RESEARCH · VARIATION CONFIDENCE FREE

Which child on a variation listing is actually selling?

How Profit Zones scores every child in a variation family using signals that belong to that child alone, and why we chose those signals over the obvious ones.

White paper · Version 1.6.2 · November 2026

The short version

Every child on a variation family shows the parent's sales rank, so rank can't tell a seller which color or size is doing the selling. Variation Confidence scores each child on its own evidence and gives it one of four tiers.

Three signals earn points: Amazon's **bought in past month** badge on the child, the child's **share of the family's recent reviews**, and how fast its **offers turn over**. Five patterns that mark a fading child take points away. The net score sets the tier, and the family's leader is always High, because a listing with any sales has at least one child driving them.

The method follows four rules. It only measures the child. It treats missing data as unknown, never as zero. It resists Keepa data glitches. And every threshold is published in Help, generated from the same table the code scores with.

The score, per child

PART	POINTS
Bought in past month	up to +60
Share of family reviews	up to +25
Offer turnover	up to +15
Marks against	up to −75
Net score	**−75 to 100**

Tiers: **60+ High**, **35+ Medium**, **15+ Low**, **under 15 Very Low**.

3 signals

Each read from the child, never the parent.

6 rejected

Popular metrics tested on a real 24-child family and left out.

1 token

Per uncached child, and nothing is spent until you approve.

THE PROBLEM

The parent's rank hides the children

Amazon ranks a variation family as one product, and Keepa reports that parent rank on most children. Twenty-seven colors can all show the same rank even when three colors carry the whole listing.

What Keepa shows. Illustrative: four colors, one rank, the parent's.

CHILD	SALES RANK
Black	#1,204
Pink	#1,204
Navy	#1,204
Seafoam	#1,204

Amazon does put a "bought in past month" badge on each child, and it's the best clue available, but it only appears from 50 purchases a month and comes in wide steps. Every child under 50 shows the same blank, and two children in the same step look identical.

Sellers usually handle this in one of two ways, and both cost money. Some skip the listing because they can't tell which child sells, and leave a profitable color on the table. Others spread a small buy across five colors to be safe, then pay storage on the four that sit while the one that mattered sells out before their shipment lands.

> The question isn't "is this listing good?" It's "which child is driving it?"

To answer that, a signal has to be measured on the child itself. That one requirement ruled out most of the numbers sellers usually look at.

Four ground rules

Each rule came from a real failure, either a test on a live family or a field report during the build.

R1. Child-level only

A signal that repeats the parent's number on every child can't separate them, however good it looks on the parent.

R2. Unknown is not zero

Missing history earns no points and no penalty. Data we don't have is never treated as bad news.

R3. Glitch-resistant

Keepa counts sometimes spike or dip for a few days. Medians and reset detection keep a data hiccup from moving a tier.

R4. Published in full

Every threshold is in Help, generated from the same table the code scores with. Tests fail if the two ever disagree.

What we left out, and why

We tested each candidate on a real 24-child family (a drinkware listing, September 2026) before choosing. These six looked useful and didn't hold up. Each one either repeats the parent's number, confuses a side effect with demand, or answers a different question.

Sales rank: the parent's number

Keepa reports the parent's rank on most children, so every child gets the same value. It can't separate them.

Offer count: misleading

One child in the test family had 13 sellers and had gained 5 reviews in 43 months. A high count can just mean sellers are stuck with inventory. Only movement in the count says anything about sales.

Lifetime review share: age-biased

An old color keeps its reviews for years after demand moves on, and a child launched last year looks weak even if it's the current bestseller. We measure recent gains instead.

Price vs. sibling median: wrong measure

It correlated well: every badged child was at or under the family median, and everything 18%+ over was dormant. But price describes the offer, not demand, so it stays out of the score.

Buy Box: a different question

Who wins a Buy Box is a competition question, and Buy Box Winner History covers it. It doesn't tell you whether the child sells.

