This is a draft of an article published in the *London Review of Books*, vol. 46, issue 16 (15 August 2024), pp. 25-27: https://www.lrb.co.uk/the-paper/v46/n16/donald-mackenzie/hey-big-spender
It was edited extensively by the *LRB*, but mainly stylistically. There was one substantive update: in July 2024, Google announced that instead of ending the default status of third-party cookies on Chrome, it would offer users the opportunity to accept or reject these cookies. The fact-checking text at the end of this draft was of course not published, but is kept here as indicating some of the sources that were used.

Hey Big Spender: What Your Smartphone Knows About You

Donald MacKenzie

Playing Candy Crush Saga involves moving brightly coloured sweets around to the sound of cheerful music. Get three or more identical sweets into a line, and they disappear, accompanied by gentle explosions. Your score ticks up, and a cascade of further sweets refills your mobile-phone screen. If all goes well, you'll soon complete a level. A warm, disembodied, male voice offers encouragement: Divine! Sweet!

The iPhone version of Candy Crush was released in November 2012, and an Android version a month later. The following December, BBC News reported train carriages in London, New York and other big cities full of commuters 'fixated on one thing only. Getting rows of red jelly beans or orange lozenges to disappear'. You've never had to pay to install Candy Crush, and it has been downloaded more than five billion times in total, which suggests that hundreds of millions of people must have played it. Over 200 million still do, according to the game's makers, the Anglo-Swedish games studio, King. Those players aren't going to exhaust its challenges any time soon: Candy Crush has over 15,000 levels, and dozens more are added weekly.

In 2015, the media scholar David Nieborg pointed out that Candy Crush was already big business. By 2023, it had earned over $20 billion in total for King and Activision Blizzard, the games conglomerate that bought King in 2016 for $5.9 billion. Activision Blizzard has now itself been bought by Microsoft for $69 billion, a consolidation of the games sector that worried the UK's competition regulator, the Competition and Markets Authority, enough that it came close to trying to block it.

How do you make money out of a game – not necessarily big money, but at least enough to repay its often high development costs – if it's free, as nearly all mobile-phone games are? Candy Crush is more difficult than it looks, and you can end up temporarily out of 'lives'. It's then tempting to spend a modest sum to keep playing or boost your future chances. In my second session, I succumbed. A half-price weekly deal was on offer for 99 pence. I tapped, was taken to Apple's App Store, which has my credit card on file, and a thumbprint sealed the deal.

Most players of games such as Candy Crush are less easily tempted to spend than I was. The prominent app-economy analyst Eric Seufert tells me that typically ’95 percent, 97 percent of all users who play a game will never monetise’. If your game is popular enough, that might still mean tens or hundreds of thousands of players spending within it. Attracting them is therefore a vitally important part of what people in the business unromantically call ‘user acquisition’.

Many of them, like me, spend only small sums, and only once in a while. The most valuable players are the bigger spenders referred to as ‘whales’. Because the basic goal in playing Candy Crush is to get at least three digital objects next to each other in a line, it's known in the business as a ‘match-three’ game. There are hundreds, perhaps thousands, of such games. A ‘match-three whale’ – that's the term that’s used – is someone who spends tens of dollars monthly in one or more of them. Find a decent number of such whales, and you’ve got a valuable income stream.

Word-of-mouth recommendations, charts of ratings and total downloads, favourable reviews by games journalists, endorsements by prominent influencers, and general social-media buzz can all bring players to a game. But building a mass user base frequently requires large-scale advertising, often on a social media platform such as TikTok, Snapchat, Facebook or Instagram.

The most immediately obvious place, though, in which to advertise your game is in another game. It's not simply that it’s there that whales are to be found. I’m a late convert to the pleasures of mobile-phone games. Until the last few months, someone advertising a game to me, however attractively, would simply have been wasting their money. But if I'm already playing a game on my phone, then I’m an a priori plausible target. More specifically, if I’m playing a match-three game, then why not advertise your match-three game to me? It’s not going to cost me anything to try it out, and I might find I prefer its slightly different format, images, colours or soundtrack. Handily, there’s a well-established way of making me want to watch a video ad within whatever game I’m currently playing, rather than ignore it or be annoyed by it: giving me an in-game boost as a reward for watching the ad.

