

Transcript:

Databite No. 137: Adtech and the Attention Economy

Tim Hwang and Moira Weigel

Recorded: 10/15/20

Moira Weigel: Hello, everyone and welcome to Databite #137, "Adtech and the Attention Economy," featuring Tim Hwang, the author of *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet* (FSG Originals, October 2020). My name is Moira Weigel, I'm a sociotechnical security researcher here at Data & Society, and I will be your host for today, supported by my team behind the curtain—CJ, Rigo, and Eli.

For those of you who don't know us yet, Data & Society is an independent research institute studying the social implications of data and automation. We produce original research and convene multidisciplinary thinkers to challenge the power and purpose of technology in society. You can learn more about us through our website at: <https://datasociety.net>.

To begin, I ask you to join me in acknowledging where Data & Society was founded, Lenapehoking, a network of rivers and islands in the Atlantic Northeast we now refer to as New York City. Today, we are connected via a vast array of servers situated on stolen land. We acknowledge the dispossession of Indigenous land by the data-driven logic of a white settler expansion and uplift the sovereignty of Indigenous people, data, and territory. We commit to dismantling the ongoing practices of colonialism and its material implications in our digital worlds knowing we interface with power differently based on our race, class, gender, and ability.

Now, I want to say just a little more about our speaker today—although Tim was a 2014–2015 inaugural fellow at Data & Society and might not need much introduction here. Tim is a writer, lawyer, and technology policy researcher based in New York. Previously, he was at Google where he was the company's global public policy lead on artificial intelligence. Forbes has also dubbed him "[The Busiest Man On The Internet](#)." So I'm very happy to welcome Tim back to Data & Society where he has a lot of history. It's a special privilege for me to welcome him here because he was also one of the inaugural authors of *Logic* magazine, a magazine I co-founded back in 2016. I first got to know Tim when he wrote an article for our first issue called "[The Madness of the Crowd](#)." It was really a brilliant analysis—and a groundbreaking analysis—of this mood shift that I think a lot of us observed around late 2016—early 2017, where a lot of the positive properties we'd attribute to social networks and digital media suddenly seemed to be flipped,

and the very same features that we thought would promote cognitive surplus and good social movements now had brought about a presidential election—and all sorts of other problems that hadn't necessarily been anticipated in earlier rhetorics. I think now this has become a kind of common-sense narrative, maybe even a cliché—and we can talk about that with the so-called “techlash.” But Tim's book, which he'll be sharing with us today, builds on that argument and really advances it and intervenes in the state of the field, arguing that both the hype and the critique of big tech companies may be more similar than we're used to thinking. It's a provocative argument that upends the assumptions of a lot of the most prominent tech critics. Tim was joking at an event earlier this week, about his ideal reader, that his goal is to make absolutely everyone mad—which seems noble to me. It's also a topic of keen interest to a lot of us here at Data & Society: it has a lot of implications for how we think about disinformation and media manipulation; also, if we're thinking about privacy, digital health, tech addiction; and speaks as well to this new initiative and conference against platform determinism. Methodologically speaking, it really highlights a core value of getting close to the machine: building specific ground-up accounts of how our sociotechnical infrastructures work. So I'm really excited for the conversation and, without further ado, ask you to join me in giving Tim a silent virtual round of applause. Tim, maybe you can tell us a bit more about the book.

Tim Hwang: Sure, definitely. Thanks, Moira. Let me begin with the origin of the book, which is really based in the two years I spent at Google, basically running public policy on AI and machine learning. And I think one of the most striking things about being at Google is a lot of the discussion on the day-to-day is about self-driving cars, or artificial intelligence, or uncomfortable partnerships with the military industrial complex. What was fascinating is that—if you look at the balance sheet, you look at the SEC filings that Google has to do on a quarterly basis—that's actually not where a lot of the money comes from. In fact, 80% of Google's funding is still from ads. And, of course, people know this, but if you ask them: “So, how does that ad system actually work? Can you walk me through it on a step-by-step basis?” It becomes “here be dragons” territory very quickly. And so there's this very funny irony that the business model—this core business model of the internet—is, in some ways, a rumor. It's an urban legend, right? We know these businesses run on ads, but we actually don't know much more than that.

And so the beginning of the book was to really take a look at this business model, which has been the rocket fuel—the financial engine—that has driven the internet over the last 10–15–20 years. What's interesting, probably from the point of view of Data & Society, is how much this specific kind of online advertising infrastructure that now exists is really built on the kind of

credibility and the power of data. The dream is that we have an incredible amount of information about consumers, and we have the tools to pinpoint with laser accuracy a message to the person, and we can build these behavioral models that allow us to do persuasion in a way that has never been possible before. And this has been the claim of industry, right? This has been the argument that Eric Schmidt made early in the days of Google to argue for why Coca-Cola and the big companies of the world should invest in and buy advertising on Google versus television, or magazines, or billboards. And I think it's really interesting that even critics have brought this narrative: critics have said “The problem with Mark Zuckerberg is he has this mind control, effectively. He has this ability to reach into the minds of Americans and push around our belief systems.” As yet—despite these business claims, these claims among the sort of most vociferous critics of the tech industry—we have these really interesting stories pop up, especially in the last few years. So I'll talk about two of them, but the book is chock full if you really like this sort of thing.

