who worked for the AdTech platform and SSP, AppNexus, co-founded by ad-exchange pioneer Brian O'Kelley) faced internal opposition because of that: 'there was a lot of internal friction about why are we spending time on this [the wrapper]', says interviewee BR.

They took their case to O'Kelley, AppNexus's CEO. 'They're like ... I know this is our secret sauce, but I think we should open-source it' (O'Kelley interview). 'I'm, like, you guys are crazy', he says, but he realized, though, that 'if everyone competes on header, it's not going to work very well', and Google might launch a more successful alternative, which he calls a 'header-bidding killer':

[S]cenario one, we get some short-term revenue but then we lose to Google. Scenario two, we open-source this, it becomes a platform, and everyone will build around it, then that will beat Google.

Only as the meeting ended, recalls O’Kelley, did the engineers reveal to him that they had actually already released an open-source version of the wrapper.

The engineers called their open-source project Prebid, the name signalling that the user’s device would gather bids from SSPs before Google’s ad server was alerted. Getting support for Prebid from SSPs that competed with AppNexus was, however, initially very hard. There was ‘a lot of drama in the first year’, says interviewee BQ, ‘lots of screaming and yelling’. Making the code open-source seemed at first to ‘make the political [problems] even worse, because now everybody can see what you’re doing’ (interviewee BR). The wrapper needed to contain scripts that would seek bids from SSPs other than AppNexus, so the latter’s engineers needed to ‘look at their [other SSPs’] code ... see how are they doing this’, and then rewrite the code (interviewee BR). AppNexus’s lawyers ‘gave [the engineers] a green light because the code [they] used is already in the public domain’, says interviewee BQ, but that did not stop other SSPs ‘telling ... publishers they [AppNexus] were stealing our code’, in one case even threatening legal action (interviewee BR). Publishers, though, began to accept the gift, adopting the free open-source wrapper: ‘it was a gradual thing’, says interviewee BR. ‘It wasn’t overnight’.

What became the canonical form of header bidding crystallized around Prebid’s open-source software. Examples of the auctioning of opportunities to display ads to the first author are shown in Figures 2 and 3, and the

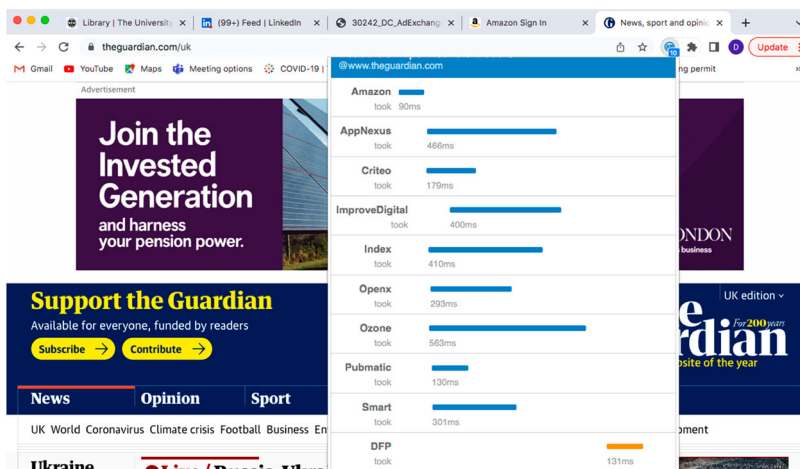


Figure 2 Calls for bids prompted by visiting the *Guardian*'s homepage, 22 April 2022, captured using the Chrome extension, Headerbid Expert. Times shown are in milliseconds (thousandths of a second)

| Bidder | Ad | \$ CPM | Won | Timing |
|----------------|-----------|----------|-----|------------|
| pubmatic | Q 970x250 | US\$3.34 | ✓ | Green bar |
| improvedigital | Q 970x250 | US\$0.05 | | Green bar |
| oxd | Q 728x90 | US\$2.73 | | Green bar |
| and | Q 970x250 | US\$0.09 | | Green bar |
| criteo | Q 970x250 | US\$0.09 | | Green bar |
| ozone | Q 970x250 | US\$2.14 | | Green bar |
| ozone | Q 728x90 | US\$0.43 | | Green bar |
| ix | Q 728x90 | US\$2.94 | | Yellow bar |
| ix | Q 970x250 | US\$1.17 | | Yellow bar |
| pubmatic | Q 970x250 | US\$1.04 | | Green bar |

