Shane Greenstein:

I'm Professor Shane Greenstein, and you're listening to the Harvard Business School Digital Initiative seminar, a premiere seminar series that hosts thinkers and scholars who are pushing forward research on the digital transformation of the economy. By conducting and connecting with cutting edge leaders, equipping leaders, and building community, the Digital Initiative seeks not just to study, but also to shape digital transformation. To learn more, check out digital.hbs.edu.

Shane Greenstein:

It's a great pleasure to have Hema from the University of Washington, whose last name I will not pronounce appropriately, so I'll let her pronounce it for you. As normal, it's good practice first to go around the room and introduce ourselves and our units. I'm Shane Greenstein, I come from the TOM unit and Digital Initiative.

Ava Ascarza:

I am Ava Ascarza. I'm from the Marketing unit.

Paul:

Hi, I'm Paul [inaudible 00:00:57], I'm also from the marketing unit.

Michael:

My name's Michael, I'm supporting marketing unit, topical student.

Speaker 5:

[inaudible 00:01:01] marketing unit.

Avy Goldfarb:

I'm Avy, marketing unit.

Natalia:

Natalia, strategy unit.

Dave:

Dave [inaudible 00:01:07], Digital initiative.

Diane:

Diane. Computer scientist, alumna of MIT and Harvard.

Fenju:

Fenju from [inaudible 00:01:27] unit.

Hang Yanari:

Hang yanari.

Ben:

Ben [inaudible 00:01:27].

Speaker 13:

[inaudible 00:01:27]

Chris Stanton:

Chris Stanton. Entrepreneurial management.

Julia Ordis:

Julia Ordis. 12.

Shane Greenstein:

Okay take it away.

Hema Yoganarasimhan:

Okay, thank you. I've never had such an interdisciplinary audience. So it's fun to come here and talk to you guys. And thank you, Shane, for the invitation. And today, what I'm going to be talking about is targeting and privacy in mobile advertising. And this is actually joined work with my student, Omed Rafian, who is basically the driving force behind this project, and he [inaudible 00:01:52] very good. Makes me feel bad because I was such a bad student to my advisors. But there's not much karma in this world. Anyway, so before I start this talk, let me get into the problems, let me give you a little bit of background about where this paper kind of comes in.

Hema Yoganarasimhan:

So as we all know, mobile valid option and usage has grown quite a bit in the last few years. There are two billion smartphone users worldwide and to put it in context, there are seven billion people worldwide. And the average user actually spends about 3.3 hours a day on phones and the bulk of this usage is actually coming not through, mobile browsers, but through apps, and now, actually, the internet usage through smartphones exceeds internet usage through desktops. So what of course, where the people are or the eyeballs are, that's where advertisers want to be, so all this has all has also seen this concurrent rise in mobile advertising dollars. And now, it's actually the largest share... it constitutes the largest share of digital ad spend. And greater than TV and so on.

Hema Yoganarasimhan:

So and this growth... this pretty rapid growth in mobile advertising is mostly attributed to an ad format which is kind of somewhat unique to mobile devices, and that is called, In app advertising. So what's in app... I have an example here because at some point in time, I used to be addicted to Angry Birds. Thankfully, no more. So this is the Angry Birds app as you can see, when you're playing this game. You might see a little ad, this little banner ad in the bottom, so that's really the in app ad that I'm talking about. And there are two reasons these ads have kind of become popular. The first is that it's a common ad monetization strategy.

Hema Yoganarasimhan:

So lets say I wrote a little app. I want to make a bit of money you could sell it. And that's why monetization strategy. But the more common monetization strategy is to show ads through this app.

And then depending on the pricing model, you make some money. And the second reason that these ads have become popular, this form of advertising has become popular.

Hema Yoganarasimhan:

Is that they actually offer very excellent user tracking and targeting properties. So what some of you might know or may not know. Actually I find that most people in the audience usually don't know is that all of our devices, mobile devices have a device ID and ad networks. Essentially platforms have a full access to this device ID. And they are able to track your device ID and track which apps you use, which ads got shown and and so they can in some way stitch together a user data across sessions, apps, and even ads unless you go and actively resect this ID.

Hema Yoganarasimhan:

So you know, so in some ways I have very persistent tracking information, which is not even available in desktop browsers. So it gives you this like from advertisers perspective or the ad networks perspective, it gives us like a huge ability to engage and use a tracking and do behavioral targeting.

Hema Yoganarasimhan:

So in some ways that is the good. So all these good tracking techniques, are good from... I'm using the word good mode in the sense of there they have a good ability to track them. I'm not making an automotive statement on whether it's good or bad. I have allowed advertisers and ad networks maybe they can to improve behavioral targeting, which of course if that it will lead to improve efficiency in the marketplace, we can make these better matches. [inaudible 00:05:28] on one. So that's on one hand. But on other hand it has also let all these privacy concerns among users on this ability to stitch together that information across apps, time and ads and this is part of his broader debate over consumer tracking and privacy.

Hema Yoganarasimhan:

So of course advertisers on one hand want for your protections, more ability to target consumers want higher privacy and limits on targeting. And there are people getting into this debate like regulators, governments and so on like GDPR regulation which kicked in about last year. They want to kind of regulate this targeting. And of course the phones and the ad industry also wants to not have regulation the one of the claims they make is that they can do self regulation and this question of like whether they can or they have incentives to engage in self regulation. In some ways it relates to this. Then you come to this broader question about this like revenue, from a theoretical perspective, right? Like this questions about revenue and efficiency trade off. So if you look at the... because all these platforms are using auctions fundamentally, right?

Hema Yoganarasimhan:

What we know in some ways from the auctions literature is that higher efficiency does not necessarily lead to higher revenues for these ad networks. So when you have fact, and I'll get into this in greater detail later, but the high level idea is that when you have said it factor distributions of valuations and you have 10 markets. So as I target more and more, that is one person who really, really wants you and is willing to pay a high price, but the next person's willingness to pay for the same impression is really low then all this.

Hema Yoganarasimhan:

So any contracting mechanism and efficient contracting mechanism is going to transfer that wrench to the advertiser rather than the ad network. Right? So this was kind of the conjecture that Lebanon Meldrum had. So they don't prove it, but they conjecture that, especially in online advertising markets when you engage in exercise targeting or narrow targeting as they call it, you can create thin markets and hurt platform revenues. But while this has been kind of discussing the theory literature, there's limited and product evidence that too much targeting can actually hurt platform revenues and so basically you want to think about what is the optimal level of targeting some the platform and if that is consistent with some of the users, preferences on privacy, then maybe at least in principle we can think about self regulation as a possibility. Yes.

Shane Greenstein:

Okay. So let me make sure I catch up there. All of you used to be... the problem was in the presence of very broad targeting like it used to be intelligent advertising.

Hema Yoganarasimhan:

Yes.

Shane Greenstein:

Yeah. There was an unending benefit to further division and targeting along dimensions of the audience. And the thought was that when we went into online, the being able to distinguish between ages and genders and even locations, even in history and so on. And there was no end to the benefits of that. And the point is, the Lebanon Milgrim point is there is an end to the benefit of that because if you get it down to one or two people, you've now... Right the intuition is.

Hema Yoganarasimhan:

Yes, yes, you know-

Shane Greenstein:

There's only one advertiser who really wants that.

Hema Yoganarasimhan:

Yeah. So you're basically, the more targeting you do, the one intuition is that you're increasing the heterogeneity in how much people value this impression. And as this heterogeneity increases, you're going to clear these large gaps.

Shane Greenstein:

Gaps.

Hema Yoganarasimhan:

And gap is basically the informational rank. You're creating value, but you're transferring this value to advertisers.

Shane Greenstein:

To advertisers, because they're just going to be one person who wants, one advertiser who wants that.

Hema Yoganarasimhan:

And they're giving the pain, the willingness to pay off the next guys.

Shane Greenstein:

[crosstalk 00:09:20] and what will be much lower, yeah, yeah I see. Okay.