The first one is: A few years ago Procter & Gamble, which is one of the biggest advertisers in the world, decided that what they would do is cut their digital advertising spending, not just by a little but by a lot. They cut it on the order of about \$200,000,000—just sliced it out of the budget. And it ended up being this great natural experiment in trying to figure out: “Does all this online advertising stuff actually make a difference?” And they waited. They waited and they waited. And they actually discovered—they reported the year after they did this—that there was actually no noticeable change in consumer behavior. And, in fact, the reach of their advertising had gone up just about 10% largely due to efficiencies from cutting out all this money from the budget. That's one story that I find really interesting.

A second one is that Google itself, a number of years ago, did a report that indicated that close to 60% of ads are never seen at all. That is to say that *they are delivered*, but there's not even the chance for them to persuade someone because they end up, you know, below the fold or oddly placed or otherwise not delivered—which is just staggering. You think about 60% of all this activity that's happening in this market—this jet fuel that has driven Silicon Valley—being basically worthless. It starts to raise some really interesting questions. The argument of the book is basically to say: “Okay, what is going on here? And is this financial engine, ultimately, a kind of bubble?” And the way it goes about doing this is basically to say: “What is the kind of advertising ecosystem that has been built?” Because, when I say advertising, sometimes people think of Jon Hamm in *Mad Men*: guys saying really terrible things in wood-paneled offices. But, actually, the modern advertising ecosystem is very much a commodified marketplace. It is a

financialized marketplace. And, in fact, the explicit goal of a lot of the people that architected the early days of what's known as the "programmatic ad infrastructure" wanted to design it in the image of the stock market. And so the game of the book—the argument of the book—is basically to try to make the case that we can actually use the pathologies of these financialized markets as a way of thinking about the future of the attention economies of the web, and that what looks like an incredibly solid—rock-solid granite—footing for the internet might be more delicate and brittle than it looks. And ask the question: What does it mean if it all disappears? That's just a quick intro, but I know we'll get into a lot more of the specifics.

Moira Weigel: Thanks so much, Tim. I already want to break script and ask you a specific question about Google. In the book, you talk about search engines as early pioneers of this new model of advertising—and AdWords and AdSense specifically. I was wondering how specific what you are describing feels to Google? Are there any differences between programmatic advertising on Facebook and Instagram and other social media platforms in contrast to search engines? Or does it all feel like part of the same bubble?

Tim Hwang: Search is really interesting because, in some ways, it was the prototypical "case" for programmatic advertising. Basically, Google had built this search engine that delivered thousands and thousands of hits a day. There was so much attention flowing to it. That, really, the rise of programmatic advertising comes from the task of trying to build an advertising ecosystem that could keep up with the speed at which Google was growing. There's actually a great history of this—actually, interestingly, a labor history—where early in the days of Google, there's this office in New York with a bunch of old-timey ad people who are still getting people on phones and saying, "We need you to put a banner ad on Google." And there's this constant suspicion that Mountain View is trying to automate them out of a job, and there's this interesting tension between what is valuable labor in the world of advertising. I think that template that Google built is similar to what you see on a Facebook or a Snapchat or an Instagram, for that matter. The main difference which is debated about in the ad industry is: What is the kind of intention that is captured by the advertisement? So in search, the adage is that search tends to be more effective because it is trying to get you when you have expressed a want into the world. So, famously, "mesothelioma" is a search term that is incredibly expensive. And it's in part because when you search for it you typically have it, and it's very powerful to advertise against that term. In display advertising it's a little less clear. You're just browsing through social media and Facebook's like: "While you're here, would you like to buy a mattress?" It's unclear if that captures intent in the same way. So the underlying plumbing is the same, but

I think the channels capture very different aspects of experience and what is attempting to be commercialized..

Moira Weigel: That makes a lot of sense. Now I want to ask you about what the other really expensive search terms are but—unless you have a really good one ready in mind—I’ll spare everyone that. You talk a couple of times in the book about this famous ad industry adage, this John Wanamaker line: “I know half the money I spend on advertising is wasted. I just don’t know what half.” I was curious, reflecting on those passages in your book, how novel this problem you’re describing of the difficulty of commodifying attention and capturing attention is, and what—if there are new sources of difficulty and opacity in a digital environment as opposed to in the *Mad Men* smoky, martini-filled room—what some of those are and if you could elaborate on them.