Figure 3 Examples of bids in header bidding, captured using the Chrome extension, BidFilter. Although what is being bid for is the right to show one ad, display advertising’s convention is to quote the price for a thousand ads (CPM or cost per mille). E.g. the actual ad-slot price corresponding to the winning bid (\$3.34 CPM) is 0.334 cents

underlying header-bidding process (which takes around half a second) is sketched in Figure 4.⁹ Directed by Prebid code in the webpage header, the user’s browser requests bids from supply-side platforms with which the website’s publisher has a commercial relationship. Each of those SSPs then conducts a ‘sub-auction’, seeking bids from the demand-side platforms (DSPs) with which it works, and sends the sub-auction’s winning bid to the user’s browser. The browser forwards the highest of the SSPs’ bids to the ad server, and Google’s AdX can beat it only if it can produce a higher bid. Hence, header bidding’s attraction to publishers: the possibility of their receiving higher bids than AdX on its own would produce.

The main source of an SSP’s income is the fee it gets from publishers if a bid received via that SSP is successful. Adoption of Prebid by publishers therefore gave initially hostile SSPs an incentive to improve how it solicited their bids. ‘[A] developer from the SSP ... they’d take a look at it and be like, you know what: you could actually do it this way. And they’d start giving us little tips and tricks to improve it’ (interviewee BR). As with other gifts, open-source software such as Prebid can indeed ‘generate attachments and habits’ (Fourcade & Klutzz, 2020, p. 5); power can come not from exclusion or coercion, but from ‘generosity’ and ‘co-production’ (Eyal, 2013, p. 875).

The second problem faced by header bidding is implicitly signalled by how, as already noted, its proponents often describe it as a ‘hack’: e.g. ‘a hack of the Google ad server’ (interviewee BF); ‘hack[ing] the [ad server’s] priority setting in order for it [header bidding] to work’ (interviewee BQ); ‘[t]he whole process is a hack’ (interviewee BQ); ‘header bidding is a hack, that’s the truth’

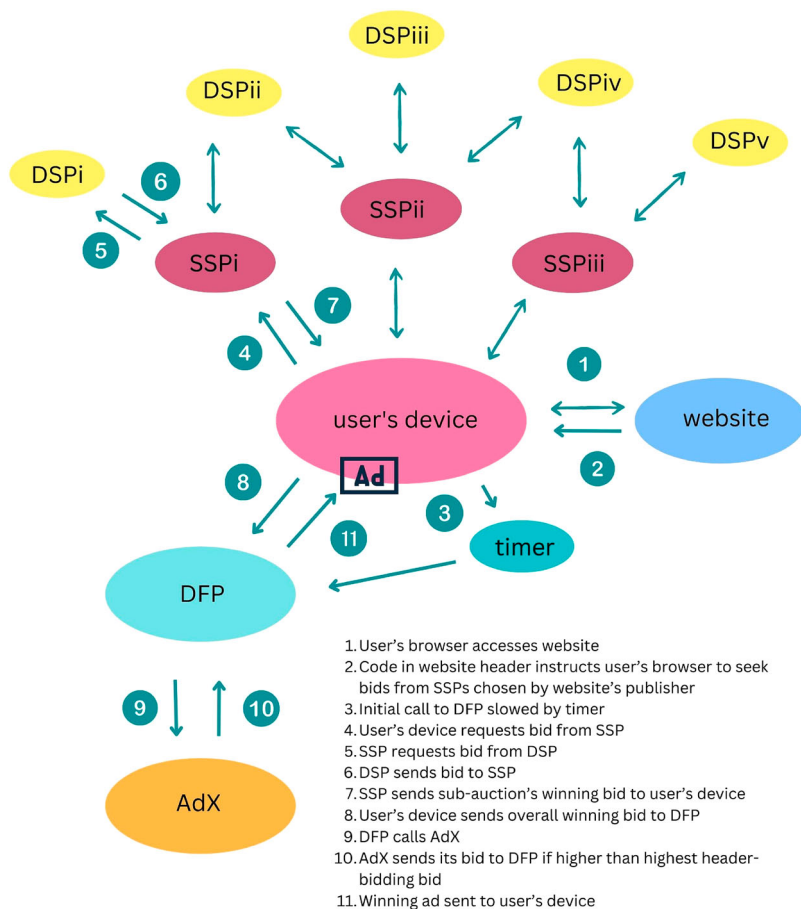


Figure 4 Header bidding (simplified)

(interviewee BR); ‘a hack’ (interviewee BS); ‘a bit of a hack’ (interviewee BY). In software engineering’s culture, ‘hack’ is a richly polysemic word. On the one hand, it signals expertise: the ability to ‘accomplish [a] task ... faster’ (Eyal, 2013, p. 869), and thus membership of what Burrell and Fourcade (2021, p. 217) call ‘the coding elite’. On the other hand, calling something one has done ‘a hack’ is self-deprecating. Prebid’s code might have been written quickly and well, but the process it triggered could be considered inelegant, cumbersome, even ungainly.