The lure of trying to acquire new players by advertising within other games is strong. Which ads get to be shown often comes down simply to how much the advertiser is prepared to pay, and, as a games executive put it to me, the owners of other games ‘are willing to pay a lot’ – more than an advertiser for a different kind of product would – ‘because that’s where they're going to drive their installs’. As a result, Seufert reports, often ‘the overwhelming majority, 95 percent, of all ads shown in games, are for other games’. Sometimes, an ad for a game is itself a game, a brief sample of the real thing, although that's an expensive form of advertising, and I’m told it has become less popular recently.

‘I’m buying users from you, you're buying users from me, a lot of revenue was materialised but actually it all got sort of negated by the fact that we’re just buying from each other’, says Seufert. The games executive quoted above tells me that ‘fundamental tensions’ surround earning money by showing ads for other games: they ‘can be competitors, which isn’t awesome’, and can cause ‘my players to churn out’. Most ads in mobile-phone games are sold via automated bidding systems, rather than face-to-face negotiation, so it is not always

straightforward to block specific unwanted ads. If a competitor ‘is determined enough, they can get an ad in your game, no question. These systems are not hard to game.’

It can, however, be hard for a games studio to say no to the revenue stream that advertising provides. Indeed, there’s an entire genre of relatively rudimentary ‘hypercasual’ games, which I’m told are in effect entirely funded by showing ads for other games. As another experienced practitioner puts it, ‘the justification that the companies make and use is, I’m going to make some ad revenue and then I’m going to spend the ad revenue to create a marketing budget [to acquire users for my game] …You’re hoping you’re giving away your less valuable players, but I'm not sure it’s really that scientific.’

_____

Following the mechanics of the advertising of games, apps and other products on mobile phones leads us deep into the tensions of the digital economy. Some firms advertise simply to keep consumers’ views of their brand well-burnished: a motor manufacturer wants to interest you in its cars, but doesn’t really expect you instantly to buy one. But much, perhaps most, advertising on mobile phones is what’s called ‘direct response’. That has an immediate goal, a ‘conversion’: a purchase, a sign-up or subscription, an install of a game.

A games studio, though, can’t ignore what happens once someone installs a mobile-phone game. Its chief concern is likely to be the ‘lifetime value’ of players (the total revenue that they will bring to the game), and whether that will exceed the cost of the advertising that led them to it. There’s a series of ‘app events’, as they are called, that it will want to monitor closely. How many people install the game but abandon it before finishing the ‘tutorial’ with which it begins? How many keep playing to at least, let’s say, level 5? Above all, how many make in-game purchases, and if so when, how often, and for how much?

But direct-response advertisers such as games studios can’t afford to wait too long for lifetime value to become evident: they want to maximise the efficacy of advertising in close to real time. As a digital-advertising specialist puts it, ‘you’ve got these bidders [for advertising opportunities] with machine learning that are saying this segment is working, bid higher here because there are conversions occurring. There’s all this automated feedback loops that are running.’

Machine-learning optimisation is on offer throughout digital advertising, but its doyennes are Google and Facebook, renamed Meta in October 2021. When I began to develop a serious interest in the advertising of apps and games, I took an online course on it taught by Seufert. He carefully explained its technicalities, such as the complications of determining users’ lifetime values, and implicitly warned against the temptation for a games studio to spend all its money advertising in other games.

An entire session of the course was devoted to advertising via Meta’s platforms, Facebook and Instagram. Meta enables you to specify with surprising precision what you want its advertising systems to optimise on your behalf. The obvious choice would be simply the number of installs of your game or other app. In 2016, however, Facebook introduced ‘App Event Optimization’, which allows fine-tuning by focusing advertising so that it generates

installs that its systems predict will be followed by player actions of the kind that you want to prioritise, such as in-app purchases. You can also choose the goal of maximising the total revenue that players will bring you. That's what Meta calls 'Value Optimization'. It can be expensive, but you would choose it if you are hunting whales.

As in most of digital advertising, auctions, in this case internal to Meta, determine which ads get shown to which users, but you don't yourself have to take on the daunting task of working out how much to bid. Specify your goal, your budget, and perhaps the minimum 'return on ad spend' that would be acceptable to you, and Meta bids into its own auctions on your behalf. You can manually set a maximum bid, but the Competition and Markets Authority reported in 2020 that only a small minority of Facebook advertisers were doing that: the others relied upon Facebook's systems to formulate bids that would achieve their specified goal cost-effectively.

