Bidding strategies

SEGMATIC

Expert Series #6

The comprehensive guide to setting the right bid every time

Contents

1. A bidding strategy to dominate the auction

A lot of advertisers don't give bidding the time and energy it deserves. We think that's crazy. Sloppy bidding undermines your paid search performance, and it hands money to Google and your competitors.

Bidding is hard to operationalise, but conceptually it's is pretty simple. The aim is to enter bids that maximise return on investment for each click. To do this, you need to understand two things: how much value you get from each click, and how that value will change if you increase or decrease your bid.

Because this requires prediction, and because no two clicks are the same, any bidding strategy centres on five models:

Intrinsic models

1. Conversion rate
2. Cost per acquisition (CPA) target

Extrinsic models

3. Max-to-actual-CPA ratio
4. Optimal price point
5. Portfolio model

Conversion rate and cost per acquisition are both intrinsic to your business. To calculate them, we don't have to factor in the actions of competitors in the auction. We call these intrinsic models. Max-to-actual-CPA, optimal price point, and the portfolio model are all impacted by other actors in the auction, so they rely on feedback loops from the auction. These we call extrinsic models.

We use these models in different ways, depending on the nature of the bidding problem we're trying to solve.

The first three models—conversion rate, CPA target, and max-to-actual—are built on solid data; so you can leverage them heavily, and can use them to make bidding decisions that involve subtle shifts between marginal positions.

However, both optimal price point and portfolio have high margins of error; so they're not really suitable for more nuanced decision-making. The margin of error with these models means they work well when you can observe a clear result from your decision. They work well when there's a swing above or below the fold. But in a contest between marginal positions on the results page, there's too much statistical noise for the price point and portfolio models to be really useful.

Sloppy bidding undermines your performance, & it hands money to Google & your competitors.

These five models work at the keyword level, and then we break each keyword down by device, location, audience, and demographics and arrive at a bid for each possible instance of the keyword.

How does the auction work?

Every company wants its ads to appear at the top of the results page. The higher up your ad appears on the page, the more clicks you'll get. The ad in number one position typically getting about a third of clicks, and a sharp drop-off in clicks for ads that appear below the fold.

Page position is determined by the auction, with competing advertisers entering bids on keywords. If the match logic determines that your keyword matches a particular search term, you enter the auction. That auction determines where your ad appears on the results page.

The higher up your ad appears on the page, the more clicks you'll get.

How does your bid influence the auction?

If you're live on a keyword, and someone enters a search term that Google's match logic decides is relevant to your keyword, your ad will enter the auction.

The position that your ad appears in is determined by two factors—bid and quality score.

Your bid is the amount that you've told Google you're willing to pay for each click on ads related to a particular keyword.

Quality score is a score out of ten that Google assigns to advertisers based on how good they think their ads, keywords, and landing pages are. It's designed to improve users' experience of using Google.

The cost per click, that is the amount that you actually pay for each click rather than the maximum amount you told Google you were willing to pay, is calculated like this:

$$\text{CPC} = \frac{\text{Your quality score}}{\text{The Ad Rank of the bidder below you}} + €0.01$$

If your goal is to improve your page position in the auction, there are many things you can do to achieve this. The most basic one is to increase your budget: if you're willing and able to spend more in the auction, that will yield results.

Really what everyone wants to do is improve performance without increasing spend. That's our holy grail. In this guide, we're going to share our secrets. We'll talk through our approach to bidding strategy, focusing on the predictive models we use to inform bidding. Our aim is to teach you an approach to bidding that delivers the best page position at the best price.

2. A framework for model building

Each of our bidding models breaks the data down to keyword level and considers all of the non-keyword factors that are relevant to our bid, like device and demographics. Then we use those factors to solve for each permutation and arrive at the correct bid.

Building blocks and factors

When you bid on a keyword, you're actually bidding on that keyword plus modifiers for device, location, audience, and demographics.

For each of our models, we begin by breaking the the keyword data down across each of these factors. Once we've done this, what we have is the basic building block of an account.

An example of a single building block might be:

KEYWORD	+SHIRTS
DEVICE	Desktop
LOCATION	Birmingham
AUDIENCE	Non-pixel
DEMOGRAPHIC	Male

Breaking things down to this level produces thousands of possible building blocks for each campaign.