Tim Hwang: As a preface to this question, there’s a great critique I have received from some people, which is: “Why have we built this enormous surveillance capitalism infrastructure if it doesn’t actually do anything?” One response is that there’s actually, in some ways, a need to differentiate digital advertising from earlier generations of advertising. And, as I talked about earlier, the fact that it is data-driven and therefore “more scientific” as a form of advertising was actually, and is still, a very powerful comparative advantage that digital advertising claims as against magazines and billboards and TV and all this stuff. In some ways, the claim of the book isn’t necessarily that digital advertising uniquely doesn’t work, it’s more that this advertising is potentially *just as bad* as all the other forms of advertising that we’ve had in the past—in terms of measurability, in terms of the ability to see whether or not it’s actually making an impact. I think there’s two cases here that are sort of unique, though, in the extra layers of obfuscation that we see in programmatic advertising. The first is in the case of brand safety, which I find is really interesting. Which is basically that this kind of programmatic ad system is so big, so complex, and so automated—it turns out that sometimes your messaging ends up next to a video that a white supremacist produced. And despite the greatest minds of a generation using AI and all these other tools to try to prevent this, the marketplace just can’t eliminate this problem from the system. And what it indicates to me is that there’s a great deal of opacity in why an ad ends up in someone’s browser—it’s actually unclear even to the people that run the system—which I think is really interesting.

A second kind of opacity I’ll point out is really what I call the “monopolist’s opacity.” So a number of years ago, Facebook said everybody needs to pivot to video—video, video, video—it’s

going to be the greatest thing in the whole wide world. And a lot of people fired a bunch of journalists and said, “We’re going to hire a bunch of video producers to do this.” And it turns out, actually, that Facebook, either through incompetence or malice or some combination of both had overestimated these statistics about how much people were watching Facebook videos by something like 60–80%. I think it’s a really interesting kind of opacity because it’s built on market power. There’s no way for advertisers to force Facebook to give more transparency, so they have to take the word of these companies. And so I do think there are special aspects that make this market more opaque than others, but I agree with you that I think there’s a lot of similarity to earlier generations of ads.

Maira Weigel: Since you’ve spoken to the problem of monopoly power—and, obviously, that’s a topic on a lot of our minds given the new news out of Congress on antitrust—do you think your redescription or rediagnosis of the state of online advertising can help us think about monopoly in a new way? Does it change the way we think about antitrust problems in the tech industry?

Tim Hwang: I think it does in two ways. One of them is that it’s unclear if the “data advantage” is actually a real one, in the end—which I think is really intriguing. There’s been this accepted wisdom that basically the data collected by Google allows it to create products that are extremely sticky, and that self-reinforces, and, therefore, there’s no way to break through the moat of these companies. In some ways the account of the book questions whether or not that “data advantage” is real. It’s interesting, I think that there was a critique of big data that was like: “This big data reveals nothing.” But we don’t apply that same logic to the data these big companies have and their market position. In some ways, we should view their claims, or even those fears, with some skepticism. The second one is whether or not there would be much harm in blocking access by these monopolies to certain types of resources. A lot of the companies have said, “Look, the social contract is [that] you get the product for free and we get the data to target ads to you.” But there’s a lot of evidence right now that, even when you eliminate all that data, advertising can be just as effective. And so it also eliminates that claim that there’s a trade going here that we’ve all “agreed to.”

Maira Weigel: That’s really interesting, thank you. Towards the end of the book, you do get into suggestions for what is to be done and what are some measures that could be taken to ameliorate or stall or control this impending crisis that you see possible in the attention marketplace. I was wondering if you could speak to that a bit. You have some interesting

historical analogies as well as proposals for things that haven't happened yet. Could you tell us a little bit about those suggestions?

Tim Hwang: Sure, definitely. I was joking with someone earlier this week that there's one way of reading the book which is a kind of attentional Marxism. Where you're basically like: "There's these contradictions in the marketplace that will just bring it down, and so, if we hate advertising, we just need to wait—and the arc of history will take its course." I guess I'm not so certain. I do think there is this time bomb, but the interesting thing about market bubbles is that they're very robust. If you think about economic indicators in 2007, they would have told you that these subprime mortgages were the way to go—collateralized debt obligations were "the thing"—and I think there's a similar case here. And my worry is that letting it grow and grow and eventually blow up will exact a really big human cost. I think it's more than just a matter of Mark Zuckerberg having a billion less dollars. You think of the entire media ecosystem and news ecosystem that rests on this programmatic ads structure, and I think there's very real concerns about what happens if that breaks all of a sudden. In the book, I advocate for the notion that we can deflate this bubble: the idea is to reduce its credibility and have a handbrake on the idea we can slow this down and eventually create the room for new models to emerge.

There's two things that I advise here: the first one is taking a look and inspiration from the capital markets. There's a really interesting Great Depression-era history about how The Securities Act of 1933 came to be—but the end result is that the government mandates a certain level of transparency in selling stocks, and if you don't offer that level of transparency there's legal consequences. Obviously, this is not a solution for every financial crisis because obviously there were financial crises after 1933, but it was actually really powerful as a way of creating more transparency in the marketplace so we even knew what was going on in the first place.