A striking moment in Seufert's course was when he described how advertising on Facebook had changed. He talked about meetings, in around 2015, of a computer-game user acquisition team, in which he and his colleagues would discuss in detail how to target their Facebook advertising. 'Maybe we should try car enthusiasts, because … we're trying to reach men … maybe we should target people that "like" Bruce Willis's page because he's like an action star … That was what you did.' Facebook's increasingly sophisticated machine learning has, however, made such discussions a waste of time, he said. 'Now, Facebook has basically internalised all that and now you're just feeding it with … inputs [e.g., multiple variants of your ads so that it can test which are most effective]. None of that stuff you used to do matters.' Now, the job of the advertising practitioner is 'to feed this experimentation machine'.

A major digital-advertising platform, such as Facebook or Google, is like an iceberg, Seufert suggested. Visible above the waterline are the characteristics of users that advertisers can employ to target their advertising, such as age, gender, and 'interests' like enthusiasm for motorcars. But below the surface, invisible to the advertiser and too copious to make full sense to human beings, is the much larger volume of other data, far more heterogeneous in nature, that the platform possesses. That's the data that can make platforms' machine-learning optimisation of advertising considerably more effective than human-guided targeting.

The iceberg's implications go well beyond the advertising of games. One of the most interesting of the nearly 90 people I've interviewed so far about digital advertising has a business selling handmade saris, and a strong commitment to preserving village handicrafts. In our first conversations, he was highly critical of Big Tech. But he has gradually learned that his best market is Tamil Brahmins, in India and in the diaspora. There's no list of them for him to work from, and on Facebook or Google 'there's no classification saying "target Tamil Brahmins".' Yet using data on 'customers who've bought … from me before, who've interacted with my products, Google and Facebook are able to find them. … My clients look for 30-minute recipes. They look for Bollywood news. They look for Tamil cinema news. … I have to trust the machine to be more effective than me to do this.'

It's a useful reminder that, in his words, 'machine learning, although it's imperfect, has the potential to amplify the reach of small businesses.' But the machine is, of course, amoral. It optimises for whatever 'conversions' it is told to pursue: installs of Angry Birds; sales of saris made by village women; voters signing up for a Donald Trump rally.

People whose political predilections resemble mine often like to think that the explanation for Trump's 2016 election or Brexit is cunning microtargeting by political consultants using platforms such as Facebook, perhaps funded by Russian money or informed by Cambridge Analytica's psychometric data. An April 2020 *Atlantic* article by Ian Bogost and Alexis Madrigal has a more convincing hypothesis in respect to Facebook: that the Trump campaign's success with it resulted simply from its use of Facebook's standard machine-learning optimisation.

Trump's ads were banal, but rather than trying to build the case for him they often encouraged a specific action, a conversion: 'Buy this hat, sign this petition, RSVP to this rally'. Researching the ads for a January 2020 Trump rally in Milwaukee, Bogost and Madrigal found little sign of the targeting of specific demographic groups. They suggest that instead the Trump campaign's use of Facebook began, just like my sari vendor's, by providing its system with a 'custom audience': a list of people who have already taken an action, such as providing an email address or phone number, which suggests that they are Trump supporters. Machine learning can then search for 'lookalikes': people who resemble the custom audience. But its search for likeness would go well beyond characteristics that a political sociologist might think of as influencing voting preferences. It would use the full power of the submerged portion of the data iceberg.

Thus the iterative, machine-led process of experimenting and optimising can begin. Indeed, the course I took recommended precisely this use of 'custom' and 'lookalike' audiences. Bogost and Madrigal report a 'source close to the 2016 Trump campaign' telling them that its use of Facebook was inspired by the successes of the mobile-game studio Machine Zone's machine-learning optimisation of user acquisition. But the exemplar might not have been needed: by 2016 the practices they report were fast becoming standard among those who realised how machine learning was changing advertising.

_____

The submerged portion of the iceberg is enormous, but its full power to optimise advertising requires a degree of order within it. The crucial issue is what practitioners call 'identity resolution': the capacity to discern, in an automated way and with at least some degree of accuracy, that two or more often very different data traces involve the same human being.

