Models for Sponsored Search:
What are the right questions?
(Panel Discussion)

Panelists: Kamal Jain, David Pennock, Michael Schwarz, and Rakesh Vohra
Moderator: Jason Hartline
Transcript: Maria-Florina Balcan, Jianqing Chen, Nikhil Devanur, and Anuj Kumar

Panel: June 11, 2006
Transcript: September 27, 2006

The Second Workshop on Sponsored Search Auctions was held in conjunction with the Sixth ACM Conference on Electronic Commerce in Ann Arbor, Michigan on June 11, 2006. The workshop closed with a panel discussion on the subject *Models for Sponsored Search: what are the right questions?*. This panel was proposed by Lance Fortnow and Rakesh Vohra. The panelists were Kamal Jain from Microsoft Research, David Pennock from Yahoo! Research, Michael Schwarz from Yahoo! Research and U.C. Berkeley, and Rakesh Vohra from The Kellogg School at Northwestern U. The panel was moderated by Jason Hartline and transcribed from a noisy recording by Maria-Florina Balcan, Jianqing Chen, Nikhil Devanur, and Anuj Kumar. The transcript is as accurate as possible given the nature of the recording and is meant solely as a research aid for those doing research in the area of sponsored search. Thanks are in order to all who participated in the panel and to those who made this transcript possible.

1 Opening Statements.

Jason: I'm delighted to introduce our panelists Kamal Jain, Rakesh Vohra, Michael Schwarz, and David Pennock. The topic of this panel is on "models for sponsored search: what are the right questions". The way that we are going to run this panel is that first our panelists are going to have a moment to make a brief statement about this topic in general and then we are going to proceed through a number of broad topics. We are going to give the panelists a chance to respond to the topics and then we'd like to encourage questions from the audience on these specific topics. Then after proceeding through these specific topics, if you have questions that don't logically fit under these categories, then we encourage you to ask those at the end. Also, I'm delighted to announce that a number of students have volunteered their time to take very careful notes during this presentation and including the fact that we are recording this discussion so that we can make sure not to miss and gems that are said [laughter], uh, gems or otherwise that are said during the course of the panel. So, we should definitely be appreciative of their efforts as well as the efforts of the panelists. So let's start with Kamal with his brief commends.

Kamal: Thanks, my name is Kamal Jain and I'm from Microsoft research. the first thing I would say about sponsored search is that it tries to price everything at a very atomic level and individual auctions are run very often. That allows a higher optimization. There is another place where the same thing happens. It's called the BitTorrent system. It is a peer-to-peer system where people exchange files where you give as many packets to others as you receive. The BitTorrent system is much better than any other peer-to-peer system. These systems are a starting point. There may be many other systems like that where we use the power of the computer to optimize and price at an atomic level. What it is lacking [in sponsored search auctions], i think, is that it excludes the user [a.k.a. searcher] from the game. It has two parties, the advertiser and the search engine and the searcher is not in any bargaining position right now. We should bring the searcher in to the picture since the search engine business is a competitive field. Yahoo, MSR, and Google each has capacity to serve all the searches in the world, it is a good amount of supply, but the competition between the search engines does not get transferred into competition on the pricing model. There is no reason for Yahoo, Google, or MSN to change their pricing to increase their market share [in terms of volume of searchers] because these two are completely independent.

Rakesh: I'm Rakesh Vohra from the Kellogg school. I have a list of things to talk about that roughly match what's up here. The first is complaints. One has to do with the fact that a lot of work has been devoted to a special static model where the private info is one dimensional. This is odd because in this model because when the private value is one dimensional, we know what the answer is. If we want efficiency we know it has to be Vickrey, and if we want revenue maximization we know it has to be an allocation by virtual values. So that leaves us with a puzzle: why don't the companies use these mechanisms? Why do they use the other ones? One possible explanation is that this is a mistake, in the same sense that those of you who have taught mathematics to undergrads have seen is something called the freshman binomial theorem that $(A+B)^N = A^N + B^N$. So maybe what's going on here is the freshman Vickrey theorem where the generalization to N goods is simply the price of the person just below you. But if you take this one dimensional type model seriously, as I've said, it does not explain why we see the mechanisms that we see. One explanation that has been offered is that, well, why not look at budgets. Well, the problem with budgets is that they are soft. The bidders can change their budgets every time. So it is not so clear that these budgets are really serious, and even if you take budgets seriously you have to be careful how you set up the model. Here is a really stupid model with budgets: if I leave you parked on a keyword long enough, I soak up your budget. So in that case, why am I not selecting bidders by their budget? All the bid-per-click is saying is the speed at which I use up the budget. so if you really take the budget seriously then it would say that the allocation should depend not just on the bid-per-click but also on the budget. In terms of the various ways of paying, for example, pay-per-click, pay-per-conversion, this deserves more discussion because this says something about how risk is being shared between the search engine and the advertisers. The nice thing about bid-per-click is that the search engine is taking on more risk that otherwise the advertiser would have done. Other arrangements result in more risk sharing. One question is which risks are the advertisers better able to take care of because they are diversified over many keywords. Another issue has to do with the fact that you've got competing search engines, so if they use different mechanisms that could induce a selection bias in terms of the kinds of advertisers that would be attracted to a particular search engine. So something very simple minded:

if one was using Yahoo's discontinued mechanisms, one would have thought that maybe it would have attracted the kinds of bidders that who generate a high expected value given a click, but generate very few clicks. [They would have preferred the auction where bids are not normalized by the click-through-rate.] They would have preferred the first price auction to the second price auction. This would have resulted to a different distribution of bidders at the web site. So in thinking about whether you switch from one mechanism to another it is not enough to look at the mechanism in isolation, in the usual classical framework, but you also have to think about how that is going to affect the pool of bidders that you are going to attract. The last issue is ambiguity. A couple of the search pages have auctions whose rules are ambiguous in certain respects. One question is as to whether this is necessarily a good idea. Apparently two arguments for ambiguity are (1) is prevents free-riding on the part of your competitors and (2) it would allow you to change your mechanisms without having to deal with complaints on the part of your bidders. On the other hand, the empirical evidence suggests that ambiguity would make bidders more risk averse and so that could potentially lower revenues.