Hema Yoganarasimhan:

Yeah. So that was kind of the conjecture. But it's kind of been hard to prove and frankly, so no one has looked at it because you need data from this platform, then you need to kind of show that what is the level of targeting where it actually breaks down. Right? Because if I don't build a strong enough targeting model, it's never going to kind of happen.

Hema Yoganarasimhan:

So that's kind of and you can see that in some ways it directly... it looks like this is one is a theoretical issue and the other is like a substantive issue on privacy, but they're like very closely interlinked. And so that's kind of what the paper is going to be about. Yes.

Speaker 13:

So there might be a different reasons consumers hate, but we have faculty to like tracking probably maybe like just like enate aversion to being like, "Oh," it's like currently targeting is just really dumb because I already like purchased something and all you do is to target me[crosstalk 00:10:22] exactly. So in that sense there's this kind of tenfold trade off in the sense that like in the future, like in 10 years time the targeting may be, most of the cases smarter and in some sense, but you get there without the dumb basis.

Hema Yoganarasimhan:

So you're talking about like the accuracy of our targeting model. So this is saying that, okay, even if I improve the accuracy and I made it better, it's not clear my revenues would improve. Is that clear? So it's possible that I have a dumb targeting model and that doesn't bring me enough revenue, that's understandable. But even if I improve, the [crosstalk 00:10:58]

Speaker 13:

Different from this.

Hema Yoganarasimhan:

[crosstalk 00:11:00] I want to be, you will see that I won't say anything about consumer welfare here because I'm not going to have any like targeting regimes which will change. And I'm not going to build a model of consumers. I mean we will be modeling them to the extent that we can... their behavior. I'll get to it in a moment, but I won't say much about consumer welfare because I don't have the variables to make that statement.

Ben:

This is raising another question. It could be offsetting in fact if people like well targeted ads, which Calgarian has argued they do. Then-

Hema Yoganarasimhan:

He works for Google. Yeah. [crosstalk 00:11:34] I totally believe that statement.

Ben:

You see it might be true if you give me an ad for something that I actually want and I learned about it that way. I could actually enjoy it. And so.

Hema Yoganarasimhan:

Yeah. So this comes through the question of, and we'll get to it hopefully in time. This question of when I think about contra-factual targeting. Right? So I'm going to be extrapolating from the randomization I have in the data. I think what the question that you're kind of getting at is that if people react adversely to more targeting or positively to more targeting. To what extent can we extrapolate from this targeting regime? Like very small randomization and then if it, so think of this as a short term contra-factual. Okay. So what's in such agenda? So basically you can think about what we're going to do in this. Oh, sorry, go ahead.

Avy Goldfarb:

I just want to make sure you understand the conjecture perspective- so that assumes there is no reservation price.

Hema Yoganarasimhan:

Good question. So this is with the prices. Of course If you impose optimal result price like the [inaudible 00:12:45] Estonian and approach, you can extract, everything. There are two things which kind of stop you from doing it practically. One is that it has to satisfy certain value a certain valuation distributions like the [inaudible 00:12:56] Estonian distribution to get the optimal reserve price. And most black platforms when you look at the advertiser valuations, including ours, doesn't satisfy that. The second is that here you'll see that optimal reserve prices have to be impression specific and no one has ever done that. Even Facebook doesn't do that. Like so I can impose some reserve price, like and improve things, but if I want to optimize at an impression level, that's going to be like, you're solving a different problem.

Hema Yoganarasimhan:

Yeah. Okay. So what's in such agenda? So the first question's about targeting and efficiency. So first they want to think about how can our networks develop targeting policies because that's what we want to then take to the counterfactuals, right? And then how can we evaluate the performance of these targeting policies, both from a factual and a counterfactual perspective. And the second is you want it, this relates to the substantive question of how do what's the value of different pieces of information. So if I think about [inaudible 00:13:53] contextual information, where is this user, like what kind of app is this person using? At what time of the day. I only knew that information. How much better can target them simply from an efficiency perspective compared to knowing who they are, which means I kind of know their behavioral history, like what they plagued on, on what they were shown in the past.

Hema Yoganarasimhan:

If you wanted to next quantify the value of different pieces of targeting information. And then finally you want to look at this question of revenue efficiency trade-off and characterize this empirical relationship between efficiency and platform revenues and think about what is the optimal level of targeting from the perspective of the ad network and the advertisers.

Hema Yoganarasimhan:

Okay. So because there are two sets of players here. And then relate this to the platform's incentives to preserve privacy. So that's kind of the high level goal and tell her you can think of this as kind of like in each part is how they like three main challenges. The first is that, when you think about like developing targeting policies, like I should be showing this ad and this impression, but when you look at the data, you won't have one ad which was actually shown, right?

Hema Yoganarasimhan:

So you actually need a framework which is able to get counterfactual, click through rate estimates, for the ads which were not shown in the data, right? And so to be able to develop any kind of efficient targeting policy, and I'll get to details of these, all of these in a few minutes. The second is you kind of need a framework which is accurate enough because if you don't have a highly accurate targeting framework, it's never going to kind of reach the boundary, right? Which is kind of happening in the real world. Otherwise if on my targeting framework as I just randomly draw and tell you like but it's going to get a click or not. It's not going to see that you're not going to see those behaviors.

Julia Ordis:

Just quickly a few question, what do you mean by value of each price of information? Like how much are they generating if you can explain or value to the?

Hema Yoganarasimhan:

Good question. So value here I'm going to be like how much improvement in efficiency I can get?

Julia Ordis:

How much you're mailed that this person might like that?

Hema Yoganarasimhan:

Yeah this impression and then the actual economic value will come in the next part. Yeah.

Julia Ordis:

Oh yeah. So this one is the discovery it had the originating in the [crosstalk 00:16:11] in amount of situations.

Hema Yoganarasimhan:

Yeah. To what extent can I explain the click [crosstalk 00:16:16] with behavioral information.

Julia Ordis:

Okay thanks.

Hema Yoganarasimhan:

And finally. So this is where like you have to just do the first, this part for example. The second part I could just use like a mission learning framework, which will do this for me. Right? But then you also want to be able... you kind of need an economic model of some kind of strategic introductions or options in this case, if you want to take these targeting policies and outcomes to some kind of revenue, right. If you want to talk about incentives, you need theory at that point. So we're going to... you need some kind of economic model to be able to do this, to kind of link these targeting policies to incentives.

Hema Yoganarasimhan:

So that's kind of the three things that we need to take care of. And here's a broad approach that I'm going to take. So the first is we're going to propose this I think procedure which is going to take into account. They're going to think about identifying ads. Think of it as like a nonzero propensity of being shown. Because in most it if you don't have a deterministic option, but most online advertising is deterministic. You can imagine that if auction is deterministic, there's only one ad which has a non zero probability of being shown. So, you actually need a setting where there is some randomization therefore. Then you're going to burn the machine learning models, click-through production, which has sufficiently high predictive accuracy.

Hema Yoganarasimhan:

And the core things which are going to happen here that's like, generate features which are less about like the behavioral information and the contextual information and then use those to measure the value of those different pieces of information.

Hema Yoganarasimhan:

And then finally we're going to burn my economic model of auctions to determine, market outcomes and the different targeting policies. And then we will use that. And the results from here. And combine it with this to categorize advertisers utilities and the platforms... the platforms that have venues and advertisers surplus and the different conquer factor targeting regimes. Empirically look at this problem of like what is the optimal level of targeting. And one of the goals here is that in some ways if that means we want to be able to coherently combine some of these like productive machine learning methods.

Hema Yoganarasimhan:

Which are very good for like, let's say targeting, but also like economic models. Which are a little more descriptive and tell us a little bit about how we can take them and say something about, play us incentives.

Hema Yoganarasimhan:

Okay. So now let me talk about the [inaudible 00:18:47] literature [inaudible 00:18:48] There are many, obviously strings of literature this relates to. At a high level, you can think about this as there is a literature and computer science on click through rate escalation and targeting are the bunch of methodological papers on that, not just for targeting, for anything like production. And there are couple of ways in which this works different from that because those are all for factual predictions. They never think about contra-factual predictions. And you'll see how one of the goals here is to see how you can bring in some methods from there and some ideas from the [inaudible 00:19:21] literature from the economics.