So one of the arguments in the book is: We should create similar transparency regimes in these attentional marketplaces. We should have more data on how things are targeted, whether or not they are effective, and how much fraud is in the system. The second one, which is a little bit more activist-y, is the notion of a research outfit that is a combination public policy research institute and group of trolls, essentially, that is working to leak documents or otherwise put pressure on the ad industry to change. And, as part of that, I have set up an initial experiment on this front called AdLeaker: it's a signal number whereby people in the ad industry can drop me things that they're seeing. But I think we need more experiments like that to put some more informational pressure on the industry.

Moira Weigel: At risk of asking you to speculate, what kinds of relationships would you see between these entities you're describing and older institutions like the Internet Advertising Bureau (IAB) or older attempts to create mechanisms for accountability or transparency in this marketplace?

Tim Hwang: What's an interesting thing to know is that the space is full of frenemies. And I think, because the space is full of frenemies it is brittle in ways that may be both powerful and sort of amusing to stir the pot in. One of the big tensions is between these buyers of ads, who basically have felt that for the last decade they've been forced to have to buy ads from Facebook, and the companies themselves who have actually not offered a lot of transparency. I think there are these sort of interesting bedfellows that can be created between, say, Procter & Gamble—maybe you should underwrite people who are trolling the ad tech industry—because it is in your interest to know what's going on. I think there's interesting games to be played there in how we can get aspects of the ecosystem to fight one another in ways they haven't in the past, but that would be quite productive.

Moira Weigel: Wow. I hope we're going to get the Procter & Gamble critical trolling fellowship call circulating.

Tim Hwang: It will be laundered through a series of intermediaries. It will be very cyberpunk.

Moira Weigel: It sounds perfect. I alluded to this in the introduction, but because as your book demonstrates and you have been talking about so much of the internet is built on advertising and on the logics and imperatives of advertising. I think that part of what's so powerful about what you've done is I really think it has implications for just so many different areas of critical work on technology, and since we have a Data & Society audience here, and a number of folks who are engaged in their own research and projects grappling with the power of big tech, I wondered if you could speak a little bit to how you imagine or how you hope this book might change how we do work in areas like disinformation or thinking about privacy, about public health and the tech addiction issue, or fiction—depending on how much of a certain Netflix film we want to talk or not. But I wonder if you could elaborate a little bit on how folks in the audience who are engaged in critical work might do their work differently as a result of your book.

Tim Hwang: So I think there's two things: One is maybe a frame or a lens for thinking about some of these issues; and I think the second one is a matter of activist strategy—which I'd love to

get into. So the lens that I really came away with, from the book, is the degree to which ads are almost the whole way of reading the surface of the web. In the sense that you can look at any feature online and say, “How did advertising cause this to be.” And, more often than not, because advertising has been so ubiquitous in many of the services we think of as just online, it becomes this great way of “reading” the internet that is. The example that I've been using a lot—just because I think it's very salient—is the incidence of the “like” button. At this point, if you had a social media platform that *didn't* have a “like” button, you would probably be slightly frustrated. You'd be like: “I can't interact on this platform in the way that I'm used to.” But, of course, the reason that you have a “like” button is because having an explicit indication that you are engaging with a piece of content makes it easier for someone to profile what you like online. It also makes it easier to figure out how much engagement occurs with an ad that you've launched. You can give a whole account of this very, very tiny building block of the web as being generated by these broader economic forces. And I think there's a worthwhile kind of intellectual exercise there—it's both a parlor game and a useful kind of research agenda—in kind of deconstructing all these UX elements in the logic of advertising. That's the sort of thing that I've personally felt is compelling and interesting, coming out of the book, that I want to think about some more.

I think the second one, though, is a kind of interesting challenge that the book throws down for tech critics in the space, which is: To what degree does tech critique end up forwarding the claims of the tech industry? And the degree to which the strategy for tech critics should be to make a fool of the tech industry or to dramatize the danger of the tech industry? And I think, genuinely, there's a battle between those two incentives. In general, you want to provoke a mass movement. So sometimes it's really critical because David needs to fight Goliath; we need to say: “Our adversary is Mark Zuckerberg and he has a mind control ray and we need to defeat him.” On the other hand, it gives an incorrect account, in my mind, as to what the source of power is and what we should really be concerned about. *Maybe we should shred surveillance capitalism and just talk about capitalism.* It is kind of to the degree in which technology gets in the way of resolving the deeper issue. I do think that one of the interesting aspects of the ad industry is how much it is a mirror image of the stagecraft of data that we see in other places. So much of it is the authority of having this huge dataset and the power of being able to target that. Maybe the goal is that we should reduce the public sense that that is really where the danger lies. It's something that's certainly been done with great success in the AI case, where it feels like a lot of the critique of AI now is that it *just simply doesn't work* and that's why we shouldn't use it—and maybe

that's a tactic that should be used elsewhere as well as we talk about how should we frame up these issues for the public and try to create change. And obviously, *The Social Dilemma* is in there and I'm happy to go into it.