Michael: Hi, I am Michael Schwarz and I am from UC Berkeley and I would be joining Yahoo! Research labs in couple of months. First I would like to thanks the workshop organizers specially Jason and Asdemir. I think this is wonderful conference and it's great that people from academia and from industry come here and talk to each other. I think it makes a difference for all of us. So let me start with a question that has one very very important advantage. Nobody mentioned this question today. So maybe if it's not THE most important question, which it isn't, at least perhaps it is controversial and entertaining enough to grab your attention for next 100 or so seconds. So on the plane to Detroit I was reading a newspaper article and here is what I learned for the first time in 73 year history the radio city Christmas spectacular will use product [placing]. So advertisement is everywhere. It's crouching. We heard this morning that the yield from advertisement in different media is dropping as there is more and more advertisement coming in. So one thing that we perhaps should be thinking about is: what are the main competitors to those Google sponsored search links and there is one thing I know for sure, it's not radio city spectacular Christmas products. So then what is the main competitor to Google sponsored search links and you may say that it's Yahoo! And actually it is not. It isn't at all. What is the main competitor for them is Google organic search results. So in fact if Google was a monopoly, if there was no Yahoo!, if there was no MSN, imagine if Google were to dramatically decrease the quality of the organic search results. Well you would think that advertisements would become more competitive and a user would click more often on the sponsored search links. Well, how do I know that, for example in the days when Overture was an independent stand alone company and it wasn't owned by Yahoo, pretty much every user who went to Overture website would click on one of the ads. What a miracle and in the absence of organic search results that would happen. The same thing would happen to sponsored search results if organic search results were not there. So of course there aren't many people who are willing to go to the site but don't want to click. So when you think about it and you realize that for Google sponsored search links the main competitor is Google organic search and for Google organic search the main competitor is Yahoo!. So that necessitates keeping [the organic search results] there. So the question is, and I think there are many questions here, first, how do we think about competition in this market? And I know that I am picking something that is somewhat relevant because completely without prior communications with Kamal, I think this question echos his comments. and the other question is first what is the future of the Industry? How does that competition play out? To what extent [will we] see product placement in the organic search results in the future just like Radio City Christmas spectacular is experiencing, and if that would ever happen, contrary to FCC's best advice, how would it be priced. How would competition play out? So I think it's both an interesting question and perhaps a plug for my future work.

David: Hello I am David Pennock and I am research scientist at Yahoo! Research and myself and my group work on a variety of market mechanisms that are important to Yahoo and, of course, sponsored search is one of the most important to Yahoo. It's one with the biggest leverage we have in terms of revenue. And one of the short and medium term targets that we can work on. In terms of models, what's the right model? So I think in this framework, the sort of worst case analysis or adversarial type analysis is of some interest, but it's not really the right model in this framework. In this framework we have huge amount of statistics, basically we have very good estimates of priors essentially so I think the Bayesian type model or expected case analysis is more relevant in current setting. Let's see some other specific [questions]. One of the questions was if truthfulness is important. To me it seems like truthfulness is not really achievable in practice. So to me that as a quality of mechanism design may not be that essential. But on the other hand stability is nice. The stability is not essential but its nice at least in terms of computational resources of the system. So let's see, what the current system does right, I think in the head terms [those that are frequently searched for], things work pretty well. And this may not be the case in the tail terms [those that are infrequently searched for]. This is partly because of the cost for eliciting bids across all these tail terms, and inexact match [e.g., broad-match] addresses this to some degree. [It is] sort of like an automated way of translating the bids from current terms to some other terms but essentially eliciting the true value across all these terms maybe just not be feasible. Along these lines there is a general tradeoff, we want to elicit the advertiser's value for everything. Right, we want to elicit their value for impressions, for conversions, for clicks, for calls, for all different terms, for different times on the day, for different demographics, their value for having guaranteed click bundle vs. just spot markets. All these things we want and that would maximize certain social welfare and maximize value but of course we can not ask all these things. So what we need to do is essentially figure out the ones that would add the most value with the least complexity for advertisers – the least amount of things that we need to elicit from the advertisers.

So that's all I'll say for now. I just want to say one other thing. I just want to encourage people. I just want to point out, while I have the microphone, at the main conference on Tuesday there is some more Sponsored Search activity. So I encourage you all to check it out. There is a session on Sponsored Search on Tuesday at the main conference and also Hal Varian would be giving a keynote address on Wednesday related to Sponsored Search.

2 Models for Sponsored Search Auctions.

Jason: Thanks you, so we proceed to our first topic which is the current models for the sponsored search and again focusing on what these models get right and, maybe, what they are missing, and how we need to proceed to get models that are appropriate. We want to consider things like reputation, multiple keywords, and a number of other aspects that are missing in the models that are considered so far. So, are there any questions?

Audience: Rakesh, you had some comments about the budgets. So we don't use virtual valuations and you don't use VCG and they are saying that in the community one reason for this is budget. But you don't believe this reason because budgets are in some sense soft and if you really care about the budget you should use that information in the ranking mechanism. So what about using it in the ranking mechanism? Do you advocate that or do you think it's completely, I mean these things exist in the market and they have existed for a long time so you can't just say they're . . .

Rakesh: Right! I guess so. I have asked various people why we have budgets. I have asked advertisers and the sense I get is that they like the budgets because it allows them to do something else. Which is to minimize risk. I want to control my exposure. So that's the real reason that they like the budgets. I don't think that the budgets necessarily reflect the hard constraints. So if you take what [they] said seriously then it says that there is a preference for something that the bidders have that they are not able to express directly in the current mechanism. Which is that I like to pay but I want to make sure that I don't pay too much over a two month period. So one could imagine [a mechanism that would say] well .. fine I would guarantee you a certain slot for a certain defined period. And whatever you make from that you are allowed to keep but then pay me upfront. You could imagine doing something like that, then you wouldn't have that problem. Then the bidder would know that, well, "this is all I am in for", so the budget constraint is not so important.

Now the other thing is that if you do take the budget constraint seriously, there are arguments [that say] it should be taken seriously because the advertising department of large companies are given budgets that determine their expenditure over a year, then if that's the case, then why aren't you selling advertising upfront because if I can capture a larger portion of your budget, there is less and less budget to go around for my competitors. But I don't see that happening in sponsored search. In advanced markets, or serious advanced markets. So again I come back to the point which is that I am in two minds about this. I can see the arguments for it to be serious but if you take it seriously then I don't understand the (current) mechanisms.

Audience: So what mechanism would you advocate. Paying upfront [involves some sort of negotiation and presents a barrier of entry] to mom & pop shops. The nice thing about auctions is that they are very transparent and low cost automated pricing mechanism.

Rakesh: So the general proposal was a system where you offer a slot for a defined period of time.

Kamal: So the issue would be that there are two type of advertiserz..1) the big advertisers 2) the small advertisers. And we don't have two different set of tools. For small advertisers [which may form a long tail] the main preference is small budget. If you are selling for the current period, you don't want to sell it for 2 seconds, you want to sell for a day. The small advertiser can not afford that slot for the day. So how do you want to involve him?

Michael: To pull off from the budget discussion, there is actually a lot of academic interest in budgets partly because there are some fascinating theoretical problems associated with the budgets. But then there is another question. Although budgets are not binding, [if they are then they don't matter,] why does an advertiser have budgets. Well a simple reason for budget is when you have an attack of fraudulent clicks, if you have budget then that couldn't damage you that much. And of course the problem is that, though you get refund in the fraudulent clicks, you may wait for a month for that refund. So you have an incentive to set budgets, perhaps a non-binding budget, just to save your self from the attack of fraudulent clicks. So then there is a question - an empirical question - to what extent are those budgets binding? What percentage of advertisers that have binding budgets? And we know the answer that Google gives to this question. The answer is that it is confidential information. And we know the answer that Yahoo gives to this question and I am afraid that's very similar to the one that we get from Google. And you can also try asking MSN. So some folks here work for i-prospect and other companies that actually handle advertiser's accounts. So they could share their experience on how constraining binding budgets are. But I think one interesting question: is it something that's actually important out there?