At least at first, header bidding required extensive, tediously repetitive work by members of publishers’ staff not in the ‘coding elite’. There was no direct way of representing the winning bid in header bidding within Google’s ad server, DFP, so ‘header bidding tricked the ad server’, as interviewee BJ puts it: header bidding was a ‘hack’ in that sense too. The winning bid is

inserted into the ad server by representing bid prices using the mechanism DFP provides to permit ads to be targeted, e.g. by specifying a target audience. Header bidding thus involves, as already noted, taking one ‘stacked’ economization process, targeting, and repurposing it to alter another, the structuring of ad trading. Doing this required the labour-intensive creation of a separate line-item – as noted, the basic component of an advertising campaign – for every possible bid price, and sometimes repeating the process for every possible bidder. At least thousands of line-items were needed, sometimes tens or even hundreds of thousands. Eventually, the process was automated, but until that happened it was a real constraint on the adoption of header bidding.

The third obstacle to header bidding – and one aspect of Google’s criticism of the technique – was the delay involved in the user’s phone, tablet or laptop becoming an auctioneer. In 2015, Google executive Jonathan Bellack publicly debated header-bidding advocate Tom Shields.¹⁰ Increased competition for ad slots could better be achieved, Bellack argued, by ‘using the server-to-server programmatic pipes that work so well’, enabling auctions to take only around a tenth of a second. Header bidding, employing a user’s device, could take five times as long.

The objection was serious: delay increases the risk of users leaving a page before ads are visible. Header bidding’s proponents responded by developing a second version of the technique, ‘server-side’ header bidding, in which a server designated by the publisher gathers the bids, faster than the user’s phone or laptop could. The user’s device, though, is politically neutral terrain, so to speak – it is not under the control of any SSP – while whatever organization ran the server would have a potential advantage. ‘It did get a little political on that piece on that,’ says interviewee BR: ‘when you move the code to the server ... You don’t really know what’s happening’. No full solution has yet emerged, so British news publishers from across the political spectrum have taken collective action, establishing a collaborative AdTech effort, the Ozone Project, which has developed and hosts a server-side header-bidding ‘wrapper’.

Material politics II: Conflicts over header bidding

As header bidding gathered support, says interviewee BR, ‘we were all curious and wondering what Google’s response would be ... just watching and waiting to see.’ This section examines (a) Google’s response; (b) the main limitation on header bidding’s success, its failure to displace Google’s ad server, DFP; and (c) the potential impact on header bidding of carbon-focused environmental politics.

Interviewees, mainly external to Google, told us that when header bidding emerged Google’s reaction was not enthusiastic, and that is consistent with the limited public evidence about that reaction, such as the debate referred to in the previous section, or Google’s (2020) assertion that ‘[h]eader bidding

is characterized by increased latency [electronic delay], reduced transparency and significant user trust and privacy concerns', the last of these because requests for bids often contain cookies or other identifiers of users, and in header bidding these requests are circulated widely (Google LLC, 2020, p. 10).

Most consequentially, however, Google also developed what is in effect an alternative to header bidding: 'Exchange Building', later renamed 'Open Bidding'. Like header bidding, Exchange Bidding 'give[s] people [publishers' systems] a way to call out to multiple auctions' (interviewee AA), but without the user's device collecting bids. Instead, this happens largely 'within the Google ecosystem' (interviewee AA). The UK Competition and Markets Authority suggests that 'a major reason' for Google developing Exchange Bidding 'appears to have been protecting Google's revenues from the impact of header bidding' (Competitions and Markets Authority, 2020, p. M10), but our interviews do not contain direct evidence of this. However, Google's internal name for its Exchange Bidding programme, 'Jedi', did have something of a conflictual, even a 'mind-trickster', connotation, and the company seems to have worked hard, in an effort that Google called 'Jedi Blue' – seemingly a reference to the colour of Facebook's logo – to ensure that Facebook, which had flirted with header bidding, embraced Google's Exchange Bidding instead. 'We thought we were going to get Facebook in Prebid [header bidding]', interviewee CB told us, but 'they ... went dark'.

Jedi Blue, and Google's reaction to header bidding more generally, are central to a number of lawsuits the company faces. In December 2020, Texas and nine other states with Republican Attorneys General began legal action against Google, accusing it of 'unlawfully foreclos[ing] competition' in online advertising (Texas *et al.*, 2022, p. 84), while the US Justice Department and eight different, predominantly Democratic, states launched a similar lawsuit in January 2023, and there is analogous litigation in Europe too. Google contests such lawsuits vigorously, and the stakes are high, because the litigants seem to be envisaging a partial breakup of Google, at least of its centralized, integrated display-advertising system. The suit led by the Justice Department demands that, at a minimum, Google divests itself of the two systems on which we have focused, DFP and AdX, and the European Commission seems to be seeking a similar outcome.