In advertising on mobile phones, identity resolution largely boils down to something basic and deceptively simple: whether it's the same phone. Suppose an ad network (a specialist firm that places ads on behalf of a games studio or other advertiser) shows an ad for a game, and shortly afterwards someone installs the game. If you know that the phone that installs the game is the same phone as the one on which the ad was shown, that's not proof that the ad led to the install (correlation isn't causation), but it's extremely useful fuel for measuring and optimising advertising.

In smartphones' early years, it wasn't difficult to tell that it was the same phone. Every smartphone, whether Apple or Android, had a unique identifier number, which the phone's owner could not alter or delete, and which was visible to the apps installed on the phone and the ad networks that displayed ads on it. Apple told app developers not to 'transmit data about a user without obtaining the user's prior permission', but there seems to have been no insurmountable technological barrier to it. Mobile-phone apps leaked data, sometimes on a large scale, and that began to become known. 'Your Apps Are Watching You,' the *Wall Street Journal* warned its readers in December 2010.

Apple, though, was not yet ready to leave advertisers without a dependable way of answering the question, 'is it the same phone?' By 2010, privacy-conscious people were regularly deleting the 'cookies' (strings of digits unique to each user) that websites and web advertisers had deposited in their browsers. Although cookies are a web technology not available within mobile-phone apps, Apple decided to give savvy iPhone owners an equivalent to purging cookies. From 2012 onwards, it denied apps and ad networks access to phones' permanent identifiers, and instead made available to them a 32-digit IDFA or 'Identifier For Advertisers', which uniquely identifies a particular phone, but can be changed whenever the phone's owner wants. In 2016, Apple gave iPhone users the additional capacity to 'zero out' their IDFAs. When an ad network or an app installed on the iPhone of someone who has done this asks the phone for its IDFA, it simply receives in response an uninformative string of 32 zeros.

Google introduced a similar GAID (Google Advertising Identifier), which owners of Android phones can similarly delete. But most people, me included, aren't savvy enough to alter or delete our phone's IDFA or GAID. In practice, therefore, IDFAs weren't so different from the permanent identifiers they replaced.

'Apple created IDFA', says an experienced advertising-technology specialist, 'so that advertisers could [continue to] "attribute" marketing campaigns to their app and have a [measurable] return on investment and run effective advertising'. But the scale on which IDFAs were used to link up diverse data created capabilities that Apple most likely had not fully anticipated. What it made possible, he tells me, was 'effective deterministic [targeting]. They would know that you use Deliveroo to get Chinese or Vietnamese food on a Saturday, they know that you use Tinder … They'd have known bloody everything.'

I was at first puzzled by this use of the word 'deterministic', because little in life is that certain. I now see what he meant. If a specific IDFA is associated with, say, repeated purchases within a match-three game, and an advertisement for another such game leads the player to take that up instead, it's very likely indeed that they will spend similarly in the new game. Being a whale is repetitive behaviour: you are unlikely to stop just because you have switched game.

And if you were an ad network or advertising platform that acted on behalf of multiple advertisers, then simply in the ordinary business of tying together ads and resultant app installs, purchases, etc., you collected lots of IDFAs and GAIDs with lots of records of actions on the same phone tied to them. The submerged portion of the iceberg did indeed gain structure, and the efficiency of your advertising could increase markedly.

If routine business did not bring in sufficient IDFAs, there were ways of getting more. An ad network will often offer games studios a fixed ‘cost per install’, which simplifies their financial planning. There was, however, frequently a quid pro quo. In the words of the experienced practitioner quoted above, the ad network would say to the studio, ‘if you want me to … sell you installs, you can’t waste my money by [me] serving ads to people’ who have already installed your game. So the ad network would ask for a ‘suppression list’, a list of the IDFAs or GAIDs of all the game’s existing users. That, however, could be an important resource for the ad network’s advertising of other similar games.

‘Whoever has the most data wins’, says this practitioner. The way in which IDFAs were being used to accumulate data and gain capabilities for prediction, targeting and ad optimisation ‘pissed Apple off’, he says, and may have alarmed it: user privacy is an iPhone selling point. At Apple’s June 2020 Worldwide Developers Conference, held online because of the pandemic, it announced a new privacy policy, App Tracking Transparency or ATT, which it went on to implement in April 2021. ATT tightly restricts apps’ use of IDFAs.