Hema Yoganarasimhan:

And then combine those two together. And there's also the paper of course relates to this very large literature on targeting and privacy and, both on that set of privacy regulation on how users respond to different, targeting. I think Avi Goldfarb and [inaudible 00:19:40] have some nice papers in that area and finally relates to this, more theoretical literature on revenue efficiency trade off, since the, [inaudible 00:19:52] conjecture, there are a lot of theory papers which have looked at it and have shown that this can happen in theory. Of course they impose some functional form assumptions on the valuation distributions and you'll see that what we try to do here, we try not to impose any such thing. And there are also a couple of empirical papers which kind of briefly look at it.

Hema Yoganarasimhan:

Like the ATN like Apollo have a case study where they tried to do this and Yow and Nailah but they show that actually more targeting house partly because they don't have like a very highly accurate machine learning model. They never get to the extremes. So that's kind of the broader literature where this paper should be situated. Now let me get to the setting and data. Any questions? Okay. So the data actually comes from major in app platform of an Asian country. It's not China. I know when I say a large Asian country, people always assume it's China. I'm like, no, it's not China or India.

Julia Ordis:

But is a large.

Hema Yoganarasimhan:

It's large, but not the largest. I never said the largest.

Speaker 16:

[inaudible 00:20:52] relative.

Hema Yoganarasimhan:

It's a relative terms. Yeah so this country doesn't have actually, Apple's a product sold. So they only have android devices and this in app advertising platform at that point had about 85% market share. In the country. So it was like the largest in app mobile app advertising platform. And so it's a pretty that percentage of data for the country. And it was also interesting in the sense they only had one format of ad, which means that, now of course there are many different formats. So now back then it was just like small little banner that you were showing the [inaudible 00:21:48] and they also had some interesting things like they had very limited targeting in the sense that advertisers could only target on broad categories like app category, like sports news and so on and, or the problems like the stage and, connectivity time, wifi or data.

Hema Yoganarasimhan:

And they also didn't engage in any kind of behavioral targeting. So the ad network itself was not doing any kind of like extreme like targeting. And the most interesting thing about the platform is the user auction mechanism, which is not very common. And it gives us a lot of power actually. It's called aquatic proportional auction. So it's not a deterministic auction and almost all auctions that [inaudible 00:22:33] you'll see that the allocation is deterministic. So what these guys did was, so this B is the bid and Q is the quality score, which is not updated regularly. So it's not like, every after every impression your quality score is getting updated. It's kind of a constant. And so if you are usually in a deterministic auction, if you have the highest bid into the quality score, you win for sure.

Hema Yoganarasimhan:

And then you pay the next person's like expected price, right? And then, but you don't always, even if you have the highest bid into the quality score, you actually have a probability of winning. So as you can see, this basically creates a situation where, even if I am in principle, so first of all, this key was not very customized. So the match is already kind of poor. But even conditional on that match, if I were going to be kind of going to still get it, it might go to the person with the lowest willingness to pay on the lowest match. So this is kind of like an experimental variation with some ways, right? Because this conditional on a set of advertisers participating in the auction then who it gets allocated to is in some ways like randomized.

Hema Yoganarasimhan:

Good question. That's a very good question. That was the first question they asked them. So there are two reasons. The first is that so such actually, the theoretical reason, is that actually, these types of auctions have been shown to have better what is that? Maximum properties. There are actually some auction theory papers which show that they perform well under the best worst case scenario. So there is that, but that was not totally their thinking. Their thinking was more along the lines, I think about this, you are on the mobile device, right? And you are spending 10 minutes in this app and they thought that I, for the 10 minutes, if it's a deterministic option running, I'm going to see the same ad for 10 minutes. Right.

Hema Yoganarasimhan:

Because it will, the quality score is not changing. Your bid is not changing. I'm going to see the same idea. So they thought people are just going to get really annoyed. So they thought it'd be a good idea to kind of hack some randomization into the allocation mechanism. So that was, that was kind of their thinking. But after we dug into this, we saw that, it has certain properties which might be useful, but that was not their thing. Yes.

Shane Greenstein:

It's, it's almost as if they lacked confidence in their own ability to allocate. I don't think you would do this if you thought your own auctions were behaving well.

Hema Yoganarasimhan:

Yeah. So if you were updating these qualities for them in real time, like after every impression, and you actually look at this and another paper, but you look at variety and so on, then you could optimally allocate it within the session. But then I'm not doing that. So they didn't want to show the 10 exact ads for you for 10 minutes. So this was like a frock say for them to kind of create some randomness. Yeah. So, that was the idea. Anyway so this probabilistic allocation mechanism, obviously it's not very good from a revenue or an efficiency or any perspective, but it think of it as you're basically learning like the scores are experiments, which I have no problem with.

Hema Yoganarasimhan:

So that was the background in the setting and the data. So the data comes from about it's a very large data set but the sample you're using for support for a one month period and the bunch of young adults in the day, he told you his four to know first we see the time stamp and the date. And the app ID, the [inaudible 00:26:17] impression happened for each impression that is, the device ID which is really your user ID because no one it seems like from what we see, not many people really reset their device ideas. I worked on this for years now. I don't really reset it. So it's quite persistent. There is ad ID. There's a targeting variables like which advertisers but kind of engaging in what types of, and that's very minimal we see and the location and finally a quick indicator whether it's like happened or not.

Hema Yoganarasimhan:

So you're going to split the data and sample it. Because obviously it's very large to work with and you also, you'll see need to generate features based on the history. They only think about, it's like all kinds of prediction kind of settings. You have a cold start problem. I can't start predicting anything for people here because I don't have any history to make any features. So what you're going to do is take the last three days of the dataset, for the actual impressions for making the predictions. But we are going to use what I call this lower data to actually come up with the features for these. Okay. And the sampling procedure is that we actually sampled or 100000 users. And because I cared about behavioral history, I'm going to sample users and take their entire history because I don't want gaps in someone's history.

Hema Yoganarasimhan:

If you didn't care about behavioral, that you could actually sample over impressions or they both would be the same. But here it's not. And after about 700000 we didn't see any improvements in any model, but an over 700000 users over these three data set or three days of training on tests and validation gives us about 27 million impressions for training and testing. But we also used about one 46 million impressions here to generate all these features, which will be used for the prediction purpose. So this is the data splitting and sampling. Okay. So let's start with this machine learning framework for targeting. And let me start with the problem definition here. So if you, so I have this matrix because it's the easiest way for me to think about this.

Hema Yoganarasimhan:

So think about an ad network or a platform which has let's say ad impressions, right? So these are all impressions and capital A ads I didn't leave. What time to predict is basically for each impression ad combination, what's the probability a click would happen. Okay, so that's why IA work once the click will happen. So we want to think of, we call this a match value matrix in some ways, right? So if I could know each element of this matrix, I can come up with any kind of targeting policy and predict what would be kind of the gains from that. So our goal is to accurately estimate this is match value matrix. And then once I have the match value matrix, then I can use different targeting policies that math impressions to ads.

Hema Yoganarasimhan:

Okay. So what are the challenges here? The first is the counterfactual click through rate estimation because it's important to recognize that. Then I look at what I have to do this for each row. I have one only one variable which was shown maybe for impression one ad two was shown and I see an outcome of zero on one for that and I don't see anything for the rest is that clear? So I want to be able to like you know, so I can maybe learn to predict factual outcome better if I want to be able to predict counterfactual outcomes, I have to be a little bit more careful. And I have all these like high dimensional categorical inputs, you'll see user add app and so on. And finally I want this framework to be accurate enough. Yes.

Shane Greenstein:

So is this app independent in the sense that.

Hema Yoganarasimhan:

The ad network you mean.

Shane Greenstein:

Yeah. So it doesn't matter what app you're using or how you've switched between or anything of that?

Hema Yoganarasimhan:

No, because they have visibility to all of them.

