


brainlabs

Paid Search Account Structures






Your new restaurant

Imagine you have set up a brand new restaurant (*congratulations!*) for which you have big ambitions. You've hired an amazing chef.

You lead the chef into the kitchen on her first day. You've laid out all the ingredients for her, but you've separated them into three boxes. You tell her that you'd like her to design her dream menu, but with a catch: each item on the menu can only make use of the ingredients within one of the boxes - there can be no overlap between the boxes.

Clearly, these constraints are not a recipe for success. They're going to leave your chef with one hand tied behind her back, and will prevent her from conjuring up the very best recipes.



Or






Welcome to the world of optimisation problems

In setting up your new restaurant, you've unknowingly stepped into the world of optimisation problems.

An optimisation problem requires that we find an **optimal solution**, subject to a set of **constraints**. A **constraint** is a limitation on the solution of the problem, which can introduce with our next example. We also have what is called an **objective function**, which is the thing we want to achieve.

Imagine now that you have a factory that can make chairs or tables, and you are constrained by labour and machine time. You make \$40 per chair, and \$30 per table. A chair requires 2 hours of labour and 1 hour of machine time. A table requires 1 hour of labour and 3 hours of machine time. How do you maximise profit?



This is an optimisation problem - the **objective function** is maximising profit, and the **constraints** are the labour and machine time.

And as it turns out, in this case there is an optimal solution, which is 42 chairs and 16 tables, or \$2,160 of profit per day. While this is a very simple example, for most real-world optimisation problems we rely on AI to find optimal solutions, as we can very quickly get into a place where we have hundreds of dimensions to consider.







Welcome to the world of optimisation problems

You can see from this example that constraints are important. They allow us to find the optimal solution within the real world constraints. However, constraints that don't match up with the real world will limit our ability to get to the true optimal solution. A solution within a constrained space may be called a 'locally optimal solution', as opposed to the true 'globally optimal solution'. Where the constraints have been added artificially, we've limited our efficiency in solving the problem for no good reason.

Artificially constrained solution space




True solution space




An artificially segmented solution space will have lots of locally optimised solutions, but they won't be able to get to the true globally optimal solution. This is what our hypothetical chef is dealing with.

Artificially segmented solution space



True solution space





Paid Search is a lesson in communication

Understanding optimisation problems, especially optimisation problems that are solved by AI, is critical for us as PPC practitioners. For instance, most have now accepted that smart bidding solutions achieve better results than manual bidding... but what techniques can we use to make smart bidding work even better for us?

In any paid search account, our **objective function** is to maximise the number of fires that take place of a particular conversion goal. There will also be one necessary **constraint** - either the budget we are willing to spend, or the efficiency target we are required to hit.

However, in many paid search accounts, practitioners end up introducing **additional constraints** that the bid strategies believe they have to work within. These constraints are often communicated by the way we structure our accounts. Sometimes, these constraints are genuinely important for the business; in other cases, they are artificial.

What is our job then, as PPC practitioners? We are **communicators**. The structures and goals we choose communicate to Google what optimisation problem we want them to solve. AI algorithms are very good at solving optimisation problems; but solving an optimisation problem is no good to you if the solution isn't ultimately aligned to benefitting your business.



"In the age of AI, our job as Paid Search practitioners is to communicate our clients' business goals to algorithms via the *structure, goals and budgets we choose.*"





Staying in control?

This is all very well, you might say, but we are Paid Search practitioners, and we like to be in control of the activity we are running because we know our business best. We like segmented structures that allow us to have varied budgets, bid strategies and see segmented data across different parts of our business. We want to decide what ads we show when and monitor all of our keywords separately.

We get it.

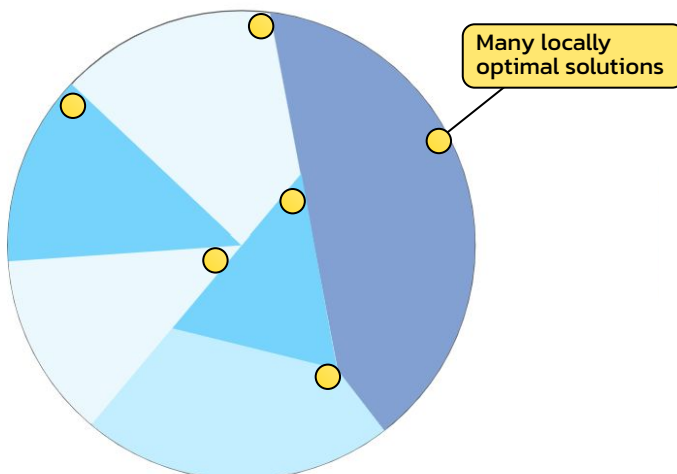
There was a time a few years ago when Brainlabs shouted about the fact that a Single Keyword Ad Group structure was the best way to run Paid Search, because of the control it offered. And many clients and advertisers feel that giving up that granular structure for a consolidated structure also means giving up control.

But we would argue that we don't actually have to give up control. In fact, by learning the language that Paid Search systems speak, we actually gain more influence over the auction than our competitors have.




A granular structure might make you feel like you are in control, but you are actually communicating to the algorithm that you have a constrained system, and stopping it from maximising your potential performance.

Artificially segmented solution space



True solution space





Brainlabs' approach to search structures

So, how do we keep the right level of control over our activity, while giving our AI chef access to a fully stocked kitchen to get us the best possible results?

We need to communicate to bid strategies in ways they understand. Now that we know a little bit about optimisation problems, we know that there can only be two ways we can possibly influence the situation:

Setting good objective functions.

Setting constraints that actually matter - now you have consciously ditched the ones that don't.

An example of this is altering your conversion goal to indicate differing value between products, rather than putting some products into their own campaign with their own budget.

