

FREE COURSE

Amazon Ads Fundamentals

How Amazon ads actually work, which numbers to trust, how much data you need before you change something, and the weekly optimizations that do most of the work.

By **Cameron Scot**, founder of Merch Jar

Read it online, one lesson at a time, with the free tools linked from each lesson:

merchjar.com/learn/amazon-ads-fundamentals

Contents

1 Why advertise on Amazon

2 Search terms and keywords

3 How a campaign is built

4 Placement adjustments and bidding strategies

5 Which numbers to trust

6 How much data before you decide

7 Your first campaigns

8 The weekly optimizations

9 Advertising books on KDP

10 What comes next

11 Running these rules every day

Glossary

This is the course I wish I'd had when I started running Amazon ads. It covers how campaigns actually work, which numbers to trust, how much data you need before you change something, and the handful of optimizations that do most of the work every week.

Most lessons end with a free tool that runs that step on your own numbers. The tools that read a report do it on your device, so your files never leave your browser.

Why advertise on Amazon

Amazon wants two things. Shoppers who find what they want fast, and as many sales as possible. Almost everything about how products get shown follows from that.

So Amazon puts the products it expects to sell at the top of the results, because that's where shoppers look first. Amazon doesn't publish its ranking formula, but what sellers see over and over is that sales move rankings more than anything else. A product that sells faster tends to rank higher. Ranking higher brings more organic orders, more orders bring more reviews, and more reviews help the next shopper decide to buy. That loop is usually called sales velocity.

The loop runs the other way too. A new competitor or a slow season can slow sales down (so can a product that's just getting older), and then the rankings and organic orders slide with it.

Ads are how you keep the loop moving. They buy the visibility that gets the sales that earn the ranking. That's why an ad can be worth running even when it only breaks even on its own, as long as the organic lift is real.

It also means ads are a long-term investment. It takes time to collect enough data to know what works, and longer for reviews and rankings to catch up.

How to think about it

Strip Amazon ads down and there are two moves. Do more of what works, and stop doing what doesn't.

Easy to say. The hard part is knowing which is which, and that only comes from data. So the whole job is a loop. Test something, collect enough data to judge it, then turn up what's working and turn down what isn't.

You'll hear "it depends" a lot in PPC, and it's true. What works for one product in one category won't for another, and what worked last quarter can stop working this one. Test everything (it's kind of the mantra of PPC). Your own account's data beats any rule of thumb, including the ones in this course.

Search terms and keywords

These get used interchangeably, but they're different things.

A **search term** is what the shopper types into the search bar. A **keyword** is what you, the seller, bid on to show up for those search terms. When someone searches "stainless steel water bottle", that's a search term. Your keyword might be exactly that phrase, or something broader that matches it.

Match types decide how close a search term has to be to your keyword.

- **Exact** shows for searches that match your keyword word for word, plus very close variants like the plural.
- **Phrase** shows when the search contains your words in the same order, with other words before or after them.
- **Broad** shows for your words in any order, plus synonyms and related searches. It can even show on searches that don't contain your keyword at all.

Negative keywords do the opposite. They stop your ad from showing for a search. Negative exact blocks one search, negative phrase blocks any search that contains the phrase.

Indexing is the other half. Amazon reads your listing (title, bullets, description, backend search terms) to decide which searches your product is relevant for. If your product shows up anywhere in the results for a search, it's indexed for it. Auto campaigns lean on this, because Amazon picks the searches from your listing. So the more relevant searches your listing earns, the more places your ads can show.

One opinion of mine from years of doing this. It's easier to double your traffic than to double your conversion rate. Coverage of relevant searches is usually where the growth is.

Try it on your numbers. Paste a list of terms into the [keyword processor](#) to clean it up and switch match types.

How a campaign is built

How Amazon structures campaign data changes how you should build and manage them.

At the top is the **campaign**. Think of it as a container. Inside it are one or more **ad groups**, which are containers too.

Inside each ad group are three sets of data that sit side by side.

- **Targets**, your keywords and product targets. This is where bids live.
- **Ads**, the products you're advertising.
- **Search terms**, the actual searches (or product pages) your ads showed up on.