Shane Greenstein:

Yeah. But it also, you also aren't going to some sense of a model what the user's doing as a way of understanding why ads get clicked. And you understand why that might be a problem. Right. Because you can have someone who's going from a game to a map.

Hema Yoganarasimhan:

And those should get captured in the features.

Shane Greenstein:

You'll have some features.

Hema Yoganarasimhan:

So we will use features for that, we'll say. Okay. How many times do you added this app? Was this like what are the different apps you have seen in the past and what are the ads that you clicked in? Which apps.

Shane Greenstein:

Well get some of it? Yeah. Okay.

Hema Yoganarasimhan:

So just think of this as an outcome prediction problem at this point. Yes.

Speaker 13:

So from the apps point of view, especially within games. Is advertisement for other games that maybe like a competitive app, so do like game apps have to reject certain kind of advertisements?

Hema Yoganarasimhan:

No, I don't think so. Once you work with the platform, you don't have any ability. You don't, you actually don't even own that space anymore. They allocate that.

Ben:

And the, and the ads that you're tracking, like it's the same ad being shown on the screen at a time.

Hema Yoganarasimhan:

The same set of ads. We work with about top 37 ads, which constitute about 80% of the traffic and the rest are kind of pulled together in one bin. So it's the same.

Ben:

And those advertisers aren't using like nudging techniques where like their ad gets progressively more interesting.

Hema Yoganarasimhan:

So the thing is they can't do that level of targeting, right? Because they give this one banner to that platform and they actually, because they are not the ones serving the ad, there's no behavioral targeting. They don't know it's you. So they don't have any data that you might have been shown this ad. So they can't actually say you've seen this ad once to take it down. You're going to see those different creative because they don't, neither the platform nor the ad network is using any kind of.

Ben:

Because that seems to be where targeting then does have go further. You can get at that then targeting you could extend further than where you are.

Hema Yoganarasimhan:

Yeah. So it's possible that the ad network could do much more to improve targeting. But one thing you will see is that one of the lesson says like maybe what they, even what they could do with this is already too much. Yes.

Natalia:

So if your quality score is also not dependent on the app.

Hema Yoganarasimhan:

No, the quality score is ad specific and kind of fact like they change it every month or so. That's yeah, exactly. Yeah. Yes. It's exactly the same. So their quality score is very primitive. Yeah. I wouldn't tell them that quality score is constant.

Natalia:

So is there any targeting there?

Hema Yoganarasimhan:

Yeah. So the advertisers can choose that. So I can say I don't want to be shown a news apps.

Natalia:

But then it's kind of global.

Hema Yoganarasimhan:

Yes. You are never shown in a news app. It's like you said, it's like equivalent to setting you a bit to zero for a certain category or a certain like state. Yeah. So the thing, cause it has bids kind of, but we see that even they, even though they can like make different campaigns for let's say news apps was a different like a sports app, a sports category, they don't change their bids. Yeah.

Diane:

I'm curious if, your framework, the apps have an option to specify whether they want to attract new users or like sort of scale up versus [crosstalk 00:34:01].

Hema Yoganarasimhan:

So the app store, you mean the advertisers because apps already are the ones who have with the impressions. These are like people playing angry birds.

Diane:

I guess in as an input into, the design of the targeting campaigns.

Hema Yoganarasimhan:

But the apps are not the ones targeting I do you mean ads? Oh, the ads, no. So the ads don't have no. So the platform doesn't necessarily allow the ads to say this is a goal and like optimize on that for me. So the ads just have like usually the targeting variables I showed you. Okay, do I want to advertise in these categories or these areas or not? And what are I want to be, what bid I want to bid for this area. So it's very basic in some ways. That's why it's helpful in the sense it creates this very poor mapping between impressions and ad. So there's a lot of randomization which is being generated. Obviously is it efficient for the, is it good for the platform? How obviously not.

Hema Yoganarasimhan:

So, okay, so this, this is basically what we want to do. So let me talk briefly about this counterfactual problem. So I've got counterfactual click through rate estimation. So but as the accurate, when can I be sure that this is that ritual shown in my, in this introduction I want to print it, the click through rate of a different that if it had been shown in the impression, right? So even if I want you to do that so people think like if you use the mission learning model, then you can just like do magic and anything would happen. But that's obviously not true because there are some strong assumptions which we are making. The idea is that if I have an impression, if I have an app, which is, are never if I never showed this ad to let's say women and I only showed this act to men, right, no model mission learning or magic or anything, but I won't help me predict how men would react if they see this ad.

Hema Yoganarasimhan:

Right. There's simply not that in the data. So intuitively, so that's intuitively like formerly what this means is that as long as the joint distribution of covariates and outcomes in my training data is the same as what's in my test data. I can print it, right? I can make concrete factual predictions and that's why it happened. And these estimates that accurate, as long as the ad could have been shown as long as there is some randomization my model sees instances that this could have happened in the data, if it doesn't see sufficient enough instances that I'm never going to be able to do it. So what does an acquirement who need some randomization and randomization here coming from intuitively you know, there's very little targeting which is happening, right? Specifically that was more behavioral targeting, which would end up being shown in a broad set of apps, users and settings.

Hema Yoganarasimhan:

And of course the probabilistic allocation. We will naturally create randomization within the set of ads. And let me show you, find you some evidence from the data because I could say that, Oh, it's probabilistic, but it could be like 99%. It gets allocated to one and one percent it gets to allocated to another. So like when you show you from any pressure from across our test data, so this one is how all the exposure number, I mean like within a session impression number like you know, so this was, it has been around for like six impressions with enough, I'm not this person, this like set of people.

Hema Yoganarasimhan:

And you can see that the median, a number of unique ads that they have seen is about three to four, right? So there is a lot of, what that means is that over 50% of the people I'd see for more than four, four or more add unique ads within like five impressions. Right. Which means that there is lot of randomization being created and they see kind of the same pattern. So there's probably the stick allocation rule is creating a lot of exogenous variation in some ways. Yeah.

Shane Greenstein:

That's gotta be something also about the [inaudible 00:37:56] don't kill you. So it's got the greatest gotta be some minimal requirements here. The pool of set of ads that are available can be too big and the combinations of men age, education, location, whatever it is you're conditioning on, can't be too big so that you don't have empty cells. Right. That has to be something like that here.

Hema Yoganarasimhan:

So that's a good question. So what, let me maybe just tell you what other strategy is. So that's good. Okay. So we call it as a first firstling procedure, which means for every impression do you want to identify, so impression already takes into account things like the user, the province and everything, right? We want to identify the set of asphalt which our conquer factual estimates would be accurate, right? Consistent, I want the accurate. Accurate also means that like other things are consistent, right? So, you want to identify the availability of an ad for an impression and that can happen through two things. Of course there's the probabilistic which creates the variation, but what else? The targeting decision, maybe certain acts are never going to be shown in this app.

Hema Yoganarasimhan:

So you cannot predict for them at this time of the day they never did all the other thing, which often happens as campaign but be like this ad their budget is always exhausted by 9:00 PM and they never show like so things like that. So we actually have data on this. So it's kind of a data driven approach. We know for what types of users and bitch ads are available for each impression, it's like literally an impression level filtering metrics and then you're going to filter out ads and buy filtering out ads for each impression, we can still print it and we do that. But when you come to dialysis, we won't use it in the like point of analysis because I can't be sure about that accuracy. So, so to basically, so this is kind of this intuition of like non-zero propensity score, right? So that basically what it is.

Hema Yoganarasimhan:

And we are going to the with what we call this filter sample of ads for each impression and it could vary from impression to impression its important to kid of keep in mind. So think about this as the way we actually do this in the paper is if you go to this you will make exactly as some of the metrics like this with zero and ones. And the ones are the ads for which contra-factuals are valid, zeros are the ads for which, and then you can do like a dot product, yes.

Speaker 13:

I'm just curious about which ads in terms of frequency of use, like I'm using an app and this when I use something I really like they actually start showing you more apps and it makes me wonder like, uninstall? Its like [inaudible 00:40:44] but then they would be making a negative thing so I'm just wondering if they determine by usage.