While it's useful to know the conversion rate for the keyword SHIRTS, it's incredibly powerful to know the conversion rate for all the building blocks associated with that keyword.

Getting to this level of granularity isn't easy. Google and Bing don't give us clean data at the building block level. They'll give you an ad group report broken down by demographics, or a report that's broken down by gender, but not one that combines multiple factors.

As a result, this can be somewhat complicated. We have to infer the relationship between these factors as they apply to each keyword.

Translating this analysis into accurate bids isn't simple. Unfortunately, we can't enter an individual bid for each building block. Rather than bidding women up for a particular keyword, we have to apply women as a modifier across multiple ad groups.

And when there's a tension between two or more factors, we have to compromise. Because of this, setting bids for building blocks is a problem that's not perfectly solvable, but you can solve it 99% of the way with our advanced predictive models.

3. Assumptions

Before we get into the models, let's take a moment to look at the assumptions that underpin them, and how we deal with these assumptions.

The two big ones are: assuming that a change in page position does not change conversion rate; and the dangerous assumptions that your models have the data they need to produce actionable insights (and in particular the kind of qualitative insights that tell you that you'll sell more flowers on Valentine's Day).

1. Conversion rate & page position

There's a common misconception among advertisers that conversion rate is the same in every position on the results page. This is generally true, but not always; and the exceptions can really trip you up. If you increase your bid and move to a higher position on the results page, your click through rate will nearly always increase but your conversion rate may in fact decrease.

The main exception to be aware of is competitor advertising, because of a group of people we call non-discerning clickers. These people basically use Google as their address bar. They google something and automatically click on the top result without really reading it. With competitor advertising, the top result on a search for `Pepsi` could bring you to the homepage for Coke, which leads to a lot of unintended clicks.

Conversion rate is not always the same in every page position.

Another interesting exception is certain finance products. People who are shopping for a quick short-term loan will often search, click on the ad in first position, apply for the loan. If they're rejected, they go back to the search results, click on the ad in second position and apply for that; and so on, until they're approved for a loan. In this example, page position has an impact on conversion rate, and a big impact on value.

While these are both interesting exceptions, the fact remains that for 99% of non-brand keywords, changing position does not impact on conversion rate.

2. Conversion rate analysis with small data

The second assumption, which applies across all of our models, is that we have sufficient data to build strong models that lead to good decisions. Most advertisers have lots of data on a couple of keywords, and low levels of data on a bigger set of keywords that are less important to their business.

When we have a lot of data on a keyword, we can predict its conversion rate with a high degree of accuracy. This tends to be true of keywords that have appeared in search terms many, many times. But how do we calculate the conversion rate for a keyword that's appeared only once or twice? There's an art to building strong models with small data.

A common challenge is trying to calculate the relative conversion rate for two keywords for which we have only small amounts of data. We do that through relationship aggregation, looking for a larger set of keywords that share characteristics with the keyword we're trying to understand, and inferring that the average conversion rate of that larger set of keywords will give us a good approximation for the smaller one.

For example, we don't have a lot of data on 'green striped double cuff shirts'. And with good reason: they're not good shirts. But we can assume that the conversion rate is similar to that of 'double cuff shirts' or 'green shirts'.

This involves a logical leap; but when your bidding model if properly integrated with your account structure, you can make this inference with a high degree of certainty.

A lot of advertisers know that the modifier 'online' always increases conversion rate. What they don't know are the millions of ways, subtle and not so subtle, in which other modifiers impact on conversion rate.

Don't call the play too early

A common mistake that advertisers make is to call the play too early, making poor decisions based on an incomplete picture.

Without thorough data analysis, a lot of advertisers are left trying to fill in the blanks with quick and dirty formulas like 'day one conversion numbers are usually about 80% of total conversions', and work from there.

We've spent years developing a bidding model that calculates predicted conversion rates accurately and makes bidding decisions that consistently win the auction without wasting a penny of spend. The model works for all products, across devices, and it even takes account of real-world events. We'll explain how this works in more detail later in this paper, when we get into calculating CPAs.