It is, of course, far from clear that these lawsuits will succeed. However, repeated lawsuits, even if they fail, can nevertheless affect their targets. An interviewee reports telephoning a contact in Google, suggesting a joint initiative to reduce advertising's carbon emissions:

[I said] 'we should work together because I'm [pretty sure that you actually are the logically best solution to this' [see below]. And they said, 'hang up the phone!' I was like, 'what do you mean? [I]t's just us'.

Our interviewee's contact explained that 'every conversation has to have a lawyer in it' if it might possibly lead to legal action against Google. 'Google

is no longer an innovation threat’, says our interviewee, ‘because they’re so constrained by lawyers’.

In assessing Google’s reaction to header bidding, though, it is also important to highlight what Google seems largely *not* to have done. As already noted, publishers using header bidding need to create very large numbers of ‘line-items’ within Google’s ad server, DFP, in order to use the mechanisms of one economization process, ad targeting, to reshape another, ad trading. Google could thus have hamstrung header bidding very simply, by imposing a tight limit on how many line-items a publisher could create. The Texas-led lawsuit does allege that ‘Google purposefully limits’ the number of line-items ‘to foreclose competition from header bidding’ (Texas *et al.*, 2022, pp. 134–135), but our header-bidding interviews do not contain first-hand reports of Google doing this.

That is a reminder that material politics may encompass shades of grey, such as gradual accommodation to a perhaps not-entirely-welcome material reality, not just out-and-out conflict. Indeed, Google’s decisions about the material configuration of its systems are compatible with header-bidding proponent BR’s nuanced February 2022 characterization: ‘[t]hey’re not involved [in the Prebid open-source header bidding project]. Google obviously knows about it, and has come to accept it’. In 2022, indeed, Google’s acceptance of header bidding seemed to be taking material form: Google was testing an electronic ‘bridge’ to Prebid (Sluis, 2022). This would allow the winning header bid’s price to be represented directly within DFP, eliminating the cumbersome creation of thousands of line-items, although at the time of writing (June 2023) the bridge seems not yet to be publicly available.

‘Now’, header bidding is ‘like it’s infrastructure’, O’Kelley told us, and all the big publishers we interviewed were using it heavily. Crucially, however, header bidding has not succeeded in displacing the central machine of online display advertising, Google’s ad server, DFP. AppNexus, for which header bidding was one aspect of what those involved saw as a ‘full-stack’ challenge to Google, did try to do this. It acquired the ad server Open AdStream, and re-engineered it to better support header bidding and operate at global scale. Several continental European publishers – most prominently Axel Springer – adopted AppNexus’s new ad server. We originally hypothesized that this might be explained by particular hostility to Google, but interviewee BW says those publishers were still using older ad servers, not DFP, and were thus more open to AppNexus’s machine.

Most publishers, though, including all the Anglo-American publishers to whom we spoke, have kept using Google’s DFP alongside header building and Prebid. Interviewees identify two main reasons not to shift away from Google’s ad server, even in the case of one publishing group that owned a stake in AppNexus. One is that, as noted, a sophisticated ad server such as Google’s exemplifies Caliskan’s (2020) ‘stack economization’, performing a variety of crucial economic roles that are problematic to disrupt. This makes switching ad server daunting: Springer’s migration to AppNexus, seen as

smooth, took three months. The second reason is the material difficulty of representing real-time demand for ad slots from Google's AdX in any ad server other than DFP.¹¹ Interviewee BJ describes this difficulty as forming a 'moat' around DFP, while publisher BY says it was '[t]he big sticking point ... there was too much revenue risk associated' with a shift away from DFP.

That display advertising's crucial machine remains Google's DFP gives Google continuing influence, and its overall accommodation with header bidding does not preclude moments of sharp conflict. A salient example is 'unified pricing' modifications made by Google to DFP in 2019. The modifications materially constrain a publisher using DFP to have the same 'floor price' for all sources of demand for ad slots: many publishers previously set a higher floor for Google's AdX, believing it had greater capacity to raise its bids than SSPs had. Interviewee BY attended an April 2019 New York meeting in which Google staff told publishers about the coming change:

[I]t was a full-on uprising. Americans can be a little bit rude and very, very straightforward ... there was maybe 10 of us [publishers] that were like, you've got to be kidding me ... we were all just so angry.