Hema Yoganarasimhan:

But that's seems like an app related question on like actually, if I show more ads, like am I just going to like disappear from the app? So here you can see that and I was always going to be shown. So I can't really say much about more words. There's less ads. But if the same ad is being shown, is that your question?

Speaker 13:

Well, a little bit of both, seven minutes has the same ad but it may be different a, so they're doing a job frequency. Maybe they just need more money. So they could just do more ad impressions.

Hema Yoganarasimhan:

That's a good question. Unfortunately it's not a question I can answer, because in my case it's all, that is an ad, that frequency and that time of the ad being shown is constant through the data. If I could have variation in all their 15 second talks, then you go for a minute without showing an ad, then that would get to that question.

Hema Yoganarasimhan:

Okay. So again, once you do the photo sample one question you might wonder is, okay like is there like so much geo targeting going on in the system and like other things like campaign like availability, how many ads end up like competing like in this auction. So I just wanted to show you like the photo sample. Like how many you can see, like pretty much all, impressions have at least like above like 10, ads competing and the median impression has about like at least like what would that number be like? Yeah. Yeah 12 so that is sufficient out of the 37 ads that we have, right? So at least 10 to 12 I had even in most cases, and there are of course some which have like way more, impressions, they're more ads competing for them. So, so began actually then you go to the country factual, that is a sufficient number of competition still. Like we can make some statements.

Natalia:

Question about the availability of an ad for an impression. So at the end of the day an impression for you is going to be a bunch of features.

Hema Yoganarasimhan:

So impression. I was going to have it to the next slide. I thought people were living by one. I just took it out to take a minute. Sorry, literally in the next slide. I can still see it's okay. Maybe I tell you when we meet cause it's so funny. Okay, now I can't make it.

Natalia:

You don't want to tell anyone.

Hema Yoganarasimhan:

Let me see. I don't know why I became so tiny in this corner. Oh my God, I did. Okay. I did put it back. So that, that's related to the future generations framework. Kind of like how he was. So in that number basically acts in the prediction problem. So we have all these like, so each impression as you are saying is categorized by what like the user creating the impression.

Hema Yoganarasimhan:

So what do we have as like inputs, right? The app where the impression happens, the ad which was shown and the time of day, we kind of put it in 24 hour buckets and we wanted him.So this is what you're going to use. So if someone chooses not to target on an app, they want to go up. So if a user is in a certain location, which they're not targeting on that for that impression, that ad is not going to be available.

Natalia:

So then you're prediction is going to be by user, app, time which ad should add.

Hema Yoganarasimhan:

We will first credit using the ad and then you would come up with a targeting policy. Let me get to that in a moment. So now think of this as just like a factual prediction system, right? Because what we is actually, you can take it to those counterfactual ads.

Hema Yoganarasimhan:

So let me think about, it's like, of course these are all very high dementia. Like you know, there are like millions of users, millions of apps. So you want to kind of like come up with some sensible way to contract this so that it's less information in there. So this is something I have done in all of my previous papers. We discovered like assumption based feature generation framework and that's kind of what we are going to do here. And the idea is that for each impression you can user app or a time and add his input. And instead of just using these like kind of a like a common metrics, you can come up with functions which compressed this information. And I, so for example, if you have an impressions function which were take these as inputs, I tell you for user, I for a given app for the given time of the day, how many, how many times have they seen this ad in app?

Hema Yoganarasimhan:

And you can and then you can like do it at any level. So the impressions function, I think they are like actually about 30 of them. I'd like different inputs. I can if I choose not to specify a user each over all users maybe. Right. So that's kind of the intuition. So let me tell you how we actually do this. So let me do you think about this? Is that like the same function? How many impressions are you a user or that the inflections function for example can capture different types of inflammation but just contextual behavioral or Adelaide. So the contextual information that for example, we're trying to capture this, which app are you in gaming up? Or like maybe a case your behavior might be different, right? Your likelihood of clicking. So that's kind of what into your delay trying to capture or time of the day of like work.

Hema Yoganarasimhan:

Was this like leisure or on the behavioral like if the there is the feature has information or what your past history, then it's telling me something about your past behavior which might help me predict what you will do now. Right? So just like these were like telling me something at the global level about what people at 10:00 PM are they more likely to play in this app versus people at 10:00 AM in a different app. And then finally you want some many, many of these features could also be at the ad level, which means that if I show you and ad. If I show you, let's say a sports ad in a sports app, maybe you are more likely to click or people who have clicked on these type of ads are more likely to click on it.

Hema Yoganarasimhan:

So essentially you're trying to capture all those information using functions which are going to map all these user level, app level, ad level time level information into one variable. And you have, we also do it over time. So some of these are aggregated over long term or one month, just one week and some within the session. So if you already clicked on something in the session, maybe your behavior is different. So we have about one 60 features. We are using this I, this is actually kind of the part where you actually are trying to extract all this information. This is where much of the work on the machine learning part happens. And you can see that the features is like overlapping, which means they could be features which are both contextual and behavioral features, which are contextual behavioral. I specific you don't have all types of information. So, so this is basically the set of features we are going to use. And think of this as the X variable you will give into some kind of prediction algorithm.

Shane Greenstein:

Did you try, did you do any analysis of whether the number of features matters in 160 seems small to me, but maybe based on your data. That's good.

Hema Yoganarasimhan:

Yeah. So this is something else I've done in a previous paper. So this set of they have, so you can actually show that the gain information gain, but you can do this analysis because these rapper models, where you're like throw one in. Like how do you build a model and you can actually show that the information gained compared to like the runtime of the algorithm basically tapers of even with much less features. Right? Yeah. Good question.

Hema Yoganarasimhan:

Very few times I get to say I have done those. Usually it's like, I hope someone else had done it. Okay. So now think of this as X variables you have for some prediction. So, before you go use any algorithm, you need an objective. What are you trying to maximize? So what we think, think of this year as a log loss, which is our objective function and this is the most commonly use low gloss because we have a binary outcome variable here and we are going to use a validation data to tune a bunch of hyper-parameters and the learning algorithm that you're going to use this called [inaudible 00:49:12]. It was interesting. You actually developed at UDaB, by our computer science people and it's a fast and scalable version of boosted regression trees. And if you go to pretty much any capital competition, it's a vending algorithm and most of them are in one of the samples.

Hema Yoganarasimhan:

And of course, no one ever believes me when I say this is the right thing to do because ended up yours, made us do lots of checks to see, Oh, why not use a random forest or why not do something else? No, none of them work as well. If you are ever in a production where you're thinking of tabular data, not like pictures or texts, this is the best algorithm to use.

Shane Greenstein:

Is it better to use more than one?

Hema Yoganarasimhan:

This is very powerful. It doesn't give much. You don't get much or think about it as it's in front areas of the data space. Maybe one model works better than others. That's the whole idea behind ensemble learning. But this genuinely works so bad and scales so well also that even if you were to gain a small amount, the scalability of those other models is very good. Yeah.

Hema Yoganarasimhan:

So not in one to evaluate this model. Okay. So there are two ways to evaluate this model. The first is RIG, which is, what I call relative information game. And the idea is that. So think about this, this is a model where [inaudible 00:50:39] happened very rarely, less than one person. Even if I tell you like namely it could never happen, right? I'm mostly right. So you want to kind of have information you wanted to see how much better you can do then like a knife reduction. So that's like going to inflammation game is very common to used and the idea is that, this is the average prediction. So the data compared to now if you use your model to predict how much better can you do, and the first modern evaluation, you're are going to actually think of this as factual model evaluation.

Hema Yoganarasimhan:

So the red element here, I think about pick of these. I want the ads which would actually shown in the day on a test data, right? And how well can I predict for ads, which would actually show up in the data and RIG in some ways allows us to, because RIG has other properties like based on Loblaws, which was our the function be used for optimization and so on. So I can evaluate the model in this way. The second way I can evaluate the model is to look at the potential improvement in click through rate. And this is based on this, targeting policy, right? So let's say that this was the app which was actually shown in the data, right? And now suppose that I come up with the, actually shouldn't target the policy. By efficiency, I mean I allocate the ad to the, sorry the impression. No. I allocate the average has the highest click through rate for the same ration to it. Okay. So I'm going to maximize the, yes. So this was not seeing my data. So this is kind of like a co conquer factual evaluation. Okay. So you can see what's the improvement in click through rate that you would see if you changed your allocation policy. Okay. So this is two separate ways to evaluated.