3. A model that knows the true meaning of Christmas

Quantitative data give us a decent sense of how our conversions are going to play out. But we can make much more accurate predictions by supplementing them with external, non-quantitative data. There are things like the launch of new products, propositions, and seasonality, all of which have a real effect on conversion rate.

Seasonality is the most common example of this, and Black Friday is a good case. As humans, we know that sales of TVs will spike on Black Friday. We can also predict that the n+1 latency curve adjustment for the conversion rate will also spike for TVs, because people click and then buy.

Similarly, as Mother's Day approaches, the latency curve for `canvas print Mother's Day` will be a bit spikier than that of `canvas print`.

It's not always about seasonality. For example, the observed conversion rate for a new product in the days after its launch is always low, and there's a risk that advertisers will over-react to this.

You have to adjust your model with human insights, especially for new models or products.

You have to adjust your model with these human insights—especially if the model is new or this is the first time you've advertised a product that's relevant to the seasonal event.

At Segmatic, our bidding model accounts for these in a couple of different ways. We have different latency curves for different types of product, like suits vs. shirts. We also have different curves for search vs. shopping, and for different devices—mobile tends to have a tighter latency curve than desktop. We even see different latency curves for different nationalities, with Germans being especially decisive.

4. Intrinsic models

How to predict conversion rate

At Segmatic, we've developed a sophisticated and effective bidding model to predict conversion rates. Here, we'll explain the thinking that underpins this model, and how you can apply it to your own bidding strategy.

At its simplest, you look at your past conversion rates and then factor in things you know are going to happen in the future that will impact on future rates.

So you think you know your conversion rate?

Understanding past conversion rates sounds pretty simple, bit it's not really that simple. On it's face, it looks like this:

$$\text{Conversion rate} = \frac{\text{Conversions}}{\text{Clicks}}$$

If you get 100 clicks and 10 conversions, that's a conversion rate of 10%.

The difficulty arises when we understand that conversion rate plays out over a long period of time for thousands of keywords.

At any point in time after a user clicks on your ad, you can say with certainty what your conversion rate is for that click up to that point, but you can only estimate how it will play out in the future.

The sad truth is that we may never know what your conversion rate is for any keyword or ad group. And we have learned to live with that uncertainty.

A user might see your ad, click on it, and immediately make a purchase. Equally, they might click on the ad, think about the product, then set it to one side, and come back and make a purchase weeks or months later. Or they might click, fill out a form with their details, receive a follow-up call from the retailer, and make a purchase following that call.

The conversion latency curve

We can graph the typical rate of post-click conversions, giving us a latency curve that looks like this:

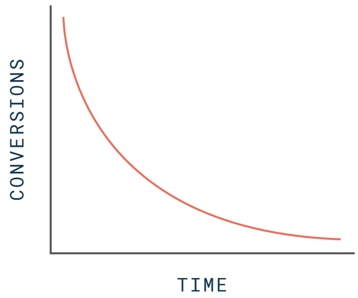

Different products have different latency curves. Pizza has a steep curve, mortgages a more gradual downward slope.

A lot of advertisers make the mistake of looking at their conversion numbers at a particular point in time and making decisions based on that snapshot, rather than trying to understand how conversion rates will play out over time and basing decisions on that.

For example, you might look at conversions on clicks over the last three days, think that your conversion rate was bad but getting gradually better over time and decide to bid down. This results in you entering a bid that ultimately under-values your customers.

Predicted conversion

To address this problem, we introduced the concept of predicted conversion. This a number that combines actual known conversions with a latency curve modifier that's specific to a particular product:

$$\text{Predicted conversion} = \text{Actual conversions} \times \text{Latency curve modifier}$$

This prediction is constantly updating, based on new click and conversion data. Every day, the new data is fed in and the latency curve is adjusted slightly.

Every week or two, we revise the latency curve modifier itself, based on the actual data.

Setting CPA targets

We'll start by looking at how we set CPA targets, and then take a step back and look at two competing CPA philosophies.

Before we even get into the bigger philosophical picture, it's worth noting that setting CPA targets is an art in itself, and very few advertisers get it right. Too many companies have 'a CPA', when in reality they should be setting lots of different CPAs for different products and situations.

