

Dear members of the inquiry panel,

Please accept this submission and consider the issues discussed and potential solutions it offers.

By way of introduction, I am employed as a university lecturer who teaches subjects in computer software development. I have once used both Google Ads and Facebook Ad Manager to advertise for a non-profit organisation on a small (\$100) budget. At the time (in 2018), I found the Google offered better value for money – Facebook actually had less precise targeting options for the particular audience trying to be reached (a cultural/arts community), but Google offered a much broader range of attributes against which I could select to target the ad.

I have two key themes that are addressed in the submission. The first is to point out an area of investigation that I think has not been adequately considered by the ACCC in the interim report, which is the role played by the publishers of websites (and apps) in deciding *whether* to provide (and if so, *how* they will provide) advertising space. The second is to propose some ways that might bring greater transparency and auditability to the whole system through mandating information about transactions involving ad technology being published on a publicly viewable blockchain.

I would also like to point out that the interim report in its analysis of digital advertising seems to have omitted a category of ‘audio only’ advertising, as may arise when people listen to Spotify and similar services (that don’t involve viewing of a screen). For information about their advertiser’s platform, refer to: <https://ads.spotify.com/en-AU/> and related web pages.

These first set of comments are primarily regarding “display advertising” as defined on page 30 of the interim report, which is what is identified as the key focus of the report.

Regarding the comment on page 14, that “High switching costs combined with the tendency for publishers to single-home may constitute a significant barrier to entry and expansion”: I don’t think cost would be a factor for a developer of a website. Single-homing might have more weight. Ultimately, the developer/programmer of a website or an app is the one who decides which “publisher ad servers” technologies to utilise, and this is often just lazily chosen by going with the one that is most well-known amongst software developers, i.e. Google’s.

The only way that you could alter the “share” of “market” that Google or others have, would be to convince a large number of software and website developers to choose to use a different publisher ad servers. **You would have to investigate what the motives are of individual software and website developers (programmers) for choosing particular publisher ad servers** (such as Google Ad Manager), as it is their decisions which are most directly forming the market’s structure. It seems wrong to criticise Google for having a large “market share” for this aspect of the chain. But perhaps actions of Google, in terms of payments that they offer to make to the publishers/developers of sites and apps, compared against payments offered by other publisher ad servers, make the Google offer more attractive to be taken up, and might be the key impediment to increase “competition” (which seems to be the motive of the ACCC).

That attraction to or favouring of Google as the publisher ad server to be used by the website and app developers leads to Google being able to promote itself to advertisers that it has the capacity over other publisher ad servers, which is the more likely cause of advertisers choosing to use Google to submit their ads. If Google never had encouraged a large network of willing space-makers to adopt their technology, Google would never have attracted advertisers. But at some point over the years the snowball effect must have kicked in to lead to its present day dominance on both sides of the advertising/presentation equation, driven by their superior ability to place content in highly targeted ways.

Secondly, it is the behaviour of individual web-users or app-users (the general public) which significantly contributes to which websites/webpages and which apps are actually used, and thus which publisher’s ad servers will be contacted for ads to serve on their site. Unless you could convince these end-users to switch

to different websites or apps that happen to use other publisher ad servers, not much could be done to spread “market share” to other publisher ad servers.

I feel that the only thing that made Google superior / preferred was because they were the first to succeed in exploiting all the data and analytics which they collect about individual people through so many channels and analyse through various machine-learning techniques. The ‘profile’ built of a person to whom ads will be targeted often comes from analysing data collected over long periods of time (months, years), combining with the more contextual moment-based data. Anyone trying to do the same starting from a clean slate now, is already a long-way behind and catch-up could take a long time. Additionally, building a network of supporting developers / publishers to create ad-space using their new publisher ad server platform would be near impossible – they would have to have a much better value proposition to the website owner/developer than what Google already offers them, which will only come from having advertisers flocking to them, which would also require a value proposition to advertisers that was better than Google’s. It may even be unwise to encourage further companies to develop such highly personal profiles of individuals as Google and a few others (e.g. Facebook) have.

Perhaps the only other significant/equivalent competitor to Google might be Baidu in the China market (and maybe other companies in other non-English markets, though I don’t know any, and I don’t have data about Baidu either).

Under these environmental settings, one could almost argue that it is pointless to advertise with anyone *other* than Google if you wanted greatest chance of exposure. Perhaps those who advertise with others are doing so to exploit niche opportunities due to others not bothering to use them. It would be surprising if any advertiser wholly avoided Google in favour of others.

Further, the websites and apps being visited may be overseas, and been created by developers who had no thought about whether their users/visitors would be in Australia or elsewhere – the technology to present the ads is designed to work based on where the user/viewer happens to be at time of presentation, which is probably one of the reasons Google’s technology is preferred by the developers and why it has found world-wide adoption.