Talk to people in digital advertising, and you encounter lots of speculation, often hostile, about Apple’s motivations. Any large corporation is a complex organisation, different parts of which may have different, sometimes competing, goals, and Apple’s internal decision-making has not become public. Its chief executive, Tim Cook, was, however, unequivocal, saying in an October 2021 call to stock analysts that ‘we believe strongly that privacy is a basic human right. And so that’s our motivation there. There’s no other motivation.’

The hostility to Apple arises because its restrictions on the use of IDFAs threatened the most important way in which the app economy’s data icebergs are given structure. Apple was nevertheless able to press ahead because, within that economy, it has infrastructural power. That term was coined by the sociologist Michael Mann in 1984, and has become a way of thinking about the leverage that can flow from you, or a system you control, being necessary for other people or their systems to do things that they need to do.

Apple’s engineers write iOS, the operating system that controls every iPhone. All iPhone apps run within iOS, and some of the new privacy policy’s rules are directly embedded in it, making them hard to circumvent. If an app breaks the rules, it could also face the potentially catastrophic penalty of exclusion from another crucial part of the infrastructure, Apple’s App Store. There are more Android phones globally than iPhones, but the latter’s owners tend to be more affluent and therefore have more money to spend on in-app purchases, and currently they can install apps only via the App Store. (That is changing within the European Union because of its Digital Markets Act, but it seems likely that most of the EU’s iPhone owners will continue to install their apps via the App Store.)

Facebook's response to Apple’s announcement of its new policy is an object lesson in the effectiveness of infrastructural power. It protested fiercely, in December 2020 even taking out full-page adverts in, among others, the *Financial Times* and *New York Times*, describing the policy as ‘devastating to small businesses’ because it endangered their ‘ability to run personalized ads and reach their customers effectively’.

For over a decade, though, Facebook's users have interacted with it mostly via their phones, so in practice Facebook is a mobile-phone app. Instagram has always been an app, with the additional twist that heavy users of Instagram tend particularly to like iPhones because of their high-quality cameras. So non-compliance with Apple's policy would have had effects that couldn't be contemplated. 'We have no choice but to show Apple's prompt', said Facebook vice-president Dan Levy. 'If we don't, they will block Facebook from the App Store.'

The 'prompt' is a screen that must be shown when an app that wants to track you beyond its own confines is first installed (it's not required if all an app does is record the actions you take within it). You tap on the prompt screen either to accept tracking or to reject it. If you tap 'reject', the app isn't allowed to penalise you by restricting the features available to you. Your choice is recorded within your phone's operating system, and if the app in question ever asks your phone for its IDFA, iOS ensures that all the app gets is the 32 zeros.

Only a minority of iPhone owners tap the button that consents to tracking: typically around a fifth, I'm told. There's a sense, too, in which that understates the depth of the problem. As I've said, the crucial use of an IDFA is to link the phone on which an ad has been shown within app A, to the phone that has taken an action such as installing the game being advertised, B. For it to be possible to use the phone's IDFA to show that it is indeed the same phone, its owner must have tapped 'accept' within *both* app A and app B.

Without IDFAs, or something to replace them, all that an advertiser or ad network has to work with is a bunch of records of the (probably very large number of) ads it has displayed, and a different bunch of records of installs of the game or purchases of the product being advertised. Connecting up the two bunches in order to measure and optimise the effectiveness of advertising is then far less easy. 'When you break that loop', another advertising-technology specialist said to me, 'it's much harder'.

Facebook stumbled, and may have lost several billion dollars in ad revenue as it rebuilt its systems to cope with the loss of finely-granular data, which had effects on virtually all forms of advertising on Facebook, not just the advertising of mobile-phone games. In an October 2021 call to stock analysts, its Chief Financial Officer, David Wehner, told them: 'I think just the retooling of … all of the targeting and measurement to basically work for aggregated events is just it's difficult especially for smaller advertisers … the areas that are hardest impacted … I'd probably call out there, online commerce and gaming …' Other factors also buffeted tech stocks, and between September 2021 (by which time the effects of Apple's implementation of its new policy had started to become fully evident) and October 2022, Facebook/Meta's stock slid from $352 to $93.