Kamal: One thing which Google says is that if you have budget limit we will try to uniformly spend your budget over time and only to give you that much. You can see that when you reload some of the ads disappear. And you can see at least from the perspective of the advertisers that the budget is a constraint.

Michael: Right. So the question is: is it something which is empirically relevant. It happens 1don't know!

Kamal: You can notice, ads disappearing if you keep pressing reload button.

Michael: Well, part of what that might be that, if it's a rare keyword, [you may want to experiment with different ads to get clicks?]

Kamal: I think the counter fraud attack argument for budget is not correct because it would be same for all the advertisers.

David: Well, it can also be sort of an estimating procedure of expected click-through-rate.. If you just saw a bunch of ads and then reload the page. You already didn't click on those ads.

Kamal: If you keep reloading, the ad goes away, and then reappears. So the most plausible explanation is that they are trying to uniformly distribute the ads.

Michael: So that seems like a nice methodology for figuring out the answer to this question with out relying on Google.

David: So you can ask Yahoo! We do have mechanisms for smooth spending over time.

Audience: But why do you do that? Is it because of the budgets or because of some other reasons, or it's just that the advertisers want that?

Kamal: You know value of clicks are different over the day. If you are letting them bid for the day and not for the hour, no advertisers gets a bad set of users. That should also motivate the desirability of smoothing out the spending.

Jason: The questions was about the existence of this mechanism at all.

Audience: It's an advertiser's preference.

Audience: I can answer part of the question. I am not from iProspect I am from their competitors. I have seen every kind of advertiser, from 1 million dollars to a few 100 Dollars. I would say budget pretty much drives what they are spending per day, for several reasons. One is that they want to spend it throughout the month and not just one day. Let's say you are running a business you want to see customers coming in throughout the month not just 5 days of action and then 25 days of silence. The second is risk. You want to limit your risk. [Also,] most of them do operate with many other constraints. [For example,] maybe they have a cash flow problem or something. In some cases in our product, we have a statistical model which predicts how much they are going to spend and how many clicks they are going to get. So we ask out customers to remove and not to set the [budget] constraints at the search engine level, because Yahoo behaves differently and we manage differently. We are able to predict how much you are going to spend to a very close approximation.. but even then the customers are very particular about their positions, [they want] 5limit. It is true that in some businesses, for example [lead generation, where you lead them to a form for a home application], they say, ok get us any number of forms per dollar. There are also other kind of constraints that come in to play. For example, a call center, they are not able to take more than 500 calls a day. there are inventory problems. We may say that get me any number of forms, but there are finite number of offices, employees, etc.,so there is this resource constraint too. There are also website constraints where your website can not handle more than certain number of visitors.

So budget is made a proxy for all these other kind of things. If you have any interesting research problems, you can contact me, we'll provide you data to research more.

David: I was just curious, how do you manage certain unexpected jumps in volume like Janet Jackson episode, etc.

Audience: That's a very interesting question. I think it happens once in a while, maybe some news event happens, [though] news event are less likely, it is more like Yahoo changes it's front page. Say it put a few baby related things. So all your baby related products are going to sell like crazy. So we constraint it because we have a totally automated software platform, looking at it every hour. So when [something] goes above the limit, we shut it down.

Kamal: So any information about clicks and impressions, do you get real time information..

Audience: As real time as it can be. Yahoo guarantees to publish this information, say, every morning, so there are some other limitations. But we are trying to be as real time as it can be. So there are also quarter end effects. Somebody wants to optimize revenue towards the end of the quarter [of the fiscal year]. Or somebody says, I don't want more revenue, I want more profit. So the budget becomes a proxy for [all this].

Rakesh: So this raises the issue that whether budget constraint is necessarily the best way to address this? Right?

Audience: There is no alternative to this.

Rakesh: I guess I wouldn't be prepared to go that far, unless 50 people thought hard about it and come back and say that there is no alternative. Then I might be persuaded.

Audience: So why not have budget constraints on number of clicks, rather than money?

Audience: it's just like financial market, [click is currency and budget on number of shares is same as on cash.]

Audience: I just want to make a quick comment. Following the previous comments that in my experience that in this you can choose how much you are getting though the bidding price, instead of controlling though the budget, because if you have high bids but have small budget, actually you are paying more for the traffic you are getting. Instead you should bid low and get a smaller slot. And still use up your budget.

Audience: That requires automating, you can't do it manually.

Jason: I got a question here.

Audience: A primary discussion question for the panel in topic 1, basic issue is about basic model we are using nascent here studying the sponsored search auction for advertising. I don't think we have something as formal as common model, as an example such as tax model, and so on. But almost all papers have a common abstract model about what an auction is that there is an inventory of advertisers' click-through rates and their estimates of click-through rates. What do you think are the one or two most vulnerable assumptions everyone is making. For example, one of them I would think might be the click-through rate is estimable property. We assume the Ad inventory is static over some period time and that general Ad response preview users against click-through static ads is going to be estimable at constant rate. But if you move to inventory of item level ads that every time someone types in Da Vinci Code, one of the Ads pops up with a different used book seller at different price, click-through rate is varying all the time, what might happen. So, essentially, what are the assumptions we are making that might be the most interesting to re-examine in the future.

Jason: Anyone want to take that?

Michael: Sure! So, one assumption that we all sort of universally made, which is probably or not always true is the assumption that advertisers don't really care who are the other advertisers on the page, or if there are other advertisers on the page. Somehow I would imagine that if I were HP, and it seems that HP is a canonical advertiser today, and I was bidding on the word cartridge. If I was there only ad for cartridges, I would imagine that I would get more clicks than I could get otherwise. I would also imagine that I would in fact get a lot more conversions than I would get otherwise. I would even eventually, perhaps, choose to charge a higher price on my web page than I would otherwise because really people who came there don't have much [other] choice. So this seems like that something that could be potentially quite important and relatively easy to address and yet, so far, we sort of conspire to ignore it. Probably we shouldn't. I think I could also mention another three thing that are equally important . . .

Rakesh: So, you can say that click-through rate is changing. The problem is that in static model, you cannot really address this issue. And so, if you want to address this issue, you've got to start thinking about dynamic model, where the keyword arrives over time and the bidders arrive over time. You saw some papers already talking about this. They were talking about the competition between two bidders and perhaps trying to drive the price down or up, there are lots of other things that could happen which are frankly just part of model. So, there are a lot of things you can think in terms of standard model. But one has also to be honest to say, the alternatives, are very analytically intractable. So, it's not clear what one could say. Perhaps there is a way out by looking at the keywords which are extremely popular in the sense of having a large number of bidders. Then you may be able to make asymptotic arguments, because the number of bidders is much larger than the number of the slots available. Perhaps one can say something about very popular keywords, but be silent about the keywords that are clicked on very rarely.