Moira Weigel: No, I promised myself we wouldn't talk about *The Social Dilemma*, despite it being my fault for having brought it up.

Tim Hwang: Talking about it is to forward the message of it.

Moira Weigel: Well, precisely. We don't want to participate in the radical change from time spent to time well spent, which I believe is now also a Facebook slogan. I imagine some of us in the audience today are scholars and some folks are maybe activists or people who are engaged in different ways or who want the force of their criticism or research to reverberate in different ways. I was wondering if there are times when it's useful to have a big narrative like surveillance capitalism, or maybe even a narrative like tech humanism? The idea that your kids are going to become addicted automatons to their phones might be more powerful in motivating certain people at certain times than the idea that massive corporations that control huge parts of our democracy need to be brought under democratic control. I am curious whether you see anything about the present moment that indicates particular opportunities or a need to be strategic, one way or another? Do you think that those big "David and Goliath" narratives have done the work of setting the stage and making a broader public concern about these issues, and now is a time for more nuanced description and critique? Or is it the opposite, we still need the David and Goliath story? Or is it always both—that you need some people doing one and some people doing the other? Any thoughts on timing right now?

Tim Hwang: Sure. I think about this a lot in the context of a lot of the debates that have happened around machine-learning fairness. For a long time, there was a productive ecosystem between the people who were hammering companies using AI—but mostly in the sense it was this crazy dangerous thing that would destroy everything and be like your nightmares from sci-fi—and then there was a group that would be on the other side of the game, basically talking to the companies and saying: "Look, those activists out there are crazy—but while we're here you should maybe implement better privacy practices with your data." And so there was a kind of productive ecosystem between the activists on the street and the incrementalists on the ground. But I think part of the trouble is that, eventually, a lot of the companies were like: "Oh, you got us! We want to be more ethical around AI now." And then the battle of co-option came, and the

question was: How far did we want to push versus consolidating the gains that we thought were gains from pushing the companies in a certain direction? I do think my experience in that lends me to the idea that there is a need, particularly at this time, to keep the standards up. The thing that I struggle with is that I do think the drama of the David and the Goliath, and the drama of the mind control ray in Menlo Park, is really necessary—because in order for some of the really critical changes to happen, you need to make it a mass movement. And, for that, I think you sacrifice a little bit of accuracy—and I think that's okay because I do think that will be the mass politics of trying to get change to occur. Because, at this point, the tools that really need to come down are the force of law or the force of the state, to the extent it can be reclaimed. Those things take more than these kinds of incremental steps.

Moira Weigel: You're someone who has worked in the tech industry, you've worked in research and policy, you've had this very interesting and very agated career, and it's given you a lot of insight into how the programmatic ad industry works on the inside. I'm curious whether you have ideas about how researchers who haven't worked in tech, or maybe they have, how different kinds of stakeholders in this space can do the kind of research that you've done or embark on developing a different language for describing the platforms. Speaking with you, I can never forget that platform itself is an industry hype term. As I mentioned earlier, there's this callout at Data & Society for this Against Platform Determinism conference or workshop, which seems really salient to a lot of what you're talking about. I think part of why many people end up using the hype language is because it is really hard to get information beyond the press release if you're not on the inside. I'm curious if you have thoughts or advice about methods that scholars or other kinds of researchers can use as we build out a different kind of account.

Tim Hwang: Yeah, definitely. I had a conversation with my parents over the weekend and I think they still want me to get a real job at some point. There was a lot of benefit to being at Google for a few years, and it's in part just because it's useful to know what parts of the company are fighting one another and the degree to which these companies are these mass nations that have these constituencies that are always warring with one another. I think that has been really helpful both as a matter of scholarly account—which is how do we understand what's going on from the outside—and then also as a matter of activism strategy, which is: sometimes you want to create messaging which actually causes parts within the company to gear grind with one another. Part of this is alliances, part of this is playing frenemies against one another. I do think that is helpful and sharpening, in terms of the research and the work. One of the continuous structural questions in the space of tech policy is: How does one get into tech policy anyways? It

doesn't seem like there has been any particular pipeline to doing this. If I have one weird tip—to use the language of banner ads—it would just be that I can't hold a job for more than 24 months. I don't know if that's a helpful response to your question.

Moira Weigel: I think it's a very practical question, this question of how to get information about how companies and technologies work, and how to maintain the inside/outside stance that lets one develop the kind of critique that you've done.

Tim Hwang: For what it's worth, the final thing that I do mention is that I do think anonymity ends up being a powerful tool here. It is a little bit like the Voices of Valley project: Can you get people—at all levels of the company—to disclose things to you? You play a little bit of a spy in making that happen.

Moira Weigel: Just to mention, Tim does have a project to help people disclose information if they want to at AdLeaker, which maybe CJ dropped in the chat. I wanted to ask you a sort of more human question or a slightly less conceptual question. I know you've been incubating and working on the ideas in this book for a while, and, as I mentioned to the audience, Tim wrote this great piece which I see as tied up with these ideas for *Logic* all the way back in late 2016 or early 2017. I was curious what new discoveries you made as you tried to put this argument into book form? Was there anything you learned in your research, once you actually started working on it as a book, that surprised you? Were there ideas that you'd been developing that changed as you wrote them out and expanded them? Could you speak to that process a bit?

