The problem that we have studied so far captures the essential features of the AdWords problem, but it can be extended in several ways.

We will shortly talk in some more detail about two of these, which relate to the idea of slates or positions, and which relate to the idea of personalization.

Aside from these two extensions, there are also many other issues that Google deals with on a daily basis.

One of these is related to click-through-rates.

In particular, how does Google know the chance that a user clicks on a given ad?

Google does this by analyzing large amounts of user data and building predictive models, similar to those we studied in this class, to predict how often users click on different ads when they're shown with different queries.

Another issue is related to the advertisers.

We saw earlier that the price-per-click depends on how the advertisers place their bids.

So understanding the behavior of advertisers and incorporating this behavior in the optimization model is also another important consideration.

Let's move on to discuss the idea of slates.

In our example with AT&T, T-Mobile, and Verizon, we assume that the search page for each query has space for only one ad.

Now typically, as we saw in Video 1 when we searched for Nine Inch Nails tickets, there's usually space for many ads.

In this case, Google has to decide which combination of ads, or slate, to display with each query.

Although this would seem to be a more complicated problem, it can still be solved using linear optimization.

Before, our variables were defined as x of a given advertiser and a given query.

But now, we would instead define them as x for a given slate and a given query.

So for example, for our wireless service provider example, if we had two spaces on our results page, then for query 1, we'd still have x_A1, x_T1, and x_V1, where here, for example, x_V1 is the number of times that we display Verizon with query 1.

But we would also have x_AT1, x_AV1, and x_TV1.

Here, for example, x_AV1, represents the number of times that we display the slate containing AT&T and Verizon with query 1.

Now, this can become even more complicated as the position of the ad within the slate is important.

For example, ads to the right of the search results might not attract as many clicks as those above the search results.

In this case, we would also consider x_TA1, x_VA1, and x_VT1.

And here, the first ad in the combination is the ad that is placed in the first position.

So, for example, here T-Mobile is placed in the first position for x_TA1 and AT&T is placed in the second position.

We would formulate our objective and our budget and query constraints in the same way as before, but making sure that slates that contain a certain advertiser use up that advertiser's budget.

And slates in a given query counts towards that query's estimated number of requests.

Let's now discuss the idea of personalization.

In addition to the query, Google can use other information to decide which ad to display.

For example, Google might know the geographic location of the user as determined from their IP address.

Google might also know other information, such as different Google searches that the user has conducted, and browser activity on Google's website.

The question then is, how can Google take this into account when deciding which ads display for which queries?

Well, just like slates, we could also incorporate this into our linear optimization model.

Rather than working with queries, we would work with combinations of queries and user profiles.

So rather than having x defined for a given advertiser and a given query, we would define x for a given advertiser, a given query, and a given user profile.

So here, a user profile just describes the type of user.

For instance, a profile might be men aged 20 to 25 who live in Boston, Massachusetts in the United States.

If we had three user profiles, we could name them P1, P2, P3.

And then for AT&T for query 1, we would use x_A1P1 to denote the number of times that we display AT&T's ad with query 1 for a user of profile P1.

The rest of the model could then be easily accommodated for this new type of modeling consideration.

We'll just now summarize the salient points of this recitation.

So, so far, we've studied a small instance of Google's ad allocation problem, where we had just three advertisers or bidders and three queries.

We saw how an optimization solution increases revenue by 16% over a simple common sense solution.

What I'd like you to remember is that in reality, this problem is much larger.

For each query that Google receives on its search engine, there may be hundreds to thousands of advertisers bidding on it.

In terms of dollar amounts, in 2012, Google's total revenue from advertising was over $40 billion.

At this scale, the gains that are possible from optimization become enormous, and I hope this convinces you of the value of linear optimization in online advertising.