David: What assumption that almost all the auction analyses make is that the auction is in isolation, sort of monopoly type of situation. So, in that case, you know you maybe do Myerson to optimize the revenue, or something like that. In fact, this world is competitive landscape. You may not want to optimally squeeze as much as possible out of. So, essentially what the search engine do is that they add value to advertisers and take some amount of that for themselves. And, so, you may not want to optimally squeeze as much as possible out of pie, because this is a competitive landscape. You know that advertisers and end users can defect very easily.

Rakesh: Actually, this is not so easy to do. The defection is not so obvious, right, which is you could say, as Pennock said, why cannot Microsoft lower the price you have to pay to advertise. But the problem is, that is not enough, you also care about the number of viewers you are going to get at MSN pages. So it is not clear that I can necessarily defect if I don't like the mechanism, because I also care about the number of viewers that your page attracts.

David: Right. If you get essentially no profit at all in one engine, you just put more money in the engines where you get defect problem.

Jason: Eh, yes?

Audience: One thing drawn from my experience from about 300 markets [is that] the story is always the same in all of these, which is you want to increase the expressiveness of the bidding language and you want to customize to the applications, and that would make market better. And in that sense, I see all of these Ad auctions as extremely simplistic. And this is really most related to what have been mentioned by David and by Rakesh. And, you know, is the budget constraints really [hard]. [There are a] number of other things we should be expressing in different ways. Let me just mention a couple of things. I am not an expert on that auction so these might not be the most important things. If I get exclusive Ford cars for the next month, you don't show Ads at all in the next few weeks, I will pay you just for not showing them regardless of whether you show me. I might say that I want to get 40,000 clicks in the next week. And things like these that you cannot see in current markets. And I suspect that Ad auctions wouldn't be different from all other auctions in this sense that expressiveness would make market more efficient and also you can control the revenue estimate. We heard an early comment that MSN is giving more expressiveness than Yahoo and Google, and Yahoo and Google have different types of expressiveness as you can see at their sites. But those are all what we call local expressiveness. So they can take a little thing and translate that to price. They are not allocation level of expressiveness. And there are a lot of things at the allocation level inherently, and then you cannot use [the current mechanisms.] that requires a completely different architecture. You got to have something that [optimizes] at the allocation level. and in these other markets what we have seen is that [at the] allocation level you need to solve NP-hard optimization problems, for determining the winner, for example. That's not a problem, you can absolutely solve NP-hard stationary problems. Then the question is can we do it in real time. And the answer is no. You cannot do it. You need do this Ad auction in a dynamic sense. Our proposed solution that David and I presented in this workshop a year ago [is that] you optimize and parameterize a dispatcher that doesn't use optimization. Then run the dispatcher in real time and the optimizer every so often. A related problem is there is a long term market and a spot market. Spot market is in real time and long term market is optimized. I think the current models are pretty naive in their approaches, they are limited inherently in that no one can support allocation and expressiveness.

Michael: So, first, I am very much in an agreement with you. You really need an auction mechanism to allow advertisers to express their preference, because if they couldn't fully express their preference, then there is no efficient outcome in optimizing the total pie. So this being said I think there is another side of story. And the other side of the story definitely doesn't want a mechanism that allows advertisers to express preferences that they don't have. Let me elaborate what I mean. For example, suppose you create a mechanism where you can express preference, I want to be number 3 on the page for a keyword "Ann Arbor" for the next 30 days. So this is a well defined preference. Should we allow advertisers to express it? I think they should not, because there is no advertiser who truly has this preference. Because being in slot number 3 is not that similar or different from being in slot number 4 or slot number 2. So once you allow it you can make preferences infinitely complex. The question is how many advertisers have this kind of preference. So if you make a mechanism more complicated than it has to be, in order to express all possible real economic preferences, so what you are doing is simply piling on advertiser increasing his communicative burden without in fact increasing efficiency. So I think an ideal auction mechanism has to be exactly as complicated as it has to be to fully express advertisers' preference but no more complicated. And I don't quite know how exactly to nail it, but that's what you're shooting for.

3 Solution Concepts.

Jason: If don't go to topic 2, we are going to be here for another hour on topic 1. So the next topic is on solution concepts. And I think for me, the thing I want to highlight about this question is that we learned that something was important when we switched from the first price payment rule to the second price rule. That wasn't necessarily truthfulness. We didn't switch to Vickrey. No one does that. So I am kind of curious of what is important in a solution concept for designing new mechanism in this area? What property is it that we need to make sure that our mechanisms satisfy?

Jason: So Yahoo just switched their mechanism a month ago from something to something else. [laughter] In the switch, there was an opportunity to adopt Vickrey. Maybe, they didn't? Maybe. I don't know, did they? [laughter]

David: One feature that second price has that is useful is more stability than first price. There is no pure strategy equilibrium in first price. And certainly, with first prices, advertisers will be constantly switching. And now they can still constantly switch even with second price. But I guess it is a little bit less burden on recording things and logging, all that stuff. And it is probably not that important, but that is one feature that second price has but first price doesn't.

Kamal: Even the GSP has such features. Allocation depends on bids, but price does not. The price could be [a continuous parameter, allocation is a discrete parameter]. If you increase or decrease your bid by one, that doesn't really affect you. But in first price, you can get immediate profit. So I would say for any auction, it is a good thing to have that the allocation is a function of the bid, but the price you pay is not.

Audience: Vickrey and truthful auction leads to academic journals in business.

Audience: I have a hard time believing [Google didn't know GSP wasn't VCG].

Michael: I would make a claim. I claim that there is witness in this room who can testify that Google when they actually introduced GSP, at that time they actually had no idea whether it was a truthful mechanism. Does the witness want to speak out?

Audience: I was the first one [to tell Google that their mechanism was not VCG].

Rakesh: We shouldn't single Google out. Twenty years ago, there was a debate on how to run treasury auction, whether they should move from discriminatory price auction to uniform price auction. In uniform price auction, the last winning bid would pay the smallest winning bid. Milton Friedman came out to support the uniform price auction because he said that that's the Vickrey auction. Milton Friedman had the noble prize in Economics, and even he got it wrong.

Audience: By the way, there were several billion dollars of revenue. And they still kept it running for a while,

Michael: Also for the record, at that time, there were no Economists who are involved with Google designing this auction. So, in defense [Economists], this was not [their] fault.

Audience: Just to follow up on the questions. Do you know any other auctions using Vickrey that are successful?

Rakesh: Be careful of the criticism of Vickrey auction. These objections of Vickrey auction take the following form. Under these assumptions, Vickrey has these properties. Once one assumption doesn't hold, Vickrey doesn't have these properties. That's not really criticism of Vickrey auction, that's the criticism of the underlying assumptions. You could say that because these assumptions don't hold, I don't want to use Vickrey. But that raise a question, what do you want to use? [On the other hand, the English ascending auction, you could argue, is a particular implementation of Vickrey auction. If you actually want the sealed bid second price Vickrey auction, it is used, for example, to sell stamps to collections. The earliest recorded version I believe is that [Getta?] sold his manuscript by Vickrey auction. And you could argue, that eBay is essentially Vickrey auction. You could make these arguments, [though] it's problematic that bidders don't seem to behave according the way Vickrey described.