Tim Hwang: I've been working on this for a little while. Some of the arguments in the book come from a kind of white paper—or working paper, if you will—that I did with a friend Adi Kamdar a number of years ago called “Peak Ads” [Full title: “The Theory of Peak Advertising and The Future of the Web”]. If you want to play the game of metaphors, the game there was can we use the fuel industry, natural resources, as a way of talking about depleting attention against ads and whether or not that has impacts on the ad industry. And so this has been kind of percolating for a really long time, and a number of people have cited this terrible self-serve website I set up with a PDF on it to publicize this paper, and it was a useful exercise to sit down and be like: “Okay, how do I actually articulate this argument in a cleaner, long format way.” I would say the biggest surprise—which is less of process but more of substance—is we have this account of Silicon Valley which is that the engineers are always in charge. And the ad industry,

or the programmatic ad industry, is an interesting case where it is actually flipped: where the engineers designed a lot of the infrastructure and did a lot of the work to make the real time bidding algorithmic trading system actually work—but the people who actually oversaw its architecture were former Wall Street traders, they were economists. There actually is this moment where it's actually not engineers in the driver's seat. I think it is interesting that not everything Silicon Valley produces is a product of Silicon Valley ideology. And that, in some ways, it ends up being a Trojan horse for *other kinds of ideologies* that are playing out. I thought that was quite surprising, because my thinking on the ad industry was previously that this was largely the tech sector disrupting traditional marketing, where it's actually finance—through the medium of technology—disrupting traditional advertising, which is a much more complicated and interesting story.

Moira Weigel: That's fascinating and brings me back to the office of Google people in New York City. Did those people lose those jobs in the end? Was that role automated?

Tim Hwang: They were. A lot of people were automated out of a job.

Moira Weigel: So a different New York finance tech dynamic persevered, but not the folks in that Google office.

Tim Hwang: That's right.

Moira Weigel: We have so many good questions coming up in the Q&A, and I thought they are meaty ones that maybe will take awhile, so I thought I'd shift over to asking you some of them. We have a question from [audience]: “I see a paradox or puzzle. On the one hand, there's a lot of research that you cite that says product ads and campaign ads don't work. On the other hand, there's been all this research—including research from Data & Society—about microtargeting and political radicalization. How do you think about this paradox? How is advertising different from the political radicalization or these other forms of persuasion or do you see it as a paradox at all?”

Tim Hwang: I think there's one way of reading the book, which is: “Tim believes advertising never works.” To be clear, I don't think that's actually my position. We have even personal cases where you're like, “I saw an ad, and I was so persuaded I bought this mattress on Instagram.” And so I don't necessarily want to make the claim that all advertising doesn't work. What I try to do is think about advertising as a marketplace and to think about the health of the marketplace

as a whole. Part of the problem here is that we do have lots of cases in which advertising can work. The question is just—structurally, as a marketplace—is it actually the case? I think in the case of political advertising, it's complicated. For example, we have this Cambridge Analytica report that came out earlier this week that said, “Look, all the psychographic targeting didn't really make a material impact on Brexit.” On the other hand, this argument was being made to me by Cory Doctorow who has also been thinking a lot about ads, who was basically like: “Look, the main benefit of online ads is that you can tell constituencies that you have terrible opinions that you don't want other people to know about.” You can basically be like, “Tim, I'm secretly a white supremacist.” That's a very effective way of mobilizing certain types of voters. It's, again, very hard to experiment with, but my response to your question would be that I think we do have cases and that not all forms of advertising are made equal. I'm sort of in a situation where I'm trying to parse out what are the ads that really matter versus what are the 90% of ads that are actually circulating on the internet every day. That's an ongoing project. I think it's an empirical question that requires a lot of work to answer correctly.

Moira Weigel: Thanks. Matt Goerzen had a follow-up question in response to Baba's question asking: How do Ad Tech systems and recommendation algorithms differ in terms of efficacy and measurability? I think that was implicit in your answer, that maybe a white supremacist video on YouTube is different from an ad for shoes or diapers or an energy drink that pops up in between, but also recommendation algorithms and Ad Tech we talk about together, but separately in this space, do you think there are important differences in how they can be measured?

Tim Hwang: I do think there are. Programmatic advertising—this high frequency algorithmic trading of attention, constitutes an enormous percentage of the ads that are distributed on a daily basis online. And there are smaller segments of online advertising that aren't traded through programmatic—one of them is “Spon Con,” or what is known as sponsored content, which is: you pay Tim, and Tim is on Twitter being like, “You should really buy this mattress.” That's very hard to detect when it's occurring, it is very difficult to ad block, and it looks a lot like content we think might actually be influential in a way that a banner ad is not. A recommendation algorithm has similar characteristics—which is how is the recommendation algorithm delivered to you, what does it look like, do you think about it as an ad or intrusion versus content you like—and all of those influence the impact of these systems. One of the questions that is asked by the book is that it is programmatic advertising that has produced the most money and allowed companies to scale the fastest. So it's unclear whether or not these

segments of advertising that don't have these same pathologies could grow or replace the automated systems that we have right now.