One Google staff member present was deeply upset, reports interviewee BY: 'I guess the publishers just really went after him hard'. But the change to DFP went ahead. Even interviewee BY cannot see a viable alternative to DFP: 'we are stuck on this ad server'.

Furthermore, header bidding's success is in one respect self-undermining: this is where the politics of advertising's carbon emissions comes into play. Many publishers use both versions of the open-source software, Prebid: the first using the user's device; the second employing a server designated by the publisher. Doing the latter reduces electronic delays, making the trading of ad slots faster, but a 'stack' consideration, targeting, keeps publishers using the first procedure as well: if the user's device collects bids, potential bidders are in direct electronic contact with it, enabling them to retrieve their cookies. Other equivalents/variants of header bidding are also in widespread use, including Google's Exchange/Open Bidding and Amazon's 'Transparent Ad Marketplace'.

Using and fine-tuning this full array of ways of gathering bids can bring publishers economic benefits, as interviewee BH reports, but at the cost of multiplying auctions many times over, increasing carbon emissions considerably. Bid requests reach SSPs via multiple pathways, each request triggering the SSP to run a sub-auction. So a single opportunity (one user, one ad slot) can easily result in several dozen auctions, and in extreme cases as many as a thousand, so Brian O'Kelley tells us. Since, he reports, there are only around 40 big DSPs likely to bid on any large scale in these auctions, that is a hugely wasteful outcome. Google's centralized, integrated system is almost certainly much less carbon-intensive than header bidding's multiple, decentralized auctions,

especially given that Google's data centres are increasingly powered by electricity from renewable sources. For header bidding to respond adequately to this environmental concern will demand thoroughgoing restructuring of trading, requiring collective action on an unprecedented scale.

Conclusion

A bid to show an advertisement in an ad slot is a material event, not just a promised price. Physically *where* such bids are collected (On the user's device? On a computer server? Whose server?), and *when* (before Google's ad server, DFP, is called, or after?) are consequential issues of material politics. At stake are advantages that some market participants enjoy, disadvantages others suffer, and – via the US and European lawsuits – perhaps even whether Google can continue to operate as a unitary corporation. And auctioning is only one of the material processes that take place in the 'longest second' between the user's action and the ad(s) appearing on her/his screen. Tracking systems are triggered; cookies are 'synced'; an electronic call to a Data Management Platform may provide additional information about the user; and one or more verification systems assess whether the apparent user is actually a human being or a fraudulent 'bot', whether the human will be able to see the ad, and whether the surrounding content might be damaging for the brand whose advert it is.

All these processes raise questions of material politics, as does header bidding in contexts other than the one on which we focus, web advertising. For example, header bidding has been much less successful in smartphone apps, where existing AdTech tools, including Google's, are materially more deeply embedded than in the web. Above all, the processes of 'platform reintermediation' and 'platform consolidation', which Langley and Leyshon (2021) rightly consider as underpinning 'platform capitalisation', are material processes, and their imposition or contestation therefore involves material politics. This entire paper is a study of consolidation being materially contested.

The conflicting players in this material politics vary greatly. Sometimes they are different teams or departments within the same corporation. How long the longest second should be is an example. One technical specialist tells us that he has experienced 'a lot of political battles' over that, between teams tasked with maximizing advertising revenues (who are prepared to have multiple AdTech systems run during page loading, even if they slow it down) and technical staff whose priority is 'page speed', which, e.g. increases the site's search-engine rankings. But material politics can also interact with more traditional forms of politics. For example, the legal politics that increasingly accompanies struggles over advertising's materiality has a party-political aspect, as in the fracturing along largely partisan lines of anti-Google litigation. In particular, under the Biden administration, the standard conceptual framework of US competition law is being challenged as too narrow, with prominent critics appointed to key roles: e.g. Lina Khan to chair the Federal Trade Commission, Jonathan Kanter to

head the Justice Department's unit that has now launched its lawsuit against Google, and Tim Wu as Biden's advisor on technology and competition policy.

What of our other argument, that platforms are *stacked* processes of economization? That helps us zoom out from the often-intricate details of material politics. One way of doing that is to compare online advertising with a superficially similar sphere, automated trading in finance. Buying and selling there is also by algorithms, and when electronic ad exchanges such as AdX were first introduced they were often explained as analogous to electronic financial markets. But there the similarities largely end, and the differences between the AdTech and financial 'stacks' explain why.