Each of these breaks out the same clicks, orders and sales a different way. The ads data splits them by product, the targets data by keyword or product target, and the search term data by search. The search term report only lists searches that got a click, so a search you showed up for but nobody clicked won't be there.

A single target can show your ad on hundreds or thousands of searches, depending on its match type. Every ad in an ad group shares the same targets and bids, so they're all eligible for the same searches.

The attribution problem

Those three datasets don't connect to each other. If an ad group has three products and five keywords, you can see which keyword got orders and which product got orders, but not which keyword sold which product.

That's why a lot of advertisers keep one product per ad group, even though Amazon's own advice is to group similar products. Then every keyword and search term result belongs to that product, and you always know what's working for what. Grouping

close variations of the same product (sizes, colors, formats) is a reasonable trade if you don't need them separated.

Placements

There's one more dataset at the campaign level. **Placements** split the whole campaign's results by where the ad appeared.

- **Top of search**, the top row of the first page of results.
- **Rest of search**, everything else in the results, including page two and beyond.
- **Product pages**, product detail pages and other spots outside the search results, like the add-to-cart page.
- **Amazon Business**, placements shown to business buyers.

Your ads can also run on some sites and apps off Amazon. Those show up in the placement report, but there's no adjustment for them.

Placement data only exists for the campaign as a whole. There's no placement breakdown for a single keyword, product or ad group, which matters when you set the bid modifiers in the [next lesson](#).

Try it on your numbers. The [campaign structure planner](#) lays out campaigns and ad groups for your products before you build anything.

Placement adjustments and bidding strategies

You set bids on targets. Two campaign settings then change those bids, for every keyword and every product in the campaign.

Placement adjustments

Each placement can raise your bids by 0% to 900% when your ad competes for that spot. A 0% adjustment leaves the bid alone, 100% doubles it. Adjustments only go up, and setting 0% doesn't remove the placement, it just doesn't add anything.

They're useful because placements convert differently. Product pages often convert lower than search, and can still be profitable at a lower effective bid. Top of search often converts better and can earn a higher one.

Bidding strategies

The bidding strategy tells Amazon what to do with your bid in each auction.

- **Fixed bids.** Amazon uses your bid as it is, plus any placement adjustment.
- **Dynamic bids, down only.** Amazon lowers your bid in real time when it thinks a click is less likely to convert, by as much as 100%. It never raises it. If a keyword gets almost no impressions no matter how high you set the bid, this is one of the first things to check, because down only can be cutting the bid far below what you set.
- **Dynamic bids, up and down.** Same as down only, but Amazon can also raise the bid when it thinks a sale is likely. Up to 100% at top of search. For other placements Amazon's own guides disagree (one says up to 50%, another says up to 100%), so plan for the higher number.

- **Rule-based bidding.** You give Amazon a ROAS goal and it raises bids to try to hit it, up to five times your bid after the placement adjustment. Amazon suggests letting a campaign run at least two weeks before you set the goal, and it doesn't promise to hit it.

How they stack

The placement adjustment applies first, then the bidding strategy works on the adjusted bid. Amazon's own example is a \$1.00 bid with a 50% top of search adjustment. That's \$1.50 for top of search, and with up and down bidding it can go up to \$3.00.

Stack a big adjustment with up and down and the numbers get large fast. A 900% adjustment turns \$1.00 into \$10.00 before the strategy touches it. Audience and video bid boosts, if you turn them on, stack on top of that too. Check every setting together before you raise any one of them.

Try it on your numbers. The [placement report analyzer](#) reads your placement report and shows what each placement earns per click, and which adjustments to lower.

Which numbers to trust

The ads console shows you a wall of metrics. You don't need most of them to make decisions. There's an order to it.

ACOS comes first. It's ad spend divided by ad sales, so a \$25 spend that brings in \$100 of sales is a 25% ACOS. (ROAS is the same ratio flipped, $\$100 \div \$25 = 4$. Use whichever you like, they always agree.) ACOS goes into almost every decision you make.

The only time you don't use ACOS is when you don't have one, because there are no orders yet, or so few that the number means nothing. With one order, ACOS is mostly luck.

Then clicks. With no orders, clicks tell you how much of a chance something has had. That's what you use to decide when to pause a keyword or negate a search term.