Moira Weigel: We have a question from Steve Perkins, who asks: "Would stronger data privacy laws like the California Consumer Privacy Act (CCPA) or General Data Protection Regulation (GDPR) in the EU slow down or change programmatic advertising real time bidding? What kind of impact do you see those regulations having on this industry?"

Tim Hwang: Definitely. One thing you learn looking into the history of financial bubbles is these bubbles continue for a very long time and there's this gap between what people think something is worth versus what it actually ends up being worth. Essentially, the bubble pops when people realize, or there's a cascading stampede of people who feel, "Oh my God, what I spent all this money on isn't actually worth very much." These privacy laws like GDPR or CCPA have the potential to create this kind of Looney Toons running-over-the-cliff experience, where basically, these privacy laws will break the ability for you to collect data. Then the question is: is your advertising any worse than it was before? And I think, in many cases, the answer will be no, which will leave the ad industry with a really big question which is: What are we collecting all this data for? And did that data ever amount to anything? And I do think that that will potentially create the precursors for a larger bubble.

Moira Weigel: There's a question from Adam Perry who asks if you could speak a bit more to the financial market analogy put out by programmatic agencies and platforms: Is it just a metaphor? Is bidding an effective price discovery mechanism? Does it actually resemble the stock market? Maybe you could elaborate more on the work that that comparison is doing in the book and how literal versus metaphorical it is.

Tim Hwang: In the early days of the programmatic ad ecosystem, it was very literal. You actually have these incredible articles that appear in *The Wall Street Journal*, which were like "Pretty soon, you'll be able to sell attention like pork bellies." That's an actual headline. A lot of people who did these companies were based in New York and they took the capital markets as an explicit thing they would use to architect their system. Hal Varian, who is this famous economist at Google implemented a lot of auction bidding theory into how ads are bought and sold online. And so I think the connection is very, very clear in the early days of the generation of these systems. Now that we're 15, 20 years on, it does feel like there are changes in the way this works. One of the arguments that you sometimes hear from ad people is: "Well, it's not

really like stocks because once you buy a stock, you hold it and then you resell it. but in this case, you just buy the right to show someone an ad at the point they upload a page.” So I think there are small differences that are like that in the comparison, but most of the arguments that I make in the book are based on the underlying psychology of market bubbles, and they do rely on “attention is like a stock” a.k.a we will have 2008 again.

Moira Weigel: That's a useful clarification. I think of you as using it as an exploratory device or grid to think through all these different issues in the book. We have a question from [audience], which builds on what you were just speaking to, who asks: How would you connect your critique or demystification of the legitimacy of data-driven optimization to what happens in the management sciences and behavioral sciences at the most prestigious and influential universities in the U.S., e.g., recent Nobel Prizes for auction theorists? It strikes me that a lot of what you are describing has ramifications beyond Ad Tech as well, and maybe you could speak to that.

Tim Hwang: Definitely. I've been thinking a lot about the idea that there are two prototypical semantic wars that are going on around technology. One is the semantic war over what technology is or what it should mean. Then I think there's a battle over what this label should actually apply to. So if you imagine a map, there's kind of a battle over the boundaries of this territory, and then there's a battle over what the territory itself implies. I do think that some of these data battles kind of play into this in a number of different ways. One of them is: how much should we consider this data-driven optimization to be something with power versus just a kind of theater, a kind of performance art, that speaks the language of power? I do think that has ramifications beyond just talking about technology. I do think, for example, the legitimization of certain types of modelling through Nobel Prizes and stuff is kind of built into that as well. But it is also, again, the question of boundary or the question of territory, which is: how much should a thing be considered uniquely data-driven development versus the data-driven development of qualitative research? I do think that is also part of the interesting battle that plays into some of the university culture and the machinery of authority around some of these systems.

Moira Weigel: Yeah, that's fascinating. I think in the academy we can talk about everything as a kind of technology—at least in my corner of the academy—but those two boundaries you talk about draw useful distinctions. One last question to pull from the Q&A: Do you see an alternative funding model—given how important ads have been as a funding model for the

internet we have or the internet as it is? Do you see alternative models that could replace programmatic advertising? Do you see that as necessary, desirable?