If your competitors are better than you at setting CPAs, if their approach is more sophisticated or nuanced, they will bid more accurately, use their spend more effectively, and beat you consistently in the auction.

At Segmatic, we've mastered the art of setting CPA targets. The idea is simple: the more CPAs you have, the more successful you will be. You should have individual CPAs for different products, different types of customer, and times of year.

We explain CPA setting in detail in our guide to setting CPAs (Expert Series #4). At the very least, smart advertisers should split their CPAs up seven ways.

1. Different CPAs for different products

If you sell different products at different price points, you need a CPA for each one. Get this wrong and you will pay too much to acquire low value customers and under-bid for high value ones.

2. Product mix shift

Demand for products fluctuates over time. A company that sells Easter eggs, Halloween costumes and Christmas trees needs different CPAs for each product in each month. The same is probably true of most companies; but it may be a bit more nuanced.

3. High- and low-value customers

Your customer base is a mix of high- and low-value customers, and revenue per customer changes over time. If customers you acquire in July are spending more than those you acquire in January, this should be reflected in your CPAs.

4. Long-term value ratio change

Your CPAs should reflect not just the value of the initial purchase, but the predicted lifetime value that they will bring you. Compare each customer's day one value to their value over one year (and beyond), and factor that into your CPA calculations.

5. Time-limited propositions

If your company offers time-limited propositions, like 20% off all ties in April, you need to factor this into your CPA calculations.

6. Margin

Identify the factors that drive changes in margin—things like exchange rates and commodity prices. Use these insights to predict future changes in margin, then factor this into your CPAs.

7. Whales

Rare but extremely high value customers cost more to acquire, so you have to decide if they're worth the expense. At Segmatic, we've developed techniques for setting CPA targets that systematically acquire whales.

5. Extrinsic models

Max-to-actual-CPC

The rough and tumble nature of the auction makes it difficult to accurately predict and model the relationship between your maximum bid and actual cost per click (CPC).

Even if you never changed your own bid, your max-to-actual-CPC ratio would fluctuate. It doesn't matter how sophisticated or accurate your model is: the moment a competitor changes their bid, or a new competitor enters the auction, your precious model goes out the window.

In reality, there is a difference between your actual CPC and your maximum bid. Actual CPC, the amount that you pay for a click, is usually no more than 95% of the amount that you bid.

Here's why: your Ad Rank is a function of your max bid and your quality score.

```
Ad Rank = Max bid x Quality score
```

But your actual CPC is the next position's Ad Rank divided by your quality score.

$$\text{Actual CPC} = \frac{\text{Next position's Ad Rank}}{\text{Your quality score}}$$

This means that, in a situation where there's a very tight auction where two advertisers with similar Ad Ranks are going up against each other, the max-to-actual-CPC ratio will be about 95%. So if you bid €2.00, you will end up paying pay €1.90

In a loose auction with competition between advertisers with very different Ad Ranks, you might only pay 40% of your max bid. In this case, if you bid €2.00, you would end up paying €0.80.

Our response to this rough and tumble existence is to be smart and quick in responding to the shifts. We integrate with the Google API to get up-to-date data that's broken down hour-by-hour, with a fifteen minute delay. With that, we can continuously runs an API query to check max-to-actual ratios, and adjust our bids to arrive at an accurate target CPC.

Price point

Imagine a scenario in which, for a given building block, each click gets you €15 worth of value. It's a scenario where you have the following perfect data.

You go live simultaneously in these four positions, and you get perfect data on each of them. It might look like this:

PAGE POSITION	CLICKS	CPC	SPEND	CONVERSIONS	GROSS PROFIT (€)	NET PROFIT (€)
3.0	100	1.0	100	10	150	50
1.0	1000	2.0	2000	100	1500	-500
2.8	110	1.1	121	11	165	44
4.0	10	0.8	8	1	15	7

The job of the bidding model is to look at data like this, only bigger, and to figure out optimal price point for our bid on each building block.

If you overpay to rise on the results page, you might lose money.