In finance, connections between ultrafast automated trading and economization processes involving corporations or governments as human institutions have been severed in time. Financial trading's timescale is now the nanosecond: a billionth of a second. The probability of a government or corporation, its voters, workers or customers, doing something of financial significance in any *particular* nanosecond is extremely low, and today's high-frequency trading or HFT is driven by quite different 'signals' (see MacKenzie, 2021). That 'severing' of the stack, the disjunction in timescale between HFT and the everyday human world, is indeed the source of much of the unease that has often surrounded HFT. Nanosecond speeds, in their turn, also dictate HFT's huge sensitivity to precise spatial location: in a nanosecond, even the fastest signal, light in a vacuum, can travel only around 30 centimetres. Location matters in AdTech too, but to nothing like the same extent.

One reason why we insist that what is stacked by digital advertising's platforms is economization processes, not just technologies, is that the AdTech stack cannot be severed in time in the way the HFT stack has been. An ad that is neither seen nor heard by a human being is not an advertisement, unless humans start fully to delegate their purchases to automated systems that can be influenced by ads. This non-severing shapes digital advertising profoundly. First, advertising opportunities are intrinsically heterogeneous, because the 'data doubles' (Haggerty & Ericson, 2000) of human users, employed in targeting, have sufficiently many dimensions that only rarely are any two exactly the same. Second, human beings' non-severable real-time involvement means advertising opportunities are perishable in a way that most financial instruments are not. An advertising opportunity 'exists for a moment in time', says interviewee CM: 'if it doesn't get filled it goes away'. An ad that has not been shown by the time the longest second ends may never be seen: the user may have left the page or switched to another app.

Third, though, and crucially in terms of how digital advertising differs from financial trading, delays in advertising are ultimately measured in human time, not the machine time of HFT. There is a lower threshold to human beings' capacity to perceive time passing, signalled by the title of Canales's *A tenth of a second: A history* (2009). Tenths of a second matter in digital advertising, and some early staff members of Google developed the near-preternatural capacity directly to sense delays that accurately (Levy, 2011, p. 185). But

time differences of much less than a tenth of a second are consequential in advertising only if they accumulate to humanly-perceptible delays.

Those differences between the ‘stacks’ of automated financial trading and digital advertising translate into quite different material politics. As discussed in MacKenzie (2021), much of the material politics of HFT concerns sources of advantage in nanosecond speed races, while the material politics of digital advertising hinges on quite different issues. This paper has begun to explore some of these issues for the case of the auctioning of ad opportunities, but much more remains to be uncovered in digital advertising’s many other aspects.

Notes

1 An ad impression is the display of one ad to one user. The figure of 400 billion is derived from the plausible, even if imprecise, estimate of 146 trillion annually in Kotila (2021).

2 The classification is porous in a variety of ways: e.g., display ads often have video content.

3 In online advertising, ‘publishers’ are the owners/originators of online content of all kinds, including e.g. firms that produce computer games and organizations as diverse as dictionary.com or the UK’s Meteorological Office. This paper focuses on web advertising via both mobile devices and laptops/desktops; as briefly noted in our conclusion, header bidding is less prevalent elsewhere.

4 DFP is, in the terminology of the field, a *publishers’* ad server. Advertisers’ ad servers are also important to an integrated full-stack AdTech offering such as Google’s, but word-count constraints prevent us discussing them.

5 DFP and AdX are now component parts of Google Ad Manager, but for simplicity we use those names throughout.

6 There is also, e.g., a helpful examination of pricing mechanisms in header bidding in Despotakis *et al.* (2021).

7 E.g., van der Vlist *et al.* (2022) focus on a crucial ‘stacking’ mechanism: Facebook’s application programming interfaces.

8 We owe this reference to Crain (2021, p. 141).

9 For clarity, we ignore, e.g., how the process has changed through time (Competition and Markets Authority, 2020, pp. M8–M12), and whether more than one bid can be forwarded at each stage.

10 See <https://www.youtube.com/watch?v=suszUQMjz3I> (accessed 30 March 2022).

11 Competition and Markets Authority (2020, pp. M109–M110) describes why this is difficult.

Acknowledgements

We are hugely grateful to our interviewees, without whom this paper could not have been written. Errors, however, remain our own.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the UK Economic and Social Research Council [grant number ES/V015362/1].

Ethical approval statement

This research was assessed, classed as ‘Level 1’ (‘negligible or low foreseeable risks’), and ethically approved under the research ethics procedure of Edinburgh University’s School of Social and Political Science: <https://www.sps.ed.ac.uk/research/ethics>. Informed consent was obtained for all interviews.