Meta's stock has, however, now more than recovered, aided by investor enthusiasm for artificial intelligence and by the increasing success of the 'retooling' efforts pointed to by Wehner. Mobile-phone games have been less fortunate. The greater difficulty and therefore the cost of finding whales and other users who will 'monetise' has 'increased the obstacle to launching a game', says Seufert; another interviewee tells me that launching has become two to three times more expensive. Resultant falls in advertising revenues, Seufert reports, have in their turn badly affected the economics of ad-dependent hypercasual games. Indeed, the previously healthy growth of the computer-game sector as a whole (not just

mobile-phone games) has slowed dramatically, with widespread job losses: 10,500 globally in 2023 (according to a report on the publishing site Obsidian), followed by more than 5,000 in January 2024 alone.

_____

Over the last couple of years, the overt controversy sparked by Apple's privacy policy has gradually been replaced by latent, subterranean conflict between two different ways of measuring the effectiveness of advertising. The first is Apple's preferred mechanism, which it calls Store Kit Ad Network or SKAN, and offers free of charge to ad networks and advertisers. If you have an iPhone, it plays an active part in SKAN. Data crucial to measurement and ad optimisation is stored not on an external server, but in your phone's memory.

No access to the ad data stored on your phone is allowed for 24 hours plus a randomly varying period of up to a further 24 hours, the rationale of randomisation being to stop exact times being used to match up ads and subsequent actions such as purchases or game installs. Once the 24-48 hours is up, the data is sent from your phone to the relevant advertiser or ad network via an Apple server that ensures the preservation of what Apple calls 'crowd anonymity', a notion that warms my sociologist's heart. In essence, it means checking that there's nothing about the data that stands out enough (an unusually big purchase, for example) to make your phone distinguishable from at least a moderately large crowd of other phones. As Seufert puts it, the ad network learns that 'this campaign delivered an install' or purchase, but Apple's system in effect says 'I won't tell you who the person is.'

Apple's preservation of crowd anonymity is a striking reversal of digital advertising's overall trajectory, which has so far been to tailor advertising to very specific audiences, and ultimately individuals, rather than the inherently aggregate viewership of traditional broadcast TV programmes. Google, too, is building a broadly analogous set of smartphone de-individualising mechanisms, the Android Privacy Sandbox, although its planned launch at the end of this year is looking likely be postponed.

The second approach to measuring ads' effectiveness tries more fully to preserve existing practices. With IDFAs not usually available, I'm told that this can involve the use instead of Internet Protocol or IP addresses. They're nothing like a full substitute for IDFAs: when your phone is connected to the internet via a mobile-phone network, it may be sharing the network's local IP address with hundreds or thousands of other phones, and that's a form of crowd anonymity.

When, however, your phone's connection is via the wifi router in your flat or house, its IP address is your router's address. The crowd is then much smaller: it's the devices in your household using the same router. I asked one of my contacts whether whale hunting continues after Apple's changes. It does, he told me, even if it is now less precise. 'They used to know that I was a match-three whale', he said, 'but my wife wasn't, and my two kids' iPads weren't. But with IP address, they still know my household has one match-three whale in it.' Add in other information that may well be available, notably the model of phone and

the specific version of the operating system it is running, and the crowd might sometimes shrink back to one.

For that sort of reason, the use of IP addresses to help measure advertising's effectiveness is contentious. Three well-informed people have told me that it's widespread, but it's hard for an outsider to determine how important it is relative to other inputs to machine-learning systems, such as data from Apple's SKAN or the behaviour of the minority of users who have agreed to tracking. There is, however, a simpler issue: timing. If you use only Apple's SKAN, you have to wait 24–48 hours for data to arrive, and possibly longer if you want information beyond, e.g., the simple fact that an install has occurred. But if you can yourself gather the data you need (and IP addresses, for example, can be captured outside of Apple's systems), advertising can continue to be optimised in real time.

Is all of this within Apple's rules? Might it take action against it? If so, what action? The App Store tells app developers that they must not 'derive data from a device for the purpose of uniquely identifying it', and among its examples of such data is 'the user's network connection'. One of my informants, though, tells me that those who use IP addresses as an input to their machine-learning systems interpret the prohibited 'identifying' of a device as 'identifying it persistently and specifically', which an IP address does not do. Elaborating the rule might not resolve the latent conflict. As this informant says, 'the second [you] clarify [a] rule, then you create opportunities to find loopholes'. And Apple can't simply 'zero out' iPhones' IP addresses, because they're the way packets of information are guided through the internet to the correct destination.