Audience: In Singapore, when you buy a license to drive a car, they used multi-unit auction, which is not a very sophisticated VCG. But they do use a Vickrey style auction. And I believe Google IPO is like [Vickrey]. We do see examples.

Audience: What about the repeated Vickrey?

Rakesh: There is a tiny literature on repeated auctions. Very tiny, because repetition means running into Folk theorem. So you end up saying that anything could happen, anything interesting, could happen. And that's not particular palatable.

Audience: If I could make one more comment on the VCG issue, so another response to why we aren't using VCG is because, you know, the VCG is not truthful in this setting for all the reasons we discussed: because of the dynamics, expressiveness, and it may be that in practice that for [the form that the] true preferences [take], the auctions that they use [e.g., GSP] are better. I find it a little frustrating the emphasis on the isolated VCG and why we aren't using it, it is missing so much of what the real problem is, in my opinion.

David: I mean another reason is that the second price auction is much easier for advertisers to understand, much easier to communicate, much easier to implement. So there are other reasons possible.

Kamal: you have to learn other things to run VCG, like the relative value of the slots, what the CTR will be, VCG requires a lot more data and GSP is simpler.

Audience: I have one question regarding the solution concept, basically the organic search is to make the user [a.k.a., the searcher] happy, sponsored search is to make the advertiser happy. But I think in the long term you really want to make the user happy. To make the user happy you really want to make the advertisements relevant to the user. How can we make the user of sponsored search more happy?

Kamal: I think long-term that what's really happening. Lets say you are delivering a web site, users want to see a web site, and . . . at the top. but you are, anyway, bidding on it. To measure the quality of the overall page, the user doesn't see what the [targetted?] bidder bid, the user sees the whole page. You paid and you are showing up so sooner or later, for the quality of the user, your rank in the organic search would have to go down. In the long term what's the natural equilibria? All the commercial sites are there and all the informational sites are there, on the right hand side. Not a single advertiser, although those that are not bidding, you . . . do you show for free . . . , but no advertiser would be able to do that because if he doesn't bid, he is not shown and all the competitors are out-bidding him. In the long term the equilibrium is: the overall quality of the page remains high, the users are happy, but a certain set of [advertisers] have to pay to be on the page. So it does take care of the user happiness. If the ad is not relevant, no matter what you want to pay it is not shown.

David: one possible way to make users happy directly is to attach a coupon to the link so if they click through they get cash back or something. So you could directly compensate people and that makes them happy, but there are other ways too.

Kamal: One think you want to do is bring the user into the bargaining process. [we expect search advertising to be more efficient than other forms of advertising because the user expresses what they are interested in. This increased] efficiency should not be kept by the [search engine] because it is a competitive market. It should be some how given back to the economy. One way is to make the user happier, to use this money. There are currently no practical ways to give the money back to the user.

Audience: Build a better search engine. That's how the money gets spent, often, in providing some free service like free email.

Kamal: If you have three different [search engines] then you have three different qualities. If you don't let companies compete on value, just on quality, then whoever has the best quality, whether it is Google or Yahoo or MSR, will start taking all the business and there won't be any competition left. In other markets, people compete by value. Cars are a good example, not everyone buys Lexus. People by Toyota too because they get better value from Toyota.

David: Also, people are inherently lazy, so [they use] whatever is most convenient. If InternetExplorer had a [search] box that went straight to MSN thats a huge ...

Kamal: That's one way to bring value competition. [Search engines] have to pay to [have their search box in the browser]. For example, Google pays Dell [to be the default search engine], and now the PC market is inherently more competitive and that value eventually gets transfered to the user. If you buy a computer and promise never to change the default search engine then you should get the computer for cheaper.

...

Jason: Are there any more questions on solution concepts? Or are there any questions on this topic we ended up on?

Audience: I have one more question. Most people think truthfulness is important. How far is GSP from [being] truthful?

Rakesh: So let's start with "why does anyone even care about truthfulness?" I think the usual answer is that, if bidding truthfully is an equilibrium, maybe that's an argument for why, amongst all the many equilibria of the game, that's the one that we will pick. But, notice that there is nothing in that argument that says that a mechanism has to be dominant strategy incentive compatible. There is nothing sacrosanct about that. Maybe we want something fuzzier like "truthiness". So here, I would build on what David said which is maybe what really matters is that we want to argue that the mechanism is stable in some sense. So we want some notion of what the stable outcomes of the dynamic process [would] be when using this mechanism? So as a first cut you could define stability as "there should be no incentive to deviate in a simple way, and do better." I say that as a first cut. I guess I'm sort of skeptical because you could simply hire someone to write code to deviate in a complicated way so maybe we end up having to worry about equilibrium anyway. but perhaps if the situation is complicated enough, we [only] need something milder which is the mechanism has the property that certain deviations will not be seen, and we believe that these deviations are in some sense the more salient ones.

Audience: As a corollary, do you guys believe that truth actually exists? Small advertisers might not necessarily know what their value-per-click would be, and that is the bulk of the market.

Rakesh: Um, so one answer is that we care about these things so that you can make predictions about whether we prefer one mechanism over another. Since we are only after stylized predictions, it is not so important that in reality they don't actually know what their values are. For the earlier question, is there such a thing as truth, the answer is yes, because I can simply define something and say, "it's true".

Audience: There are about 75 academics in this room, so that's about 75 factorial definitions of truth.

Kamal: It is a stochastic problem, that advertisers don't know what to bid, at least for small advertisers. They don't know that their probability of acquisition will be, so at least to get that number they have to do this analysis long enough.

Audience: I want to go back to the business model where users should get paid and the other side where content providers get paid. I go to Google not necessarily to read something Google is paid for, I go there because I believe the listings that come up will be high quality. There are some people who are not comfortable with that: large content providers, newspapers, people who have syndicated content. Do you think in the future that there will be an issue with content providers saying that they want to get paid.

David: There is a search engine that is trying to do this. If you have some content and you go to the search engine and you prove you own it by changing something on the HTML then they actually are paying the content providers. I can't comment on [what Yahoo is doing]. [But, this company] indexes the entire web, they still are a regular search engine, but they allow content providers to go in and say "i own that page and i want to get paid" and then they get a share of the ad revenue.

4 Should bidders understand the mechanism?

Jason: So i guess we can move on to topic three, which is "should bidders understand the mechanism." This topic is being posted because a lot of the details of the rules of current implementations of sponsored search auctions aren't precisely specified. There are little fudge factors that Google has, and I couldn't say about the other search engines. Yet, classically in mechanism design, a paramount assumption is that bidders know the game they are playing in. So, this is at odds, and the question is: what is the right way to approach this issue?