ORCID

Donald MacKenzie  <http://orcid.org/0000-0002-1860-8116>

Koray Caliskan  <http://orcid.org/0000-0001-9876-9405>

Charlotte Rommerskirchen  <http://orcid.org/0000-0001-7054-5629>

References

- Alaimo, C. & Kallinikos, J. (2018). Objects, metrics and practices: An inquiry into the programmatic advertising ecosystem. In U. Schultze, M. Aanestad, M. Mähring, C. Østerlund & K. Reimer (Eds.), *Living with monsters? Social implications of algorithmic phenomena, hybrid agency, and the performativity of technology (IS&O 2018)* (pp. 110–123). Springer.
- Amoore, L. (2020). *Cloud ethics: Algorithms and the attributes of ourselves and others*. Duke University Press.
- Barry, A. (2002). The anti-political economy. *Economy and Society*, 31(2), 268–284.
- Barry, A. (2013). *Material politics: Disputes along the pipeline*. Wiley.
- Beauvisage, T. & Mellet, K. (2020). Datasets: Assetizing and marketizing personal data. In K. Birch & F. Muniesa (Eds.), *Assetization: Turning things into assets in technoscientific capitalism* (pp. 75–95). MIT Press.
- Beuscart, J.-S. & Mellet, K. (2013). Competing quality conventions in the French online display advertising market. *Journal of Cultural Economy*, 6(4), 402–418.
- Birch, K. & Muniesa, F. (Eds.). (2020). *Assetization: Turning things into assets in technoscientific capitalism*. MIT Press.
- Brockhoff, J., Jehanno, B., Pozzato, V., Buhr, C.-C., Eberl, P. & Papandropoulos, P. (2008). Google/DoubleClick: The first test for the Commission’s non-horizontal merger guidelines. Retrieved from https://ec.europa.eu/competition/publications/cpn/2008_2_53.pdf
- Burrell, J. & Fourcade, M. (2021). The society of algorithms. *Annual Review of Sociology*, 47, 213–237.

- Caliskan, K. (2020).** Platform works as stack economization: Cryptocurrency markets and exchanges in perspective. *Sociologica*, 14(3), 115–142.
- Caliskan, K. & Callon, M. (2009).** Economization, part 1: Shifting attention from the economy towards processes of economization. *Economy and Society*, 38(3), 369–398.
- Canales, J. (2009).** *A tenth of a second: A history*. University of Chicago Press.
- Cluley, R. (2018).** The construction of marketing measures: The case of view-ability. *Marketing Theory*, 18(3), 287–305.
- Cluley, R. (2020).** The politics of consumer data. *Marketing Theory*, 20(1), 45–63.
- Cluley, R. & Brown, S. D. (2015).** The dualised consumer: Sketching the new mask of the consumer. *Journal of Marketing Management*, 31(1–2), 107–122.
- Competition and Markets Authority. (2020, July 1).** Online platforms and digital advertising: Market study final report. UK Competition and Markets Authority. Retrieved from <https://www.gov.uk/cma-cases/online-platforms-and-digital-advertising-market-study>
- Cotter, K. (2019).** Playing the visibility game: How digital influencers and algorithms negotiate influence on Instagram. *New Media & Society*, 21(4), 895–913.
- Crain, M. (2021).** *Profit over privacy: How surveillance advertising conquered the internet*. University of Minnesota Press.
- Despotakis, S., Ravi, R. & Sayedi, A. (2021).** First-price auctions in online display advertising. *Journal of Marketing Research*, 58(5), 888–907.
- Dourish, P. (2017).** *The stuff of bits: An essay on the materialities of information*. MIT Press.
- Eyal, G. (2013).** For a sociology of expertise: The social origins of the autism epidemic. *American Journal of Sociology*, 118(4), 863–907.
- Fourcade, M. & Klutzz, D. N. (2020).** A Maussian bargain: Accumulation by gift in the digital economy. *Big Data & Society*, 7(1), 1–16.
- Gabrys, J., Hawkins, G. & Michael, M. (Eds.). (2013).** *Accumulation: The material politics of plastic*. Routledge.
- Goldfarb, A. & Tucker, C. (2019).** Digital economics. *Journal of Economic Literature*, 57(1), 3–43.
- Goldman, E. (2006).** A Coasean analysis of marketing. *Wisconsin Law Review*, 2006(4), 1151–1221.
- Google LLC. (2020).** Online platforms and digital advertising: Comments on the Market Study Interim Report. Retrieved from https://assets.publishing.service.gov.uk/media/5e8c8290d3bf7f1fb7b91c2c/200212_Google_response_to_interim_report.pdf
- Haggerty, K. D. & Ericson, R. V. (2000).** The surveillant assemblage. *British Journal of Sociology*, 51(4), 605–622.
- IAB. (2022).** Internet advertising revenue report: Full-year 2021 results. Retrieved from https://www.iab.com/wp-content/uploads/2022/04/IAB_Internet_Advertising_Revenue_Report_Full_Year_2021.pdf
- Kotila, M. (2021).** How much does an ad impression emit carbon? Retrieved from <https://medium.com/art-technology-how-much-does-an-ad-impression-emit-carbon-d31daea57184>
- Langley, P. & Leyshon, A. (2017).** Platform capitalism: The intermediation and capitalisation of digital economic circulation. *Finance and Society*, 3(1), 11–31.
- Langley, P. & Leyshon, A. (2021).** The platform political economy of fintech: Reintermediation, consolidation and capitalisation. *New Political Economy*, 26(3), 376–388.
- Law, J. & Mol, A. (2008).** Globalisation in practice: On the politics of boiling pigswill. *Geoforum; Journal of Physical, Human, and Regional Geosciences*, 39(1), 133–143.
- Lebow, S. (2023).** Worldwide digital ad spend will top \$600 billion this year. Retrieved from <https://www.insiderintelligence.com/content/worldwide-digital-ad-spend-will-top-600-billion-this-year>
- Levy, S. (2011).** *In the plex: How Google thinks, works, and shapes our lives*. Simon & Schuster.
- Mackenzie, A. (2006).** *Cutting code: Software and sociality*. Peter Lang.