Reading through the interim report Executive Summary chapter, it almost seems to suggest that the ACCC has approached the situation backwards to how I feel things actually are designed and appear to function¹. If there were no publishers (websites or apps) providing space for the ads to be displayed, there would be no market for advertisers to sell their ads. That is presumably why websites hosting ad-space are paid. It is not quite as though advertisers have an unfettered right to have some guaranteed number of impressions of their ads, since impressions is entirely dependent upon end-users actually viewing a website or using an app which has advertising included in it. The website owners / developers are probably attracted to the publisher ad service which they consider is most likely to give them the best deal for income. Though it does seem odd to me that publishers are charged fees by publisher ad servers, for making available space to advertisers.

To further complicate the matter of availability to display an ad, there are plug-ins people often install on their web browser to block ads from being fetched from the publisher ad servers. So even if someone visits a website, where the creator of that website has provisioned space for advertisements to appear, the ads will not actually be collected or presented to the user having that plug-in (the browser will never send the HTTP request to fetch the content), thus a bid will not take place, (and no advertiser will have had to commit a portion of their budget, and the website owner won’t earn anything from it). An example of such a plug-in is uBlock Origin (see <https://addons.mozilla.org/en-US/firefox/addon/ublock-origin/> for the Firefox browser). Obviously, use of such a tool would frustrate all steps in the chain from the website owner through to the advertiser who potentially loses out on chance to even be considered for an ad to be presented. As more

¹ Although I have significant skill in software development, I have not yet ever actually coded the provision of advertising space on a website or in an app, but can interpret the various documentation explanations on how to do so. For example: <https://developers.google.com/publisher-tag/guides/get-started> for Google Ad Manager.

people become aware of such plug-ins and install them, the volume of ad presentation opportunities overall will decrease.

There are a great many websites hosting information of dubious quality, and sites which draw small audiences for information of various quality levels, (in addition to high-value and high-volume sites), where the owner of the website may use an ad tech on the site to gain some income and they are not really concerned with whether they are getting the best price they could (because *anything* is worth more than *nothing* in their view). Similarly, many apps which choose to include ads, are doing so merely to get *any* income they can. The number of such apps and sites, in comparison to ones which we might call reputable and truly reliant on wanting to achieve highest income from advertising presented on their websites, may be imbalanced (towards the dubious ones making up a greater proportion of the overall number of “opportunities” that may be served in a given hour by any particular publisher ad server). I am not aware of whether the advertiser has any option to exclude their ads from being considered for presentation by those less reputable sites – but even if it were a possibility, I am unsure that advertisers would care, since they usually just want their ads to be seen. In other words, I think there are much fewer website and app owners/developers than advertisers, who would be caring about the choices they respectively are making to use particular technologies.

Setting aside those issues, which are harder to control or change, I now turn to other matters that are more related to the specific matters on which comment has been sought and which may address in some degree proposals 4 to 6.

I would certainly commend any effort to make more information available to parties (advertisers, publishers) about the operation of the ad tech services. This is possibly a place where legislation would have more chance of succeeding to impact change (where you could, for example, require certain acts/processes to be done, in order for an ad technology to be compliant with a licence granted for its operation in regards to ad impressions served to Australian devices). I would suggest that blockchain-like information mechanisms would be the most useful way to achieve it.

Although initially developed to support Bitcoin, the idea of a blockchain is that it is a publicly viewable record of transactions over time, which through linked and uniquely generated identifiers can show the progress of tokens or other pieces of information over time. It provides information visibility for audit. The blockchains of the numerous ‘crypto coins’ usually provide various blockchain *explorers* and other ways to extract the data from the blockchain for analysis about the tokens or events that are recorded in the blockchain. For example, one of the explorers for Bitcoin is provided by BlockChair at <https://blockchair.com/bitcoin/blocks> - clicking on a block reveals its data including the transactions that were included in that block. Clicking on a transaction reveals the sender(s) and receiver(s) and the amounts transferred and other data included in the transaction. Clicking on a sender shows all the transactions to or from that sender over time, and thus the evolution of their balance of Bitcoin amounts held. It is possibly these explorers, rather than the blockchain itself, in the way in which they **provide traceability and thus audit capability**, which has inspired the following description of how a blockchain information economy regarding ad tech transactions, might be applicable in the market of digital advertising to solve some of the issues raised in the interim report²:

The primary focus of interest is a request for an ad impression (on a device showing a webpage or in an app), since this is the goal of an advertiser³, and of a provider of advertising space⁴. Each individual such request would be the basis of an entry included on the blockchain. I will call this an *impression entry*.

² This is by no means a final or thorough architecture of a system, just an overview of what could be possible. The details would need much more thought.

³ The advertiser wants to know that their ad was presented, and they care about how much it cost them.