Then impressions. With no clicks, impressions tell you whether the ad is even showing.

Everything else (click-through rate, cost per click, impression share) helps you understand a decision, but shouldn't make it. If a call is close, they can break the tie.

Target ACOS

ACOS needs something to compare against, otherwise it's just a number. That's your target ACOS.

A good starting point is break-even. Take what the product sells for, subtract product cost, Amazon's fees and shipping, and what's left as a share of the price is your margin. If your margin is 40%, a 40% ACOS breaks even on the ad sale. Every \$1 of ads brings back \$1 of profit to cover it.

That sounds like you're not making money, and on the ad sale alone you're not. But you're getting the sale and the ranking that feeds the organic side, which is the point of [lesson 1](#). From there, go lower if you want the ads themselves profitable, or higher if you're launching and buying ranking on purpose.

Your target can change by product and over time. Just make sure every product has one.

Try it on your numbers. The [break-even ACOS calculator](#) and the [profit calculator](#) give you the target, and the [ACOS to ROAS calculator](#) converts between the two.

How much data before you decide

The question I get most is some version of “how many clicks before I pause this?”

You want enough data to be confident, but not so much that you’ve paid for a pile of clicks that were never going to sell. Waiting for real statistical certainty would take hundreds of clicks per decision, and you’d run out of budget first.

The answer also depends on what you’re about to do. Pausing a keyword or negating a search term is hard to undo, because once it stops getting traffic it stops collecting data. That deserves more evidence. Changing a bid by 10% is small and easy to reverse, so it needs less.

For bid changes, count orders

Bid changes use ACOS, so the data you need is orders. I want at least two orders before I move a bid, and three or four makes the ACOS a lot more reliable.

For pausing and negating, count clicks

Pausing and negating are for things spending money without orders, so you count clicks.

A simple way to set the number is to start from the lowest conversion rate you’d accept and double it. At a 10% conversion rate, it takes about 10 clicks to get an order on average. Double that and you get 20. So a search term, an exact keyword or a single product target with 20 clicks and no orders has had a fair shot. (For a version based on your actual conversion rate and how sure you want to be, use the calculator below. At 10% and 90% confidence it says 22.)

Broad targets need more. A phrase or broad keyword, an auto campaign target group or a category target spreads its clicks across lots of different searches, so 20 clicks might be one click each on 20 searches. I'd want 40 or more before pausing one of those, and often the better move is to negate the searches that aren't working and keep the target.

Date ranges

When you don't have enough data, use a longer range. If the last 7 days has two clicks, look at 30. If 30 isn't enough, look at 60 or 90. For pausing and negating, look at 90 days or longer, so a target in a temporary slump doesn't get cut.

When you do have plenty of data, use a shorter range, because recent data is more relevant. If a keyword had 20 orders last week, what it did a month ago doesn't matter much. That means different ranges for different keywords, which is more work, and it's also where a lot of the performance comes from.

Leave today and yesterday out of bid decisions. Amazon credits an order to the click for up to 7 days after it (14 for vendors and KDP authors) and reports it on the day of the click, so recent days always look worse than they'll end up.

Try it on your numbers. The [clicks before a negative keyword calculator](#) gives you the threshold for your conversion rate.

Your first campaigns

Start simple and make it more complicated as you go. You want a structure small enough that you'll actually check it every week.

Discovery, then testing

An auto campaign finds the searches. Amazon picks where to show your ad based on your listing, across searches and product pages. This is where your listing's indexing gets put to work. Because it shows up in so many places, keep it efficient by negating searches that collect about 20 clicks without an order.

A manual campaign tests what the auto campaign found. When a search term in the auto campaign gets a couple of orders, add it to the manual campaign as an exact keyword. When a product page (an ASIN) gets orders, add it as a product target. Now that search has its own bid, and you can push it or pull it back without touching everything else. You can copy them over by hand, or use Target Promotion in the ads console to move them for you.

Keep keywords and product targets in separate ad groups in the manual campaign. Amazon requires it, and it keeps the data clean anyway.

Don't cut off the source

A lot of guides say that once you promote a search term to an exact keyword, you should negate it in the auto campaign so the two don't compete. I don't. If the search is still selling in the auto campaign, it's working there, and negating it throws away sales to tidy up a report.