- MacKenzie, D. (2021).** *Trading at the speed of light: How ultrafast algorithms are transforming financial markets*. Princeton University Press.
- MacKenzie, D. (2023).** Short cuts: A puff of carbon dioxide. *London Review of Books*, 45(2), 16.
- McDermott, J. (2014).** Google's display advertising dominance raises concerns. Retrieved from <https://digiday.com/media/google-bundling-ad-tech-inventory-raising-anti-competitive-concerns/>
- Mellet, K. & Beauvisage, T. (2020).** Cookie monsters: Anatomy of a digital market infrastructure. *Consumption Markets & Culture*, 23(2), 110–129.
- Nolet, M. (2010).** The challenge of scaling an adserver. Retrieved from <http://www.mikeonads.com/2010/04/04/the-challenge-of-scaling-an-adserver/>
- O'Meara, V. (2019).** Weapons of the chic: Instagram influencer engagement pods as practices of resistance to Instagram platform labor. *Social Media + Society*, 5(4), 1–11.
- Page, L. & Brin, S. (2006).** 2006 founder's letter. Retrieved from <https://abc.xyz/investor/founders-letters/2006/>
- Sluis, S. (2022).** Google ad manager builds a bridge to Prebid – but don't call it a two-way street. Retrieved from <https://www.adexchanger.com/platforms/google-ad-manager-builds-a-bridge-to-prebid-but-dont-call-it-a-two-way-street/>
- Srinivasan, D. (2020).** Why Google dominates advertising markets: Competition policy should lean on the principles of financial market regulation. *Stanford Technology Law Review*, 24(1), 55–175.
- Stark, D. & Pais, I. (2020).** Algorithmic management in the platform economy. *Sociologica*, 14(3), 47–72.
- Texas, Alaska, Arkansas, Florida, Idaho, Indiana, Kentucky, Louisiana, Mississippi, Missouri, Montana, Nevada, North Dakota, Puerto Rico, South Carolina, South Dakota & Utah. (2022).** In re: Google digital advertising antitrust litigation. Third amended complaint. Retrieved from https://texasattorneygeneral.gov/sites/default/files/images/child-support/20220114_195_0_States%20Third%20Amended%20Complaint.pdf
- Viljoen, S., Goldenfein, J. & McGuigan, L. (2021).** Design choices: Mechanism design and platform capitalism. *Big Data & Society*, 8(2), 1–13.
- van der Vlist, F. N., Helmond, A., Burkhardt, M. & Seitz, T. (2022).** API governance: The case of Facebook's evolution. *Social Media + Society*, 8(2), 1–24.
- Ziewitz, M. (2019).** Rethinking gaming: The ethical work of optimization in web search engines. *Social Studies of Science*, 49(5), 707–731.
- Zuboff, S. (2019).** *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile.

Donald MacKenzie teaches sociology at Edinburgh University. His books include *Trading at the speed of light: How ultrafast algorithms are transforming financial markets* (Princeton University Press, 2021).

Koray Caliskan teaches strategic design and management at the New School's Parsons School of Design. His books include *Data money: Inside cryptocurrencies, their communities, markets and blockchains* (Columbia University Press, 2023).

Charlotte Rommerskirchen teaches international political economy at Edinburgh University. Her monograph *Free riders on the storm: EU fiscal policy coordination in times of crises* was published by Oxford University Press in 2019.