Apple could obfuscate IP addresses by encrypting them and routing messages through a relay system of computer servers. That's what's done (for both good reasons and to hide from law-enforcement) in the 'dark web'. Obfuscating all the world's iPhones would, however, most likely spark strong government opposition, particularly in China, an important iPhone market. It would also involve a great deal of additional processing and electronic traffic, and could palpably increase the internet's already high energy consumption and thus carbon emissions.[1]

Tensions such as this haunt our attitudes to the digital economy. We desire privacy, but also want free information and entertainment, the economics of which often depends upon targeted advertising. We are excited by the capacities of giant-scale, electricity-hungry artificial intelligence, while also knowing that we have to reduce carbon emissions. We value Big Tech's sophisticated services and protected digital environments (the App Store, for instance, is good at blocking malware), but we also want to open them up to healthy competition.

Must the balancing of priorities such as these remain in the hands solely of the private sector? A glimpse of what might be possible is the current role of the UK's Competition and Markets Authority in monitoring and evaluating Google's phase-out of another of the central mechanisms of digital advertising (the tracking of users across websites using cookies) and its replacement with a Privacy Sandbox for Google's Chrome, the world's most widely used

[1] Donald MacKenzie wrote about advertising's emissions in the *LRB* of 19 January 2023.

web browser. The competition-law concerns that have swirled around Google, and fears within digital advertising that the change will increase Google's market dominance, prompted a legally-binding agreement between it and the CMA. This in effect gives the CMA the power to stop the Sandbox being rolled out if it has features that unduly advantage Google.

It's an intriguing policy experiment, albeit not an easy one. The Chrome Privacy Sandbox is a complicated set of mechanisms; the CMA's team is small; and resolving the now-familiar underlying tension between preserving privacy and fostering competition is always difficult. But tensions, astutely handled, can be productive and creative, and it's too early to write off the possibility that public policy might be central to this.

Fact checking links etc:

2013 BBC News report: https://www.bbc.co.uk/news/magazine-25334716

Candy Crush levels: https://www.pocketgamer.com/candy-crush-saga/how-many-levels-are-there/

Candy Crush download and user numbers: https://investor.activision.com/news-releases/news-release-details/celebrating-20-years-gaming-excellence-kings-milestone-journey

Activision Blizzard acquires King: https://investor.activision.com/news-releases/news-release-details/activision-blizzard-completes-king-acquisition-becomes-largest

Microsoft acquires Activision Blizzard: https://news.microsoft.com/2022/01/18/microsoft-to-acquire-activision-blizzard-to-bring-the-joy-and-community-of-gaming-to-everyone-across-every-device/

Nieborg, David B. 2015. "Crushing Candy: The Free-to-Play Game in its Connective Commodity Form." *Social Media + Society* 1(2): 1-12.

'Over 90% of UK advertisers on Facebook use the default automated bidding feature, which does not allow advertisers to specify a maximum bid': Competition and Markets Authority (2020), *Online Platforms and Digital Advertising: Market Study Final Report*, p. 17. Available at: https://www.gov.uk/cma-cases/online-platforms-and-digital-advertising-market-study.

Ian Bogost and Alexis Madrigal, 'How Facebook Works for Trump', *The Atlantic*, 17 April 2020: https://www.theatlantic.com/technology/archive/2020/04/how-facebooks-ad-technology-helps-trump-win/606403/

Apple – apps 'cannot transmit data': Scott Thrum and Yukari Iwatani Kane 'Your Apps Are Watching You', *Wall Street Journal*, 17 December 2010. https://www.wsj.com/articles/SB10001424052748704694004576020083703574602

Tim Cook quote and Facebook's $8.3bn ad revenue shortfall estimate: https://www.ft.com/content/4c19e387-ee1a-41d8-8dd2-bc6c302ee58e

Facebook 'no choice' quote: https://www.facebook.com/business/news/ios-14-apple-privacy-update-impacts-small-business-ads

Facebook call to analysts, Oct 25, 2021**,** quoted by Eric Seufert: https://mobiledevmemo.com/unpacking-atts-impact-on-facebook-revenue/

Facebook/Meta stock price: https://investor.fb.com/stock-info/

App Store's prohibition: https://developer.apple.com/app-store/user-privacy-and-data-use/

10,500 layoffs in 2023: https://publish.obsidian.md/vg-layoffs/Archive/2023
5,600 layoffs in January 2024: https://techcrunch.com/2024/01/25/microsoft-lays-off-1900-employees-in-activision-blizzard-and-xbox-divisions/