Shane Greenstein:

So you're sort of, I understand correctly, you're sort of mixing up two things. So one is sort of better targeting, but if you're suggesting that they show that one, you're also assuming they get rid of this rather naive random. Can you compare it to what if they just took the one with the highest probability of going to their algorithm and see how well that did. If they, a lot of that with certain things.

Hema Yoganarasimhan:

So they don't have an algorithm now, right? They're just doing a random algorithm. Randomized.

Shane Greenstein:

Oh, you don't have the probability.

Hema Yoganarasimhan:

So I do have the probability that's the baseline.

Shane Greenstein:

Okay, good multiplied by the qualities.

Natalia:

The probability if being shown.

Shane Greenstein:

Because they actually showed one, but that was a random draw. What if they showed it to whichever one had the highest quality score?

Hema Yoganarasimhan:

Oh yeah. I could easily yeah, we haven't done that. That's easily doable. We can show it to the bid into the quality scores.

Shane Greenstein:

Doesn't that get you a large part of the way there? This randomization could just be.

Hema Yoganarasimhan:

Could just be so bad. Okay. That's a good idea. So we could do that. That's a good point. Yes.

Shane Greenstein:

So you don't do any conditioning on a previous click through. Right. So and so in other words, yeah, for the user. So if you cause this approach, every click is independent of every other click, click through is not so, so in a session if a user clicks through an ad, you would think the probability of them clicking through on the next ad is low or you might think it's very high.

Hema Yoganarasimhan:

So if you update the features in a real time.

Shane Greenstein:

No, but you'd want a serial correlation of behavior. That's the issue.

Hema Yoganarasimhan:

So if you, as long as it's captured in the data, right? Like in your training data, if you had people who clicked on an ad once and you see weather, what's the likelihood they're going to click on the ad next to that extent you will be able to predict. Also in a test setting. To what extent people who have already clicked once, people are already clicked twice and so on will click again.

Shane Greenstein:

But you wouldn't worry about it. About Sera correlation in clicking is either telling you there's an unobservable propensity to click, which is in some sense not being measured.

Hema Yoganarasimhan:

But that's in the user level barrier.

Shane Greenstein:

Yeah. At the user level. Yeah. Or the other opposite that they are only going to do one click per session in the fact that they did then predict they'll never do it.

Hema Yoganarasimhan:

Yeah, but that should be in the historical data. Right? So you have their past history session, so this is like an outcome prediction. So you should be able to see, Oh this person, whenever they clicked in the app, they don't come back and click

Shane Greenstein:

So you don't talk. You see why I'm asking. I understand what you're saying that you think the conditioning variables are going to be good enough predictors.

Hema Yoganarasimhan:

Yeah. So as long as you are able to predict the outcome and that outcome prediction is correct and it would be as long as the randomization was sufficient then this, you should see, you should see this in the data.

Shane Greenstein:

So as long as the randomization is sufficient.

Hema Yoganarasimhan:

And you should have seen instances like this in the data. Of course you're are going to shrink it across all these variables too much. You're not going to have seen it enough to be able to predict it with any accuracy.

Shane Greenstein:

But yeah. You still understand my point that if you thought people behaved a certain way, only one click per session and so on, you might not pick it up with this kind of [inaudible 00:55:36].

Hema Yoganarasimhan:

No, we have session level variables, right? We are, we have a bunch of features which are updated real time in the session. Yes. So that will take care of it. Oh, sorry, I didn't clarify. Yeah. Okay. So, okay. So let me talk about the the evaluation of these models. First on our ideas, just factual evaluation. I'm here we run three different models. One is we have the only contextual information and then using of course all the ap specific features, [inaudible 00:56:05] specific and the full models, like everything. And you can see that a behavioral model gives a much higher, improvement, which means just using behavioral information as much better than using contextual information. And of course using both is even better.

And, so this tells us, at least from an efficiency perspective, behavioral information is more valuable than contextual information.

Hema Yoganarasimhan:

Right? And the full model of course, is even better. So from a privacy perspective, even though behavioral information is bad, maybe from user's perspective, it certainly increases matches, in the system. So that's good. I can also share, I've also want you to briefly show you factual click through rate improvement. And this is, we find that there is about about one point the click through rate goes from our 0.6 to 1.9. if you use an efficient targeting policy, but you're right, that maybe a better comparison would be like if they had used a deterministic allocation policy, but their quality scores.

Hema Yoganarasimhan:

So you'll see a pretty big improvement in efficiency both from a factual and a counterfactual perspective. So, so that's, this was sparked by, it'd be we didn't talk anything about revenues and anything. You're always like, Oh yeah, like can you actually make better matches? And now I'm going to talk about the revenue efficiency trade off. So, to kind of this is the Lebanon middle ground conjecture in a picture. So, hope to convey the ideas. So let's say that I have to use this and to advertise this to impressions and ad network with two impressions. And did I do something?

Natalia:

No I think its just [inaudible 00:57:47] no?

Hema Yoganarasimhan:

Lets see, that usually works.

Hema Yoganarasimhan:

Okay. That's an easy enough solution. Okay. So let's say there's an ad network with to users and to advertisers. And what I'm showing you here is, so let's say here's this guy who was, I can actually use the clicker. That's cool. Okay. So no, is that.

Hema Yoganarasimhan:

I don't think it's working. Yeah, I think I'll use the old fashioned way. Right. So let's say that to two users. One is, interested in traveling, one's interested in, let's say health, and if they're two advertisers, let's say an airline company at a gym. Now you can see the I for the airline company, this user is more valuable. And for the health company, there's other user that's more valuable. So these are pros and valuations. But if you use any kind of efficient auction allocation mechanism, like he had actually for a second price auction and you know, this was advertise, it was going to win this impression, but the platform where they're going to pay this price. So the entire park is going to be the advertisers surplus, right?

Hema Yoganarasimhan:

So the platform is only gaining so much, but now think about if they bundle these two impressions and it didn't reveal to them who was who, right? The advertisers can distinguish them. Obviously this sound by who the airline is has generally a higher valuation and expectation. So they want you to build both impressions. So they have, the allegation is inefficient as you can intuitively see. But the platform's going to make more money, because the informational rank that is being given away to the advertisers is now lower.

Shane Greenstein:

It's the same impression. This is the same intuition from, I'm not, I'm not playing favorites in an auction. Yeah.

Hema Yoganarasimhan:

Yeah. And this is, also relates to this target in some ways. Yeah. So, in some ways so this is kind of the revenue efficiency trade off you increased you, in this case you had higher efficiency but lower revenue for the platform and other cases, other ways [inaudible 00:59:59].

Shane Greenstein:

Yeah. The intuition for this is from government procurement. When you have a near monopolist and you're actually seeing your interest in inviting a competitor by giving a disadvantage to the monopoly.

Hema Yoganarasimhan:

Yes. Yeah. I've seen some of those papers. Yeah.

Shane Greenstein:

Which then induces a competitive response from the right. That's the same.

Hema Yoganarasimhan:

It's a similar thing except that information disclosure.

Shane Greenstein:

Yeah. This information here, what is going on here is you want to induce the competitive response from the advertiser.

Hema Yoganarasimhan:

Yes. Exactly. And here you're doing it through information disclosure and that you do it by like having higher levels of competition here, more information. It's playing a anti competitive. So, that's basically the intuition and what you want to understand is what's the optimal level of targeting from the platform's perspective. So now you'll have seen a similar matrix before, but now think about this as a valuation matrix, right? Where this is every time he says evaluate value for each of these impressions. I normally make any functional form assumptions on this, but even without that you can actually be conceptualized. Targeting has you know, bundles of information like Mundos you can give you, so within the bundle, I can't tell who is who, but I can tell the average for this bundle, let's say. Okay. So, and then you can define that elect of granularity of targeting levels.