Depending on your business goals, each of these page positions will be more or less attractive. Some people will pick option 1.0, to maximise gross profit. But if you wanted to optimise for net profit, you'd go for 3.0. But to break even while maximising gross profit, you'd go with 3.8.

There is no objectively optimal price point, there's an optimal price point for any given business goal.

This price point is determined by the level of CPC increase required for each jump in page position, and the number of extra clicks that jump will get you. In its most basic form: if you can jump to a higher position for a small increase in CPC, you can make money; but if you overpay to rise on the results page, you might lose money.

The portfolio model

The portfolio model is more than a simple model: really, it's a philosophy of CPA setting.

It aims to achieve the optimal balance between different CPAs across multiple ad groups. So, if you're bidding on three ad groups, or three thousand, you'll consider each of them relative to each other and set individual CPAs for each.

The main opposing philosophy is CPA maximisation, which aims to achieve 100% CPA across the board.

Let's look again at our price point example, but this time we'll imagine that that there are two different keywords—keyword A and keyword B.

KEYWORD	PAGE POSITION	CLICKS	CPC	SPEND	CONVERSIONS	GROSS PROFIT (€)	NET PROFIT (€)
A	**1.0**	1000	2.0	2000	100	1500	-500
	2.8	110	1.1	121	11	165	44

KEYWORD	PAGE POSITION	CLICKS	CPC	SPEND	CONVERSIONS	GROSS PROFIT (€)	NET PROFIT (€)
B	**3.0**	100	1.0	100	10	150	50
	4.0	10	0.8	8	1	15	7

Now imagine that you want to spend €100 more across both keywords, and you have to decide how much of this extra spend goes to each of the keywords.

Applying the portfolio model to this problem, we would look at the problem holistically: considering the information that we have about each of the keywords and setting different CPAs for each of them. Then, we would use those CPAs to decide how much of our £100 to commit to each keyword.

This is a simple example with two keywords, in reality we do this across hundreds of keywords, breaking the data down into building blocks and then solving for each keyword.

6. Reflexivity

The actions that you take in the auction will provoke reactions from your competitors, which you will have to interpret and respond to.

The nature of the auction means that most of the time you'll want to respond by upping your bid. This force drives up total spend for every keyword, but in our experience it rises until a sort of equilibrium is reached.

Imagine a scenario where four competitors are fighting it out in the auction, resulting in a results page that looks like this:

PAGE POSITION	COMPANY
1	D
2	C
3	B
4	A

The CEO of Company A is frustrated, he wants more conversions, so Company A pushes their bids and by the end of the month they've risen from position 4 to position 2; like this:

PAGE POSITION	COMPANY
1	D
2	A
3	C
4	B

Company A's rise in the results page was the most obvious result of their push.

This in turn led the CEO of Company B to grow frustrated wonder where his conversions have gone. Company B responds by pushing their spend.

Your actions in the auction cause other actors to react.

Our point here is that your actions in the auction cause other actors to react. Every time this happens, your price point model and your portfolio model need to be re-started. This presents a big challenge for both models. Over time, people tend to reach an equilibrium where everyone in the auction is fairly happy: after a period of each company pushing their bids, all the CEOs decide this is a bit expensive, and everyone retreats to a more comfortable position for them.

7. What's the best approach?

With the rise of automated bidding platforms, the auction has become a more dynamic place.

The number of bid changes per day has exploded because it's become so simple for an advertiser to adjust their bid, and for their competitors to adjust theirs in response, and on and on.

For the most part, these bid changes are minor tweaks, but big swings happen too. We've all experienced the sort of CEO we described earlier who's desperate to boost page position.

A lot of advertisers are looking at this more dynamic auction environment and asking whether price point and elasticity modelling is really possible.

The answer to that question is that all five models are important. We have five clubs in our bag—conversion rate, CPA target, max-to-actual, price point, and portfolio—and part of the skill is knowing when to use each of them.

We take a hybrid approach: applying the more solid models for nuanced problems with fine margins; and using the price point and portfolio models for problems with big, clear, obvious outcomes, like the very top of the auction where we're competing with Google rather than our competitors for actual max CPC, and above and below the fold.

Get that right and you can make big wins; and you can make millions of tiny wins that add up to big differences in return on investment.