Only negate it in the auto campaign if it's selling there above your target ACOS while the exact keyword does better. Promoting and negating are two separate decisions

with two separate rules.

One product per campaign to start

Start each product with its own campaigns, so its results are its own. Close variations of the same product can share, if you don't need them split apart.

Starting bids

Your first bid matters less than you think. It's your way into the auction so you can start collecting data. Higher bids collect data faster and cost more. Lower bids are cheaper and slower. Amazon's suggested bid is a fair read on what the competition is paying. Start near it and let the data move it.

Budgets are guardrails

A daily budget is there to stop a runaway day. If a campaign hits its budget every day, it's missing traffic while you're limiting it, so control ACOS with bids and negatives and leave the budget with some room.

If money is tight, test fewer things at once (fewer new campaigns, keywords or products you don't know work yet) instead of starving all of them.

Name campaigns so you can read them later

A campaign's name should tell you what's in it and why, at a glance, a year from now. Include the product (ASIN or SKU), the ad type and targeting, the goal, and a short product name. Spell things out rather than using acronyms you'll forget.

B0EXAMPLE1 - SP - Auto - Discovery - Insulated Water Bottle 32oz

B0EXAMPLE1 -

SP - Manual Exact - Testing - Insulated Water Bottle 32oz

Try it on your numbers. The [campaign structure planner](#) builds this layout for your products, and the [target CPC calculator](#) tells you what you can afford to pay per click.

Run this rule every day

Merch Jar Library templates, saved switched off until you turn them on: [Adaptive Budget Management](#) .

The weekly optimizations

Every change you make either expands or contracts. Expanding is raising bids, adding keywords and testing new campaigns, to find more sales. Contracting is lowering bids, pausing and negating, to waste less.

You'll do both every week on individual keywords. The account-level question is which way to lean. If ACOS is well under target and you want more sales, lean toward expanding. If it's over target, lean toward contracting.

Make the changes at the level the data lives. Bids and pauses on targets, negatives on search terms, pauses on ads. A campaign's totals blend everything in it together, so a campaign that looks bad can have great keywords hiding inside it, or one search eating most of the budget that a single negative would fix. You'll rarely need to pause a whole campaign.

The rules below use 30 days for bid changes and 90 days for pausing and negating, unless [lesson 6](#) says otherwise.

Raise bids on winners

Look at	Keywords and targets with at least 2 orders
When	ACOS is below your target by a few points (I use a 3 to 5 point buffer, so 22% or lower for a 25% target)
Do	Raise the bid 10%
Watch out for	Leave anything close to target alone. Set a max bid so a keyword that keeps qualifying can't climb forever.

Lower bids on losers

Look at	Keywords and targets with at least 2 orders
When	ACOS is above your target plus a 3 to 5 point buffer
Do	Lower the bid 10 to 20%, more the further it's off. With a 25% target: 10% for 30 to 40% ACOS, 15% for 40 to 50%, 20% above 50%
Watch out for	Small bids. A 10% cut on a \$0.25 bid is under 3 cents.

Get impressions on targets that aren't showing

Look at	Keywords and targets over the last 7 to 30 days
When	Zero impressions
Do	Raise the bid 10% a week, with a max bid
Watch out for	If a target still has zero impressions after 90 days of weekly raises, the bid probably isn't the problem. Start it over in a new campaign or ad group. Check the bidding strategy too (lesson 4).

Percentages barely move small bids. A 15% weekly raise takes a \$0.25 bid about 10 weeks to reach \$1. Give low bids a floor (anything under a set amount jumps straight to it) or a minimum step of a few cents.

Pause targets that spend without selling

Look at	Keywords and targets over 90 days
When	0 orders and more than 20 clicks (exact and product targets) to 40 or more (phrase, broad, auto and category)
Do	Pause the target
Watch out for	Targets with 1 order get the same treatment at about double the clicks, 50 to 80. On broad targets, try negating the bad searches first.

Negate search terms that spend without selling

Look at	The search term report, 90 days
When	0 orders and about 20 clicks (use your own number from lesson 6)
Do	Add the search as a negative exact keyword in that campaign or ad group
Watch out for	Look for words that keep showing up across bad searches. A negative phrase on the shared word can catch dozens at once.