Michael: I'll give you an economical answer. According to economic religion there are two major truths about bidders in auctions. First, bidders are infinitely wise and they basically instantly figure out the equilibria even if it's not computable. Second, to do that trick, bidders really don't have to know the rules because they just evolve to them. This is just a religious belief. So the question is how do you view that religious belief. Do I subscribe to it? Actually, since I am an economist, I do subscribe to it. I think it is actually pretty reasonable that people would tinker with things, learn the outcomes, and figure out reasonable ways to bid. It is also reasonable to think that when they are tinkering they don't have to read the manual. In fact you've probably done it many times with your VCR without having to reading the manual. So you can think of the manual as the rules of the auction. So I think a good mechanism is one where bidders can relatively easily tinker their way to a good outcome for themselves.

I think there is a good argument to be made that you want the system to be transparent, because if anything it should help the users ... to lower the cost to figure out the mechanism. So I think there is a good argument to be made that we would like to make the users able to read the rules and know what is going on, but at the same time they shouldn't have to [read them] to figure out what's going on; just like your VCR, you've got the user's manual, but if it is a well designed VCR, you don't have to use it.

David: This religious belief is only true if there is no one who goes and screws with the environment all the time. So in Google or other search engines, so you tinkered with it for two months and you figured it out, and well, guess what? they've changed the rules. They do it all the time, there is actually no reason to believe ... so you can either avoid the question or you can rephrase the question: should Google keep changing the mechanism and not telling anyone. Then advertisers never figure it out, or take a long time to figure it out, or not?

Kamal: They should be telling each advertiser what their statistical parameters are. They don't even tell the [advertisers] what their quality score or click-through-rates are for different key words. Maybe if the advertisers knew what these parameters they could optimize better, maybe he would know that one keyword is not good, or something. These things can have impact if they are known to the public like in Yahoo's [auction]. Yahoo shows all the bids publicly and an advertiser can see how they would need to bid to chance the outcome. So, at the individual level, should they be providing this information? the fudge-factors?

David: Maybe they should get an email saying "you should improve your quality score" or "you should bid higher."

Michael: I'm sure you would get both of those emails every day.

Kamal: It's easier to just tell the users what their ad status is, what the fudge factors are, not just whether they should be improved. He maybe thinking "bidding higher will cost this much, or improving quality score will cost this much" and he needs to know how to spend his investment.

David: You have the model where the search engine and the advertiser have the same goals that they are trying to optimize. The search engine and the advertiser have almost the opposite goals. It is not clear that they should be providing all this information to the advertiser.

Kamal: There goals are to use the [advertising space] effectively, and in some sense those goals are aligned.

Audience: So I think that posting the bids is different from posting the rules of the mechanism.

Kamal: They are posting the rules, they just aren't posting the CRT [or fudge-factor]. The rules are: its your bid times some fudge-factor.

Audience: Posting the rules is one thing, but I think it is not very common to post the bids in an auction. For example in sealed bid auctions in Economics, you don't typically open all the envelopes after the fact. Maybe because the companies think this information very is private. Yahoo posts the bids and Google does not.

Kamal: It is the GSP, but there are lots of other parameters . . . when you want to post the auction rules you have to know the definition of those parameters and those parameters are software code. They have to say: this is the software we are using to learn CTR, and this is the software we are computing the quality score from.

Jason: But that software might be dependent on their hoards of data.

Kamal: Yeah, so they'll have to post the data in that case. What I'm saying is that the rules have been posted, but these parameters allow them to do anything they want to.

Audience: The first thing a judge is going to do when they look at a dispute is look at the contract. If they find an ambiguity in the contract, it is going to be used against the entity that writes the contract. So there is a possibility that if you don't write the rules, then the judge can make the rules. There is a possibility that regulation can come into the industry. So from a legal perspective, on not defining the rules, I think it is a big problem in the industry. For example, one of my friends is working for a credit card company. They have to be able to explain why a person didn't get a credit card. So on the same lines, Google must be able to explain why an advertiser didn't get shown.

Audience: I have a suspicion about this. I was wondering what people's reaction to it are. There's been a lot of talk about how this pay-per-click model is a risk-sharing mechanism between advertising and the search engine, where the search engine has adopted some of the risks (rather than selling on a pay-per-impression basis). It seems natural that the search engine would want to rank [the advertisers based on relevance] and the search engines have been already in the business of determining relevance of web pages anyway; [they use] very complicated formulas. So, it seems that now [they are] not looking just at click-through rate as a proxy for relevance, but using this whole complicated mechanism the search engine already has, for determining relevance. But now, it becomes very important for the search engines to keep the mechanism secret, because otherwise they give away the whole game, they give away the way that they rank all the web pages, not just the ads. And I wonder if this is the reason that it's so important that the search engines [mechanism] must be a secret.

Kamal: CTR [at a very vague level] is converting a pay-per-click system into a pay-per-impression system. So, in that case pay-per-click and pay-per-impression would be the same system. Eventually, what the search engine's goal is that he is getting inventory in terms of impressions. He wants to make sure that this inventory is sold on an impression basis. And that's what's really happening in some sense. You are trying to optimize the placement of ads for each impression, for each ad slot.

Jason: Comment?

Audience: It seems to me that only half of [question] number 3 is "should the bidders be told what the mechanism is?" I think the other half is "do the bidders understand what the mechanism is, having been told?" or "do [they] know what to do, how to act, having been told what the mechanism is?" And [this seems like a] problem . . . This morning, just before lunch, I thought of these great papers [appearing] in this workshop, on mechanisms [for maximizing?] revenue and so on . . . But the bottom line in each one of those papers, as far as I can recall at this point, was that the mechanism that maximizes revenue is very intricate. Can we use such mechanisms? And if we use such mechanisms, are the bidders going to understand to do the right thing, or do we even have to get them to understand in somehow a deeper way?

Rakesh: And why don't you [let software tell that] to you?

Audience: I might be able to build the software agent, and you want everybody to use the software agent then . . . I don't know.

David: I think a lot of small-size advertisers may or may not understand, but the big advertisers, who are the largest chunk of the revenue, if enough money is involved, they'll figure it out. That means, the incentive is there for them to figure out how to do the right thing.

Audience: You don't think that they'll understand it? They'll just try it out and eventually would figure out what to do?

David: I mean, there's a limit,

Audience: This is a really big question. Does Blue Gene understand it is playing against Kasparov? [laughter]

Audience: The computer scientists have the answer. You build a software agent and you don't care whether they understand or not. [laughter]

Audience: Why don't you build a software agent they'll understand?

Audience: Question 3 taken literally, implies that through the answer it might actually be happening in practice. If you're going out to the field and asking advertisers what they want, what they think they need, some of that has to do with perception. [Another question is that] should the users, the ad readers and so on, of these search [engines] understand the mechanism? Because there is a very mild, default understanding that for organic search results, the good sites make it to the top for some mysterious reason. [For] prime-time TV ads, it might be a good [one] or it might be a bad [one], but it's not an illegal one. It's not something [pretending to be] on TV. Is it important the mechanism we use, if we're talking about the average person, when they're told an analogy they'll understand it. a) they advertise this, yeah these are ads and b) yeah, people would pay to be up there; it's not something like they climb up there.