Hema Yoganarasimhan:

When you say that our targeting strategy is more granular than another, if you can, if you can distinguish two impressions within one targeting strategy, they also remain distinguishable and the other targeting strategy. So think of it as like if I do app time level targeting, it's more granular app level targeting, right?

Hema Yoganarasimhan:

So that's kind of the intuition. Then you can actually show for the second price auction. But again, for because of revenue equivalence principle, you can show it for pretty much any efficient auction mechanism. It satisfies those property. You can actually prove which we do in the paper that uh, without imposing any functional assumptions, you can prove that the total selfless efficiency increases as the granularity of targeting increases. But platform revenues, you can't actually sign them. They can go in either direction. So, so that's really an ethical question then. Right?

Hema Yoganarasimhan:

And then the consider for targeting scenarios in the data which relate to the models and substantive questions. We had no targeted contextual targeting, behavioral targeting and for targeting and get it because of this theory results we know that for targeting is going to generate the highest surplus on welfare [inaudible 01:02:33]. But what contextual and behavioral will be somewhere in the middle and I can't even literally find them in surplus because they are orthogonal pieces of information. Okay. So there's not like one is subsumed by the other and but glass on revenues, there is not much that theory can offer here. So you had a question.

Paul:

You talked about how you back out.

Hema Yoganarasimhan:

I will talk about it in a moment. That's the next slide. And this time I didn't do anything back to it. Okay. So okay, so what is the, so how can you do the center equally?

Hema Yoganarasimhan:

So this is evaluation matrix. So now if I want to think about concrete factual targeting strategies and then impose my theory model on that, I need to have known each element of this right? So this is basically the most evaluation or that, but what's the advertising value for a given impression? So you can think of this as the architect says valuation for the, and this could be heterogeneous across advertisers, right? That is also the match value. So think of this as a CPM auction now, right? Any kind of efficient auction, it's the same whether it's CPM or CPC. So you already have this, do you guys request, you already have these max value estimates from the machine learning framework, right? Yeah. But you don't have this, which is the advertisers valuation for the click, right? That's the heterogeneity, right? And that could be heterogeneity there.

Hema Yoganarasimhan:

So we want to estimate that this may make it so we need both of these. And then, so that's a challenge we need to estimate advertisers click valuations from observed data and we can estimate match valuations for any targeting level. Okay. So how would we estimate a click valuation? So for those of you who are working on structural options, probably seeing this a lot, but thankfully in this case, there is, as I said, some theory work, which are [inaudible 01:04:24] proportional auctions and you chose that you can actually, we can get an approximation of the click valuations as twice the bid, but we also try many other alternative methods which are much more complex or essentially think of this as there an auction mechanism. It's just being used in the data now. Right?

Hema Yoganarasimhan:

And from that you can back out valuations and advertisers know that from that you can back out valuations and you're going to use that. And valuations are primitives here, right? And assuming the valuations don't change, there is no reason they should under a different auction mechanism how much and advertise their value as a click you're going to, so there are two primitives here, one's the match valuations and the click valuations, the match valuations came from the outcome predictions like valuations come from this model.

Paul:

The value that I get from user clicking on my ad might depend quite a bit on that particular user. And you said that you have these groups of heterogeneous consumers who are lumped together.

Hema Yoganarasimhan:

Very good points. So, the thing is that in the data though, we don't see, that variation. So I can't really, so think of this as the average valuation and if anything we added that in the result will become even stronger because I'll get even more heterogeneity.

Paul:

Depends on which way it goes. Ready. The person who clicks a lot, it might be like a 13 year old, he doesn't buy anything. They click on everything.

Hema Yoganarasimhan:

No, but these are independent. I'm not saying that those batch value has nothing to do with this. Right, but this is coming from not from the click valuation. From the auction mechanism.

Paul:

Yeah, but it's, so let's say.

Hema Yoganarasimhan:

That's coming simply from the bidders,

Paul:

There's two different types of consumers who are grouped together, in the current system. And one of them might have a very high value for you. One may have a very low value. You're going to take the mean value. I had to both people. So you have a 13 year olds, he doesn't click on anything. You're going to assume that when you can target to him and he clicks through with a very high rate, you're going to see that that's a great match.

Hema Yoganarasimhan:

That's a half way because I'm going to, you get only half the valuation to that, right? That's the average. Yeah. So the average would still be right, but if I could impose this, if I could get the actual click valuation, it's only going to increase the heterogeneity in the valuations even more. And you can show that.

Paul:

Okay. So assuming that the bidder is still bidding on that group of consumers, that will, then it'll tell you how the revenue would change. Okay. So they're not bidding individually?

Hema Yoganarasimhan:

No. Okay. So I think I then do I need to end? 15? Okay. I have enough time. Yeah, go ahead.

Michael:

So the sprint, that is the be all [inaudible 01:06:57] those come from the actual advertisers on this platform. They provide this number?

Hema Yoganarasimhan:

No, this is the bid. We see the bid.

Michael:

They bid on the whole category. Right. So they come in presumably like once a week. Once a day.

Hema Yoganarasimhan:

Yeah. And whatever. Whatever bid we see for the impression. We can use it. Yeah. Go ahead. Sorry. You didn't want to finish. Yeah.

Michael:

So then that needs to be sort of somehow grounded in reality for them times. I worked for [inaudible 01:07:26] for a long time to do some of this stuff. And I saw many, there was like a massive dispersion of how logically it would actually fit. And there'll be many of which would be happy take a loss because they think they're going to gain market share or they just, they've got this budget and they have to spend it.

Hema Yoganarasimhan:

Yeah. So, okay. It's true that in some ways, so the initial version of the paper we actually the to show that are then your efficiency trade off you actually the ideal cases, homogenous valuations. Okay. So because even with homogenous valuations, if this is happening, it's purely your natural use which are driving it.

Hema Yoganarasimhan:

And that's what we had initially in the paper. But of course when it went through the interview process that our viewers are like, Oh, you need like be diocese would have different valuations and it is true in principle that they would have different valuations. And then when you impose that, if anything is going to increase, the heterogeneity valuations and make the results stronger. And which it did. I'm thankful. So, but when you want to back out valuations in any kind of like a theory mode, a lot of structural model like theory taken to data your basis, you are making fundamental rationality assumption that the bed is reflective of be a bit as a combination of the auction mechanism and they're true valuations if they are using many other things, which I'm sure they are.

Michael:

But you could at least check that there's like a really strong correlation between the split value from each advertiser and say like CTR.

Hema Yoganarasimhan:

No. So the thing which is no, it's not clear there should be, right because it's again a function of the auction mechanism. It's not, the bid is a fundamentally a perimeter. Whatever structure you assumed about the auction there is no way like to get it, yes.

Paul:

Just to go back to the 13 year old example. You're assuming that the bids don't change. But if I had differential evaluations for these three different consumers and suddenly almost all the clicks I'm getting are from these 13 year olds consumers who I think have a very well, then you will change the yes.

Hema Yoganarasimhan:

And all I was saying is that if you do that, you will only see more heterogeneity because you are now moving, your clicks to your bid is going to increase for the people who are actually clicking on your valuation is high. Right? And the guy who are clicking but not actually buying anything, your bed is going to decrease for them if anything is going to increase the heterogeneity and make the results stronger.

Hema Yoganarasimhan:

Yes. Yeah. If you guys have readily bid, so in some ways the strongest case where you can show this is when you assume completely homogenous valuations, which is what way? And then as you add more heterogeneity from one valuations that could start driving these results in some ways. So that's fundamentally the idea. Okay, so next, okay, how about? Okay, so let's, we come to match evaluations and the targeting. So for any targeting level, think of any targeting level as the way you bundle impressions, right? So what you can think of this is that the fundamental, let me just show you an example. So how do you aggregate them from the machine learning [inaudible 01:10:34] so let's say you have two ads. I now I think of any kind of aggregation strategy. So some aggregation strategy. Let's aggregate these two, these two, these three.