You **do** make more money on boots vs sandals (*a real objective function communicated to the algorithm via your conversion goal*) but you **do not** need to only spend 20% of your budget on sandals (*an artificial non-real-world constraint that is limiting your optimisation problem*).

While we always tailor structures to our clients' businesses, there are **three fundamental principles** that sit behind the structure methodology we choose, that we've seen work well time and time again:

01



Driving a high quantity of data flowing through each RSA

02



Keeping relevance high between our ads and the search queries of our users

03





Directing investment via bid strategy targets, rather than campaign budgets





Show me the data

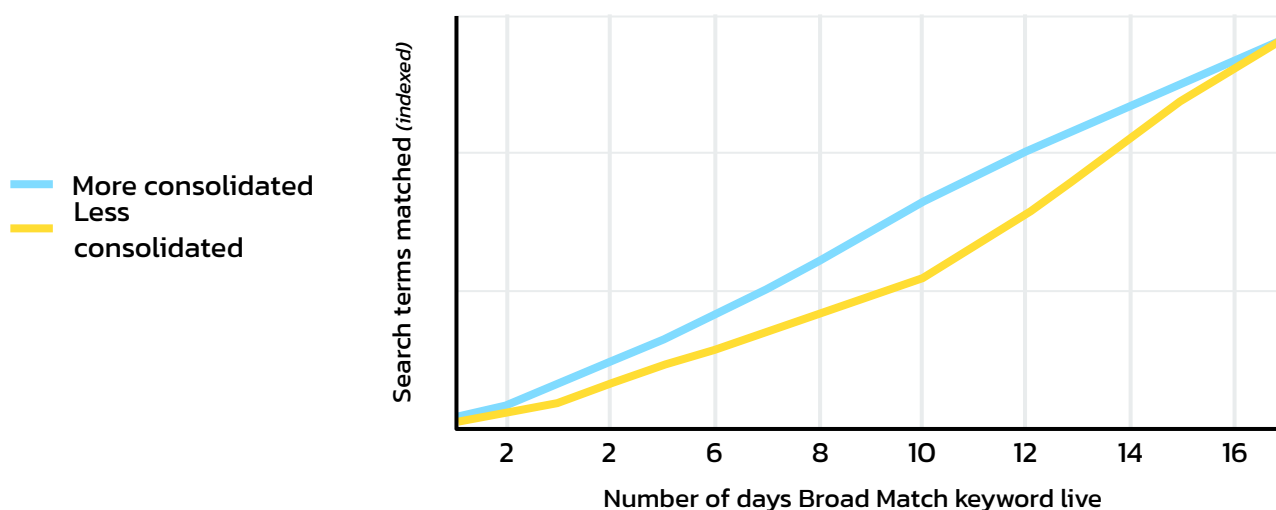
We have a proprietary structure that abides by these three rules we've set out, whilst still allowing us to be cognizant of clients' business goals. We have tested this widely, and it's led to an average efficiency improvement of 22%, when tested against more granular alternatives.



Whilst these results give us a high degree of confidence, it is worth openly acknowledging that, within the topic of account structures, it's difficult to find truly like-for-like data points; when launching a new campaign you tend not to be editing the campaign & ad group structure in isolation, but rather you're likely to combine this with a broader spring clean of the campaign; potentially introducing new match types, new keywords, and refreshed ad copy.

So what underlying factors were driving these results? We've identified two specific ways in which our consolidated structural approach actively improved the optimisation decisions being made by Google's AI.

Firstly, Google's keyword-to-query matching algorithms can learn more efficiently when they have access to the whole solution landscape. This is evidenced by data we found from our own client base, which showed that advertisers using a more consolidated structure saw the number of search terms matching to broad match keywords increase at a faster rate from the day they were launched.







Show me the data

Not only that, but Google's bidding algorithms make more accurate decisions when they are provided with greater volumes of data. This is illustrated by these anonymous charts from some of our biggest advertisers which show that bid strategies with more data going through them often have lower CPAs than bid strategies belonging to the same advertiser but with less data.

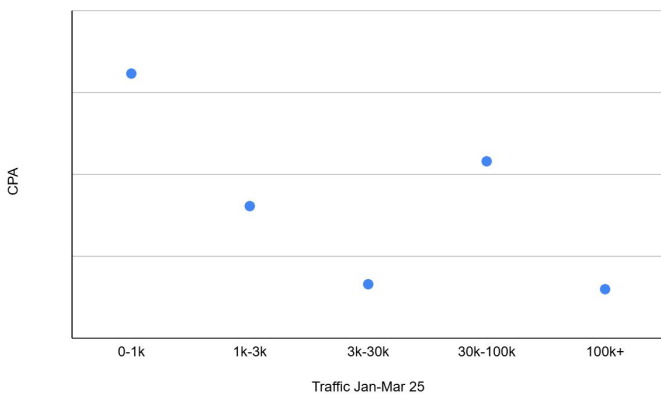
Advertiser 1




Advertiser 2



Advertiser 3



Advertiser 4



In conclusion

Campaign structures are a fundamental part of any Paid Search set up, but they can be elusive to test and understand. It's expected that you will be skeptical of any change where it feels there is less freedom to pull levers.

However, as Brainlabs' CEO Daniel Gilbert once said in 2019:

"The rise of automation doesn't steal jobs, it just changes the game."

While algorithms are doing the job of solving optimisation problems, your job has changed to communicating with them effectively. Our hope is that you understand a little bit more about how algorithms understand structures, and feel empowered to take on your role as an effective communicator of your business.

Lead author



Holly Chetwood

Senior Director, Paid Search

Contributors



Andy Goodwin

Head of Paid Search



James Freeman

Engineering Manager