So isn't it important that the mechanism be transparent to people who use search engines, too?

Michael: According to FCC, the answer's unambiguous: Yes. They have an [advice] that strongly encourages search engines to clearly designate advertisements, and all search engines do comply with this advice. Certainly FCC feels that way, and I have no argument. . .

Kamal: I think that one of the things we want is [that the average advertiser should feel that he had justice with the system]. There are currently, let's say . . . about six million small businesses and about 250,000 are advertising online. And there are potentially a huge number of advertisers who could be there, but they are not as smart as let's say, HP people are. They need to be told, like, very simple things. And if you tell them very simple things, then it means you should be implementing only simple things. Maybe you want to have some freedom to make sure that you can optimize on top of that . . .

David: Yes, you could tell them very simple things ... You are a direct marketer type advertiser. You could say "I want a pay-per-conversion and ... I'd like the search engine to optimize for me ... ", and you know, what keywords to buy, and everything and that's a very simple state ... but yes, the mechanism becomes very complicated

Audience: I have a question, maybe not to the panel, but to the search engine optimizers. Do they get these small users, do they actually use search engine optimizers, that we have to worry about them?

Kamal: You get users that spend a few hundred dollars.

5 Pay-per-what?

Jason: I think we should move on to the next topic. [We seem to be having a] very long panel discussion ... So the next topic is "pay-per-what?" What's the right way to try to charge for sponsored search? In the past it has been pay-per-click, it's been pay-per-impression, and there is talk on pay-per-acquisition, and pay-per-action, through all the spectra of ways to charge, are there right ways to think about it? What are the trade-offs between these ways? etc.

David: I think that the answer to "pay-per-what?" is "Yes". [laughter]. Essentially, you should elicit all of what you are advocating, or a subset of what you are advocating ... Essentially, they [, the advertisers,] should be able to say this however they want. Whatever's important, the direct marketers want pay-per-conversion, big branding-type companies want pay-per-impression, and they should be able to state in whatever way they want, and the search engine can convert it, do the optimization to pick which is best in terms of relevance.

Kamal: You could have them in many combinations, because I have certain values just for the impression, I just want to show my name, and certain values when people visit my homepage, so in that case I should be able to bid for both; I'll pay 2 cents per-impression, but if you give me a click, I'll give you 2 dollars more.

Audience: There could be an issue, [if the advertiser actually wants impressions, but says he wants to pay per click, then the search engine will try to show the advertisement as many times as he wants but does not charge him any thing.]

David: As long as you rank by revenue, you can convert everything back into a cost-per-impression, and rank by that, because they're not getting any clicks. so therefore they're paying for clicks, and so they're going to be dropping in their rankings ... or they can convert back to impression. Essentially they won't be getting any impressions after a while. I mean, this is the same thing as if they paid for conversion, and they're actually lying about how many conversions they're getting ... That's fine too, in the long run, because they'll just start dropping in rankings, or stop appearing ... So, if they want to keep appearing, they have to be telling you the truth about conversions.

Audience: Another question is how you rank.

David: You combine all these things, you rank by ... I mean you have some optimizer which is picking the best way to place things in order to get the most money. And that's essentially like a revenue ranking. It's essentially like converting everything into impression and ranking by that.

Kamal: There is another benefit to let advertisers choose pay-per-what. If advertisers choose whatever they believe is of utility for them, then the chances of fraud, like click fraud or impression fraud are lower. For example, in pay-per-acquisition the fraud is almost zero. The only way you can do fraud is actually buying their product.

Audience: But how do advertisers prove the acquisition?

Kamal: That is a technical question. [laughter] If they are reporting a lower rate of acquisition, then their bid per impression goes down.

Audience: This argument about [self-supporting?] acquisition is very weak. What you're saying is that there's an optimal number for them to report, and there's the number they're actually getting. It will only happen by luck that those two numbers coincide.

Kamal: There are [ways to verify?] acquisition data, but unfortunately I can't share it. [laughter]

David: Actually, the search engine doesn't really they care if they tell the truth or not ... They can report whatever they want, all they care about is revenue.

Audience: Then you basically converge to pay per impression and all those other things are a waste of expression.

David: Well ... presumably, the advertisers find it easier to figure out their value for conversions, and ... so, it's just a convenience for the advertisers, but the search engine doesn't really care ... I mean, assuming their estimates, this is the biggest [assumption?] that they can estimate things very well.

Jason: It's also risk-mitigating, too. You paid for acquisitions, paid for [money you got?] a lot less risky for advertisers.

Audience: It's a little bit of a big jump [from pay per click to pay per] acquisition, I think, because click/impression is controlled by search engines [and acquisition by] customers. There's a big gap there.

David: There is a short term window when they could take advantage.

Audience: On the acquisition issue ... all the industry people, are you [ready to track the entire process and make it part of your mechanism?]

David: Yahoo does have a search optimizer product where they track conversions. So, the merchant puts a little javascript on their page. Yes I think all of them have some kind of tracking and some of them allow bids [per-action?]

Kamal: But some advertisers may not participate.

David: Right.

Audience: And also there's a [correlation?] effect. People discover [something?] and they click some of them, and then [later] they go to the advertiser. So each of these components [has a synergistic effect on the next and has to be taken into account. It's a very complex issue. It is very difficult to keep an account for acquisitions.

Audience: Mixing different payment mechanisms in the same auction seems to be a big [data?] problem. [You have all these different data, and applying different] conversion factors [to them] seems like a losing proposition to me. So, what about trying to finesse this by doing the opposite by having a single criterion auction.

David: I think you still need to convert it, because essentially the search engine is selling impressions. That's what the search engine has to allocate.

Audience: We have data on impressions, we have data on the aggregate income and in fact the query income from the different businesses. The individual bidders don't have that.

Kamal: So, whatever crude estimates you have, we have to normalize certain things, if we want them to compete against each other. I mean, if there's a keyword, and there are only four advertisers, but one of them is bidding on by acquisition, and one of them is bidding on by impression, and one of them by click or something, then they are not competing against each other.

Audience: I see that there's a scalability problem there, if there are a large number of advertisers.

Kamal: There are keywords where the number of advertisers is 4 or 5.

Audience: Another related issue or direction that the market seems to be moving is to increase the amount of targeting. That's definitely increasing the value of the advertisers for the impressions they are buying, but totally fragmenting the market and decreasing the competition on any particular real estate the search engines are selling.

Kamal: But there is a difference that this fragmentation is potentially increasing the value.

Audience: How are the search engines to extract the rent from this increased pie? What mechanisms can be used to price these keywords better?

Kamal: Basically, some mechanism of choosing different reserve prices.

Audience: But you have to set the reserve prices.

Kamal: Yes, they may not be uniform, they may be uniform in certain [cases] but not uniform per click.

Audience: But this is a question that is going to become a very real issue very soon but I don't see a satisfying answer.