Hema Yoganarasimhan:

So let's say that I'm doing some kind of upload or targeting and these are the three apps which this happens. And so what you, what this kind of tells you what, what we do is in some ways you can distinguish these impressions within this wonder and you're going to kind of get some kind of an average. Okay. So that's kind of the adoration strategy. That's how you aggregate match valuations under, that's how you come up with a target. That's the, match value estimates for counterfactual targeting scenarios. So now we have that. You can, so we have, we needed two perimeters to kind of solve this, right? So one was the match value estimates and we know how to aggregate them and the click valuations and then you can, put you any final thoughts, efficient auction mechanism. I'm not going to walk through this, but basically for each impression you can get a link for this targeting scenario.

Hema Yoganarasimhan:

Who wins this impression, you can determine the winner. What's the total surplus for this targeting scenario. And what are the platforms that have the news and the advertisers surplus. I think we can all see that intuitively just to ignore the math. And that's why I like the matrices. Okay. So, so let's talk about the counterfactual results. So these are shown in like average, improvement in a currency. So think you can interpret them collectively. But the actual numbers are in the local currency. Okay. So I still can see the first thing which is nice to see is that Surplus has a monotonic relationship. You're targeting granularity. There's the total welfare or the surplus in the system. And, that's nice. And you can also see that behavioral targeting has higher, surplus then contextual, but of course full has the highest.

Hema Yoganarasimhan:

But if you look at our platform revenues so the [inaudible 01:12:36] revenues, you can say that it's actually maximized under the contextual targeting. Anytime they start using behavioral information. Interestingly, the platform revenues don't change much from, behavioral to default. So what are the, even going from here to here, what about extra value is being created. It's going to the advertisers and you can see that advertisers surplus actually is the highest. Then there is more information. So now they're able to place all these like targeted bids and make, so all the value creation when you do more targeting after a certain point is being appropriated by the advertisers. So does that make sense? It saves that platform has some incentive to preserve privacy and that's good. So the goal of this exercise is not so much to say that context, the platform should engage in contextual targeting rather to kind of fees out the incentives of the different players.

Hema Yoganarasimhan:

And also to say that from the platform's perspective, if they want to actually think about revenue, they should not be thinking about efficiency maximization, which a lot of platforms do to rather think about revenue maximization. And then there are many mechanisms I'm sure, and they can put in place to actually extract back some of this revenue. Right? For instance, platform. Some of the platforms are already starting to do that. If they can actually provide some information on what not, what the click valuations or the match evaluations that as a service and extract back some of the value, then they can get back. But if you are just going to maximize efficiency without thinking about extracting that surplus back, you're going to be worse off.

Shane Greenstein:

Can I mark something. So is it just the byproduct of this sample of data or do you think there's generality to it that most of the surplus did not require contextual or behavioral targeting that you generate quite a lot of value in the absence of, and then the incremental part as a percentage of what was already there as it turns out. I mean, don't get me wrong.

Hema Yoganarasimhan:

Yeah, no, that's a good question actually. And this is also something we talked about because if you look at much of the theory literature in this area, they have very strong functional form assumptions on valuation. And they have these like really fat tails which we don't see in the data and results on improvement. Like I think they are like the Hamilton McAfee paper and so on. For instance, there was also in theory are like so strong in terms of how much, Oh like you're going to die. It's like really bad [inaudible 01:15:07] shape curve. But in the data we don't actually see that. So I mean there are then your efficiency trade off exists but it's not.

Shane Greenstein:

So another way to ask the same question is I can I think of the contextual advertising here as the person is near the store cause that's really what you're picking up.

Hema Yoganarasimhan:

Oh your person is in this app at this time of the day.

Shane Greenstein:

Yeah. And so I liked the location because that's a real, that's so you throw the right ad right at them when they're right next to the night. Right when they're next to the target. That's a good pun. In this context.

Hema Yoganarasimhan:

That's funny.

Shane Greenstein:

Oh God. You didn't even get it.

Hema Yoganarasimhan:

They are laughing now.

Shane Greenstein:

But the point is it doesn't generate very much. Yeah.

Hema Yoganarasimhan:

Because if you think about the click through rate improvement still right in the most efficient state. Right. It went from 0.66 to one point something.

Shane Greenstein:

It's because it's still so bad.

Hema Yoganarasimhan:

It's still so low.

Shane Greenstein:

Oh, that's why it's this way.

Hema Yoganarasimhan:

Yeah. So how much have you improved? If I could somehow improve clicks by like[crosstalk 01:16:24]

Shane Greenstein:

That makes sense.

Hema Yoganarasimhan:

By like 500% then I would, I mean that's still a bunch of people clicking. I target them. I mean it's still, I think about like three or four. I forget the exact percentage. So you need the system to be really large to see the improvements. Right.

Shane Greenstein:

Okay. Yeah, no, that's very helpful. Thank you.

Hema Yoganarasimhan:

Okay. Yeah, go ahead.

Diane:

I guess I wonder if you could tell us a little bit about how to think about this. Keeping in mind the fact that advertisers have choices so they could take their business elsewhere. So I guess I'm trying to sort of zoom out and think about these results, but in that context.

Hema Yoganarasimhan:

So that's a good question. So in some ways we were lucky because we are working with like pretty much a monopolist. Like they had 85% of the market share. If they actually could go somewhere else with this advertising budget to like a different app, this value creation could then be, you could say that you're getting on if some of this value to the advertisers because like you are trying to generate loyalty or something and then that helps. Or you could make it more explicit mechanisms that, if you have it as a service where they get this information and then they know this value is coming from it. Yes. So, the goal is not so much just stall the platform, exactly what to do, but rather say that this exists and you want to be aware of it and depending on the context you might want to think about, it's almost like a mechanism design question.

Hema Yoganarasimhan:

Do you want to extract the surplus? How do you do it? And so on. Okay. So briefly about advertisers surplus. So here we actually one of the things you saw that advertised surplus is increasing on average, but it's not always the case. So what I do in showing you here is that going from like lets say behavioral to full, trying between 23 advertisers out of 37 benefit, but going from four to behavioral, which is less, right? Actually about 14 advertising advertisers benefit. So it's not that all advertisers have the exact same preferences in terms of how its the optimal level of stuff I just gave you because I actually don't know the identity of the advertisers. It's very hard for me to fully reason why a specific advertising benefits and doesn't. But let me just give you an example of like, you know why, like between contextual and behavioral.

Hema Yoganarasimhan:

So if you have an advertiser, let's say you have a contextual advertising regime and I'm a nutritional supplement ad and I advertise in a fitness or a health app, obviously no one maybe else wants to advertise in this and I get these things at like a low enough price, but now let's say I start allowing for behavioral targeting, maybe there are people in there now I suddenly valuable to a bunch of other advertisers which was not visible to them before, right? Maybe these are people who are spending a lot of money on other apps and suddenly they become valuable. And to me as an individual advertiser, I'm suddenly competing with those people for these users. And that's why it's not like a uniform result, but each single advertiser then if it's equally, and it was kind of nice because again, we tend to think of the advertising industry as like one big like group with exactly the same incentives.

Hema Yoganarasimhan:

And it turns out it's not always the case. But it's heterogeneity in how advertisers, what regimes kind of advertisers prefer. Okay, so that was pretty much what I have to say to conclude very briefly, the paper times to kind of has two types of contributions if you might call it that. The methodological front warfare proposing is like a scalable machine learning framework for tiny thing that is actually compatible with contra-factual analysis in a competitive environment. Right?

Hema Yoganarasimhan:

So you can think about all these targeted contra-factual and from a substantive perspective, we kind of do this extensive comparison of behavioral and contextual targeting both in factual and contra-factual competitive settings. And there are some implications from a manager perspective, you find that there is this non monotonic relationship between revenue and targeting more granularity. And from a policy perspective there are some evidence at least to show that ad networks have some incentives to self-regulate and advertisers incentives on what's the optimal level of targeting. You there is some heterogeneity and that's pretty much it. Thank you so much.

Shane Greenstein:

Thank you.