Promote searches that sell

Look at	The search term report from auto, phrase and broad campaigns
When	2 or more orders. Ask for more (3 or 4) before adding it as a phrase or broad keyword.
Do	Add it to your manual campaign as an exact keyword (or a product target, if it's an ASIN)
Watch out for	Don't negate it where it came from while it's still selling there (lesson 7).

Adjust placements

Look at	The campaign's placement report, 30 days
When	A placement's ACOS is above target, or below it with room to grow
Do	Move that placement's adjustment down or up 10 to 25 points (from 50% to 30%, say)
Watch out for	Adjustments can't go below 0%. If a 0% placement is still too expensive, lower the bids instead.

Pause products that don't sell

Look at	The advertised products in each ad group, 90 days
When	A product has plenty of clicks and no orders while others in the ad group sell
Do	Pause that product's ad
Watch out for	Check the listing (price, stock, reviews, the main image) before blaming the ad.

How often

Weekly is a good rhythm by hand. Pick a day, run the rules, and you're most of the way to well-run campaigns. Smaller changes more often respond faster to what the market is doing, which is where automation earns its keep, because nobody wants to do this every morning.

Try it on your numbers. The [negative keyword finder](#) lists the searches past your threshold, the [n-gram analyzer](#) finds the words that keep showing up in bad searches, and the [search term report analyzer](#) shows wasted spend and searches ready to promote.

Run these rules every day

Merch Jar Library templates, saved switched off until you turn them on: [Dynamic Bid Management](#) , [Pause underperforming keywords and targets](#) , [Search Term Waste Elimination](#) , [Product Ad Waste Elimination](#) .

Advertising books on KDP

If you advertise books on KDP, everything above applies, with one big problem. The ACOS in your ads console only counts order sales. The console does show estimated page reads and royalties from Kindle+ (Kindle Unlimited until October 2026), but ACOS leaves them out, so a book that earns most of its money from page reads can show a terrible ACOS while making money.

That's why authors get told to ignore ACOS and just compare royalties to ad spend. That's too blunt to manage keywords with. You can fix the number instead.

Blended profit and blended ACOS

You can't add order sales and page-read royalties together as they are. Royalties are money in your pocket. Order sales are the list price, and you only keep a share of that.

So first turn sales into profit with your average margin. Total royalties from orders divided by total list price of those orders gives you that margin.

Average margin $\times$ ad sales + estimated page-read royalties = **blended profit**

Ad spend $\div$ blended profit = **blended ACOS**

Because margin is already taken out, break-even is 100%. A \$100 spend that brings in \$100 of blended profit pays for itself. Below 100% the ads make money, above it they cost money.

Page reads after the attribution window

Page reads only count toward your ads if they happen within 14 days of the click. A reader who borrows your book and starts it next month, or finishes it in week three, earns you royalties the console never credits to the ad.

Nobody can tell you exactly what share of reads happens in those 14 days. It depends on genre, length and how fast your readers binge. A romance audience might finish 80 to 90% within two weeks. A fantasy reader with a long backlog might only get 20 to 30% done. Pick a percentage you believe and lean conservative (higher).

Estimated royalties ÷ your 14-day read share = **extended royalties**

If half of reads happen in the window and the console shows \$100 of royalties, that's about \$200. Use the extended number in blended profit.

Series read-through

Someone buys book one from your ad, likes it, and goes on to buy book two, three and four on their own. The ad gets credit for book one only.

Your series read-through rate fixes that. Divide book two's orders (or page reads) by book one's, book three's by book two's, and so on. If 80% of book one readers buy book two, then 100 ad orders of book one are worth about 80 more orders of book two, and some of those carry on to book three.

Add those up and you get your **ad impact multipliers**. If every ad order of book one leads to two more orders down the series, each ad order is really worth three, a 3x order multiplier. Page reads get their own multiplier the same way, and the extended-reads adjustment stacks on top.

None of this is exact. You can't measure cross-series buys or how much organic ranking the ads bought you. But a reasonable estimate beats ignoring ACOS completely, and it tells you which keywords are actually earning their spend.

Try it on your numbers. The [KDP read-through calculator](#) works out your series multiplier, what a book one reader is worth and your break-even ACOS on book one.

