

A weekly PPC review you can act on.

Find the data problem before changing the account. Then turn checked evidence into a short decision list.

Output: one review, up to three evidence-linked exceptions, the questions that matter and a human decision log. No automatic bids, budgets, negatives or bulk uploads.

1. Agree the scope and question

Select one account, marketplace, currency and reporting period. Record the campaign IDs and the question you want to answer. Keep brand, discovery and other campaign objectives visible; they may not share the same target.

2. Gather matching exports

- Keep original campaign and search-term exports unchanged. Record file names, dates, generation time, time zone, filters and row counts.
- Record attribution window, exact sales-field definition and whether the report includes advertised-only or broader attributed sales.
- Define the grain: daily rows, period totals, target rows or another actual report level. Do not mix overlapping views.
- Bring the owner's goal, approved review criteria, economics, stock/listing issues and recent changes. Unknown context stays unknown.

3. Decide whether the data is ready

Check blanks, zeroes, invalid values, missing campaigns, duplicates and reporting maturity. A partial period or unsettled attribution is provisional. A tool that cannot read all selected rows must say so. Conflicting scope or unresolved overlaps block performance conclusions.

Give SKILL.md only authorized data in your chosen AI tool. No account password or customer/order personal information is needed.

Calculate once. Reconcile before judging.

A plausible ACoS is not enough. Keep the source rows and the denominator.

4. Resolve duplicates with evidence

Define uniqueness from the real report grain. The same search term in two campaigns is not a duplicate. Exclude a copied row only after its origin is confirmed, and keep the original plus a record of the decision. Do not silently deduplicate or replace missing values with zero.

5. Reconcile the two views

Campaign and search-term files describe overlapping activity. Do not add them together. Reconcile each campaign to the source, then compare term metrics only where scope and definitions permit it. Explain coverage differences; do not force impression totals to match.

6. Use weighted totals, not averaged rates

- $\text{ACoS} = \text{total spend} / \text{total attributed sales} \times 100$.
- $\text{ROAS} = \text{total attributed sales} / \text{total spend}$.
- $\text{CPC} = \text{total spend} / \text{total clicks}$. $\text{CTR} = \text{total clicks} / \text{total impressions} \times 100$.

Use a calculator, spreadsheet or code for arithmetic. Sum first; divide second. A zero denominator is N/A with a reason, not 0%. Missing is not zero. Keep purchases/orders distinct from units. Round only displayed results.

7. Prepare no more than three decisions

For each exception show source rows, observed values, a labeled interpretation, missing context and a proposed human next step. Check relevance, data maturity, economics, inventory and recent changes before proposing account edits.

No universal cutoff. Do not invent a good ACoS, a stop-after-X-clicks rule, a sales forecast or a profit claim. Attributed sales are not profit, and attribution is not proof of incremental sales.

Practice on a problem worth catching.

The included fictional fixture plants a duplicate that changes the diagnosis.

First pass: \$120 does not reconcile

Campaigns show \$420 spend and \$1,600 attributed sales. The raw term file shows \$540 spend and the same sales. Stop. Row S6 repeats S4. The fictional owner confirms an accidental copy, so exclude S6 from calculations without changing the original file.

Corrected result

320 clicks, \$420 spend, \$1,600 attributed sales and 32 purchases. ACoS is 26.25%; ROAS is 3.81x; CPC is \$1.31. Campaign impressions are 35,000; retained term impressions are 32,500. The exercise explicitly accounts for the 2,500 omitted zero-click impressions.

Three questions, not three automatic changes

- S4: \$120 spend, 80 clicks, no attributed sales. Is the shopper intent relevant, and what review tolerance did the owner agree?
- S3: \$180 spend / \$600 sales = 30% ACoS. Does this fit the objective and matched contribution economics?
- S5: 20 clicks, \$20 spend, no sales. What is the learning goal, and is the data mature enough for a decision?

Finish the review

Independently recompute the headline and exception values. Record the owner's decision and who may implement it in decision-log.csv. Agree a next review date only if needed. The skill does not make live changes or start monitoring.

All values and products are fictional. See worked-example.md for the full answer and the supplied Python exercise for reproducible arithmetic. It is not a general Amazon importer.

Official references checked October 20, 2026: [Amazon Ads: ACoS and ROAS definitions](#); [Amazon Ads: search-term report help](#). Verify actual report definitions in the account; access may require sign-in.