⁴ The provider of advertising space probably has some interests too.

Each such impression entry is given some unique identifier (a UUID, possibly time-bound so they can be recycled e.g. daily recycling, monthly recycling.) Other information entities included in the entry (described below) would also be given unique identifiers. As to who “knows” the meaning of identifiers, is also described below.

The impression entry on the blockchain regarding that impression would also need to include an identifier of which ad was actually served, so that it could be reconciled back to information the advertiser can see. If the same ad was presented across various impressions (to different people for example), each impression could use a unique identifier, or they could share the identifier, depending on the need for clarity or privacy.

The impression entry should include an identifier of a bid event. This would be a cross reference to a separate entry on the blockchain which would describe such things as who the candidate bidders were for that bidding event, which I will call a *bid event entry*.

The impression entry should disclose the amount paid to various parties such as the publisher (website owner, app owner), and other parties involved in that impression having arisen.

The impression entry might somehow indicate whether the recipient who saw the ad was considered to be real/trusted, as opposed to a fraudulent impression by a robot program.

The impression entry should also include identifiers to any other information which the architecture of a blockchain for the purpose might want to include (this is up to ACCC to decide what is needed beyond what I might outline here). These various identifiers are what would enable examination of the blockchain and the transactions/events it records, but preserving to some degree some anonymity where needed.

In the case of an audio or video advertisement, the impression entry might also need to state for how long the advertisement was played / whether it was played in full, if it is possible to record such information.

The *bid event entry* might include such things as cross references to bidder entries (outlined later) of the bidders who participated, scope definition entries (e.g. do we want an ad for children, do we want an ad for a Sydney citizen), information about target price ranges for bids (does it need to be at least 50 cents worth, or at most \$5 worth), perhaps whether any particular bidders were to be excluded, etc.

There might even be a separate *presentation opportunity entry* which might represent information about the opportunity to bid, as opposed to the actually bidding done for that opportunity.

A *bidder entry* would provide some sort of information regarding bidders participating in a bid event. It might describe/indicate who the participant is, in some way that preserves the privacy of the individual but perhaps reveals an ‘agent’ that represents them in the bid. (By ‘agent’ I mean some sort of intermediate step or service which generates that one-off id to correspond to an entity/party, and which communicates that to the party as a secret). Maybe the same actual bidder has a unique or otherwise anonymised identifier for each individual bid event they participate in. It might include the range of prices they were willing to bid to be considered for that bid event.

There could be entries for other elements, other events, etc., however the ACCC saw fit.

The information pertaining to a set of entries of these various types, need to be gathered together into ‘blocks’ that are added to the blockchain (to form the chain sequence), at which point they become visible and known, and software can read the blockchain to analyse the events and entries to give information to advertisers. So they might be able to see for instance how many bids their ads were involved in where they actually lost or actually won; publishers could see how much was paid for the ad and compare it to how much they were paid to present it on screen. It would bring greater transparency than the current “opaque” mechanisms mentioned in the ACCC interim report. The architecture developed by ACCC/government would need to set rules such as how soon the information had to be placed into the blockchain after the event is finalised (e.g. a few seconds).

To enforce participation, perhaps a licencing scheme could be introduced. Providers along the Ad tech chain might need to obtain a license from the ACCC to operate in any context where an Australian-located device was to be involved in requesting an ad for presentation on its screen. Or perhaps licenses only apply to ‘declared regulated’ providers. But the licence could have conditions that require the licensee to place information into the blockchain, and to provide certain information such as identifiers to the advertisers (so the advertisers can verify/reconcile transactions by examining the blockchain themselves or via third-party providers of analytic services). There could then be penalties for failing to meet the conditions of the licence.

In such a case, I would recommend that the blockchain be one operated by the government (e.g. via the ACCC), as a trusted independent party. Thus the node(s) maintaining the data would be all owned by the government⁵ (whereas the blockchain for bitcoin has no precise owner). You might even charge a fee for each entry that a licensee places into the blockchain (to support the government’s costs of running the blockchain), or the fee may be based on volume of entries place into the blockchain so that lesser participants in the ecosystem don’t incur fees. Alternatively, if multiple countries agreed to the same mechanism, perhaps the governments could jointly operate a single blockchain shared over their jurisdictions.

The architecture just outlined only addresses a way of recording what the outcomes are. An alternative architecture might include additional events such as an advertiser making available an advertisement (irrespective of whether it is ever selected for presentation) or placing of amounts of money ‘on the table’ to be used in bidding. I’m sure a group of experts could come up with a design and business rules that addressed all the relevant matters appropriately, the above is just a sketch.

Thankyou for considering this submission.

⁵ A so-called private or in-house blockchain, though it may be made publicly readable in this application of the concept. The government-operated nodes would create the blocks, there is no need to “mine” for rewards to external parties.