What comes next

Once the basics run every week, the next step is more structure, added only when a product has the sales volume to justify it. Complex structures without the data to feed them just split your clicks thinner.

That's what the Advanced course covers.

- **Splitting auto campaigns by target group** (close match, loose match, substitutes, complements), so each gets its own bids and placement adjustments.
- **Lottery campaigns.** Your whole catalog in one campaign at low bids, to catch cheap sales.
- **Category campaigns,** to find new product targets.
- **Search term ASIN targeting (STAT).** Targeting the products that rank at the top for a search, often cheaper than bidding on the search itself.
- **Brand defense.** Your ads on your own product pages, so a competitor's ad isn't the one there.
- **Single keyword campaigns,** for your best keywords, so placement adjustments and bidding strategy apply to one keyword at a time.
- **Moving keywords between match types** to find new variations, then bringing the winners back to exact.
- **Branded and non-branded traffic,** and why mixing them hides how the ads really perform. The [branded vs non-branded ACOS tool](#) splits them for you now.
- **TACOS,** total ad spend against total sales, to see whether the ads are feeding the organic side. Try the [TACOS calculator](#).

Running these rules every day

Every optimization in [lesson 8](#) is a rule. A condition, an action and a time window. Doing them by hand once a week works, and it's how I started. It gets hard when you have hundreds of products, or more than one account, or you want the changes made every day instead of every Monday.

That's what Merch Jar does. You write each rule as logic you can read (or have your AI write it with you), preview exactly which keywords it would change, and save it switched off. You turn it on when you're happy with it, and from then on it runs every day, the same way every time.

The [Core Six](#) in the Merch Jar Library covers most of [lesson 8](#).

Rule in this course	Library template
Raise and lower bids toward target ACOS	Dynamic Bid Management
Pause targets that spend without selling	Pause underperforming keywords and targets
Negate search terms that spend without selling	Search Term Waste Elimination
Pause products that don't sell	Product Ad Waste Elimination
Budgets as guardrails	Adaptive Budget Management
A proven keyword that suddenly stops getting impressions	Impression Recovery Boost

Promoting searches that sell is handled by Smart Campaigns, which turns converting search terms into keywords and product targets.

Every plan has a 14-day free trial. [Start your free trial.](#)

Glossary

ACOS (advertising cost of sale). $\text{Ad spend} \div \text{ad sales}$. Lower means the ads cost less per dollar of sales.

ROAS (return on ad spend). $\text{Ad sales} \div \text{ad spend}$. ACOS flipped.

TACOS (total ACOS). $\text{Ad spend} \div \text{total sales, ads and organic together}$.

Impressions. How many times your ad was shown.

Clicks. How many times shoppers clicked your ad.

CTR (click-through rate). $\text{Clicks} \div \text{impressions}$.

CPC (cost per click). $\text{Ad spend} \div \text{clicks}$.

Conversion rate. $\text{Orders} \div \text{clicks}$.

RPC (revenue per click). $\text{Ad sales} \div \text{clicks}$.

AOV (average order value). $\text{Sales} \div \text{orders}$.

ASIN. Amazon's 10-character ID for a product.

Search term. What a shopper typed into the search bar.

Keyword. What you bid on to show up for search terms.

Match type. How closely a search has to match your keyword (exact, phrase or broad).

Target. Anything you bid on. Keywords, product targets, category targets and auto target groups.

Product target. Bidding to show your ad on a specific product's page.

Category target. Bidding to show your ad across a product category, optionally narrowed by price, rating or brand.

Negative keyword. A search you've told Amazon not to show your ad for.

Indexing. Whether your product shows up for a search at all, based on your listing.

Placement. Where the ad appeared (top of search, rest of search, product pages or Amazon Business).

Top of search impression share. The share of top of search impressions you won out of those you were eligible for.

Target ACOS. The ACOS you're aiming for. Break-even (your margin) is a good start.

Blended ACOS (books). $\text{Ad spend} \div \text{blended profit}$, which combines profit from orders and page-read royalties. 100% is break-even.

KENP. Kindle Edition Normalized Pages, how Kindle+ (formerly Kindle Unlimited) pages read or listened to are counted and paid.