Tim Hwang: Yeah. I think there's a big question, and reasonable minds can differ about whether or not replacement is desirable. Imagine we basically came up with the capitalist version of free energy. We basically were like: "We're going to pull this advertising heart of the internet and then we're going to just plug it into this thing that prints money." I don't know, it's like some blockchain thing. Would we want it? Because, arguably, one of the problems—one of the pathologies—of technology platforms is their incredibly fast growth and the incentives that puts on organizations that need to have that level of growth. In some ways, I think programmatic advertising has given us very unrealistic expectations about how companies should grow and develop because the idea is: "Look, if you're not turning billions and billions of dollars within five years, it's not worth it." And that's a kind of growth pattern that is only possible with this kind of advertising flywheel you've created. So I think, first, there's a question of values, which is: Would we want to replace this with something that grew and scaled in the same way? One of the things in the book is not to necessarily argue that all advertising should be removed. The real question you should ask people who take that position is: Do you think advertising should be the monoculture that drives the biggest companies in the world? And I would humbly submit: No, I don't think that should be the case. Particularly when you consider advertising has stifled a lot of other potential business models in the space. I think I told this anecdote a few nights ago at City Lights, but I've talked to a couple of friends who have pitched companies in Silicon Valley, and the response has been: "Well, why don't you just do it on an advertising model? It would scale much faster and we know it works." And that ultimate conservatism of venture capital has actually shoved a bunch of businesses in that direction. My hope is that we have a diversity of more business models—that it is something more robust that the biggest companies making their money from this one thing. I do think you see some interesting models. The obvious one is subscription—which has its access questions—but in some ways would force the question of what's actually fundamental here. Do we want to subsidize access or require that you provide free access to certain types of people? I do think there's a lot of benefits to subscriptions. I do think also that the media is experimenting more with co-ops—which is also an interesting model—not just the business model, but then what is the distribution of power and wealth that occurs within these organizations, which is equally key. There's innovation on multiple levels is what I would say.

Moira Weigel: I had said that would be the last question, but I saw one more that I want to pull out of the Q&A and ask you from [audience] in Berlin. It's a methods question, one [she] likes to ask all kinds of tech researchers and writers: how do you study and write about something that changes a lot? What's your strategy for keeping on top of your narrative, even as it changes, or is tech not actually changing as much as we think?

Tim Hwang: Again, I'll drop another weird tip, I suppose. One of my passions in life is the collection of obscure or niche trade journals. And it is really interesting the degree to which there is this professional shoptalk—which is actually very hard to find online or in the mainstream news or in the scholarly literature—mostly because it's very obscure and very niche, and when you encounter it, it just looks like complete garbage. One of the ways I've kept up to date on programmatic advertising is to read the trade journals, just because some of them get into these very niche discussions but I think they're really useful for trying to figure out what's going on here. This is a community that largely kind of doesn't want to be found often. And so there's kind of a security through obscurity that they've enjoyed for a really long time, but I think you can break through that wall by trying to dip into their professional literature, for whatever that's worth.

Moira Weigel: When you get a tip from Tim Hwang, you want to take it. We're in our last few minutes here together. I wanted to ask: Is there anything else you want to speak to in closing or anything we didn't get a chance to discuss that you want to share before we sign off?

Tim Hwang: No, I think this is great. We covered so much ground. It was really, really fun.

Moira Weigel: In that case, since we have three minutes left, I'll pull up one more question. [audience] asks: How can one engage in this talk, learn, and share resources? Are there forums or Discord channels or other places where you're sharing? Others in the group can't speak, sorry, the undemocratic Zoom webinar. Tim, are there other places you'll be publishing or sharing information about this work?

Tim Hwang: I'm going to continue to write about this. There's a number of other weird side articles—like b-roll—that I'm seeding out through op-eds. Twitter, such as it is, is a good place to find that. I'm just @timhwang.

Moira Weigel: Would you like to tell us about your ad campaign?

Tim Hwang: Oh, sure. This is maybe a final funny amusing anecdote to end on. FSG came to me—my publisher—and was like: “Normally we do paid ads. Do you want to buy paid ads to promote your book?” Of which, I don’t know—the whole point of the book is that it’s garbage and it doesn’t work and I don’t want you to waste your money. But they insisted. So a friend of mine, Helen Zhang, who’s an amazing designer, put together the *ugliest, ugliest* ads she could find. And it was very much an homage to an earlier generation of banner advertisements—like the “one weird tip,” or “linguists hate him” which is another one which I’m a big fan of. So they ran. They’re running on Facebook now. You might see them. *Intelligentsia* mag and also *n+1* was also forced to run it, thankfully. If you get their newsletter, you can also see a terribly animated GIF there.

Moira Weigel: Featuring a photo of you, if I remember correctly from my *n+1* newsletter. Thank you to everyone for joining us, and thank you to Tim for sharing your expertise today. We’ve been posting links in the chat window, you can buy Tim’s book—It’s a very attractive book, I’ll hold up one more time like a cornball—which is part of this new series of four tiles with Logic books out this week. And I hope that you can use the tags from the Q&A to keep the conversation going and before you leave, if you have a moment, please do fill out the short three-question survey. Thank you so much, Tim. Thank you, again, to everyone for joining us. And to CJ, Rigo, and Eli behind the scenes, for making it all look good.

Tim Hwang: Thanks everyone.