6 Other Mechanisms.

Jason: All right, it may be time to move on to our last topic. Maybe [we should] do it rather quickly, since it has gotten pretty late. So the last topic is "other mechanisms", and this is considering, maybe, not running auctions, [and instead] doing negotiations, or [running different mechanisms for different kinds of things instead of a single rule to sell] different kinds of keywords, [e.g.,] the frequent keywords, the infrequent keywords, the tail, etc. Maybe dealing with big players who have these budgets, like HP, and dealing with the small mom and pop shops, maybe there should be different mechanisms for different kinds of people. So that's the last topic.

Kamal: One thing it should provide is, let's ask people to take the [risk]. In that case, the search engines can distinguish between big advertisers and small advertisers. If there is a mechanism in which they can take a risk, maybe big advertisers want to take the risk, because they can absorb the risk, whereas small mom and pop advertisers may want to pay a premium to be risk free.

Rakesh: you pretty much said it Jain. On the less frequently used key words you could imagine an auction and letting people choose, that would give them the competition. You could imagine posted prices for the [static] market, where the customers would be willing to pay a premium to guarantee its position over a certain period. There is no reason why the same mechanism should be used for everyone.

Michael: I will make a prediction, it'll probably not come true but i will make it anyway [laughter]: I think they will move away from auctions. And I think we will move towards pricing mechanism, advertisers will get prices. Why will people move away from auctions? Because once you call something an auction, you are on the hook legally to respect the laws that govern how you conduct an auction, and what you could and couldn't do, there are laws about bidding, But anything could be interpreted as bidding. But nobody tells Walmart how to do the pricing. Well they could respond to the demand. But once you rename something as a pricing mechanism, you pretty much rule out about half of the law suits that could be filed against you. And on the strength of this reasoning, I'd think they'll move away from an auction.

I'd argue that in some sense Google already moved away from an auction because in my mind it is not an auction when the bidders don't know the rule. That is what makes the difference between an auction and a pricing mechanism is that you tell the bidders the rule, that is what makes it an auction.

Michael: [It is] very simple [to do]. You just do find and replace, you just replace the word auction mechanism every where with [pricing mechanism]. Say currently, on their [webpage?] a particular keyword, is running for 11 cents. That is the current going price. That becomes the pricing mechanism. Just do find and replace and you get a pricing mechanism [laughter].

Kamal: Doing find and replace is not answering a mechanism design question, it is answering a legal question.

Michael: I just made a prediction

Rakesh: In the find and replace, you can say replace the word auction by pricing, or you can say replace the word auction by dynamic pricing. Currently we have the bidders bid and they accept the price, but you can imagine the auctioneer sets the price and the bidders accept what they want at that price. The auctioneer can vary the price depending on how many people want the top slot or the second slot. One advantage to that is that if there is enough stability then at the going rate for that slot, bidders may be much happier with that because then there is no need to optimize [for the advertiser], over how much to bid. That's the going rate for the top slot.

Audience: Isn't that variable reserve pricing?

Rakesh: There may be many different phrases to describe the same thing. Auctions can be run when you have the bidders tell you what the price is, and you decide what they get. or you can have the auctioneer announce the prices, and the bidders say what they want, right? They are two sides of the same coin. The advantage of the second is that if there is stability then there are going prices for everything.

Kamal: I'd say there won't be any stability because the number of searches, [no matter how well] you estimate, is not exact. In demand equals supply even if the supply changes a little bit, it can reflect big time in pricing. So how will you manage demand equals supply, if the whole supply is changing.

Rakesh: No, no, that's it, you adjust the prices.

Kamal: You adjust the prices, but that is what this auction is doing, trying to adjust the prices based on demand and supply. People express their demand by mentioning their budgets and it is the price that changes day to day. It can't be the same.

Rakesh: I'm sorry, I am not disagreeing, I'm just arguing that there's some benefit to the auctioneer announcing prices and then the bidders saying what they want, rather than the other way around which is what we currently have.

David: That pricing will have to be quite complicated because for example prices go up in the evening.

Rakesh: The thing is, I think it is easier for Yahoo or Google to bear the computational burden than the bidder. Let the mechanism designer take on the computational burden. Leave as little as possible for the bidder to worry about.

Kamal: That's what the auction is currently do. Just bid, and if [the auctioneer says] that if the prices are below the bid then you will buy it, and if the prices are above the bid you will not buy it. Anyway the advertisers have to give this number. How does it resolve the issue?

Rakesh: I'd argue that there is a qualitative difference. I have to come up with a number, or in the second format, I have to look at a number and agree.

Kamal: But you don't because this number changes every second. You don't, you have to predecide automatically what the level would be. So that's essentially, call a bid.

Audience: [You don't want a situation where] everybody wants to buy the second slot, nobody wants the first.

Audience: more important thing [is they are] perishable products. Compare it with airlines, they are priced, there are mechanisms to price and there is a pricing problem. It's very hard, dynamic pricing is hard.

Rakesh: So the airlines have that pricing

Audience: But the airlines are going bankrupt [laughter]

Rakesh: But that's not the reason

Kamal: You go there, you are not buying a hundred thousand tickets. You just go there, and then they decide whether your true value is above the [price] or below it but you decide it after looking at the price, which is essentially the bid. If you are buying hundred thousand tickets, you can't look at the price and then decide what your true value is, whether you should buy it or not. You have to tell your software upfront that if the ticket price is this much, then buy it, if the ticket price is that much, then don't buy it. OK? And that's essentially you are giving your bid to your software.

Audience: I have a comment. [It's about] search engine [getting into] bid optimization. Today, they [the search engines] are trying to design a bid optimization to [game their own] mechanism. I heard that in the next stage you may have something like Google telling that, OK, give me [your dollars and your preferences] and I'm going to give you the maximum competitive number of clicks per dollar. So you delete the middle industries, which I think so far have [added] value. [They had a role to play in the price going from cents to dollars.]

David: Actually, Yahoo acquired one of these SEOs and that became the search optimizer product.

Audience: If you want to use that, then you [should not] use an auction. You have to use something like [what] I'm proposing, number of clicks per dollar. Because you are gaming your own system, that has more of an incentive issue. If you want to use your own bid optimizer, you have to change the rules. You can't use that auction. That also relates to this comment earlier [about expressiveness]. Right now, the end user does have expressiveness. He can tell the bid optimizer, spend thousand dollars for something today. He can say different things. So the bid optimizer tries to do all those things for him. There is expressiveness this way. If you put the bid optimizer in the black hole ...

Jason: Ahhh, huh, comments? [laughter]

Audience: [Nobody can do efficient pricing?] if they don't control the clearing algorithm. So the clearing point, like Yahoo and Google and MSN, can do more [than someone who is outside the system].

Audience: True, I think right now the efficiency in the system is being determined by the bid optimizers and not by the auction.

Audience: The thing I really want to say is that you don't want to segment the market completely because every time you split the market into submarkets you take a heap off the [efficiency available].

Jason: All right, I think it is time to bring the thing to a close, so I was planning on taking questions that didn't fit in this model but instead we should take dinner [Laughter]. So let's thank our panelists, [Applause] and apologize to our scribes for talking for so long [Laughter]. And hope that at some point we manage to get some notes published.