Keepa CSV export: no history

Four exports, including Graph Range = All with the History boxes ticked, were all snapshots. Review and offer history only come from the API or Keepa's local cache.

What survived. Three signals passed every test: each is measured on the child, reflects purchases (directly or with a short lag), and can be read from Keepa history without scraping. Their weights follow how directly each one reflects a sale: 60 for Amazon's own badge, 25 for reviews, which confirm it with a lag, and 15 for offer turnover, the weakest signal but the only one that still works below the badge floor.

POINTS FOR · SIGNAL 1 OF 3 · STRONGEST

Bought in past month: up to +60

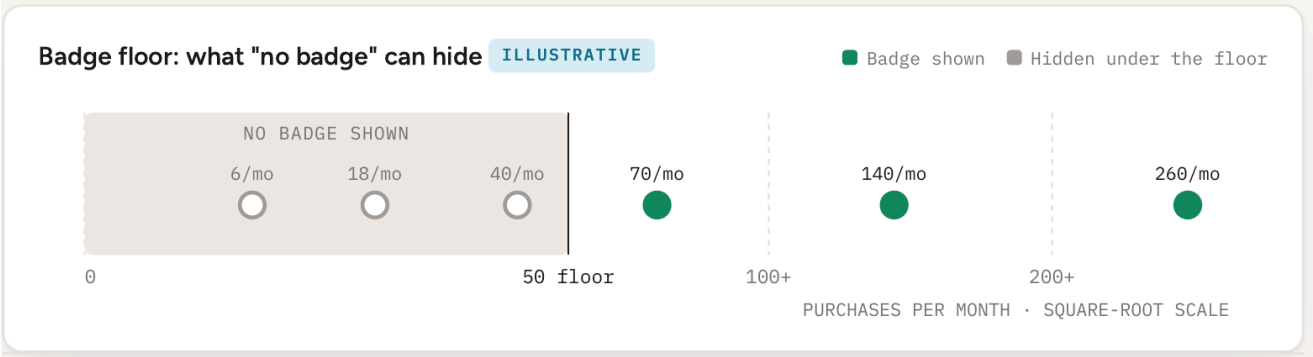
This is Amazon's own "bought in past month" badge, read from the child and not the parent. It's the only number in the model that Amazon states directly about purchases, which is why it carries the most weight: the three steps are worth more than the other two signals combined.

BADGE ON THE CHILD	POINTS
200+ bought	60
100+ bought	50
50+ bought	40
No badge	0

Why a missing badge isn't a penalty

The badge doesn't appear until 50 purchases a month. A blank means "under 50," not "zero." Treating a blank as a failure would wipe out the middle tier of children selling 15 to 49 a month, which is often exactly where a smaller seller wants to be.

From the test family. Navy had no badge, but 251 reviews and 5 active offers. A rule that read "no badge = not selling" would have written it off. Under this model it scores 0 on this signal, takes no mark against, and earns its tier from reviews and turnover.



Badge floor: what "no badge" can hide. Illustrative. Purchases per month on a square-root scale.

What the steps mean in practice

The steps match Amazon's own badge buckets, so the score never claims more precision than the badge gives. A child showing "200+" gets 60 whether it sells 210 or 2,000 a month. Separating those is the job of the other two signals and, ultimately, of your own sourcing math.

When the badge has lapsed

If your 3-Step velocity reads a period, not today, a badge seen anywhere in your chart range still counts, labeled with its age, and counts less as it ages: full points to 30 days, half to 60, a quarter after. Ridgeline, a color in the test family, had 200+ 74 days ago and none today, so it earns 15, not 60. A live badge keeps its own points.

LAPSED BADGE, LAST SEEN	POINTS KEPT
Up to 30 days	Full
31 to 60 days	1/2
Over 60 days	1/4

POINTS FOR · SIGNAL 2 OF 3 · CONFIRMS WITH A LAG

Share of the family's recent reviews: up to +25

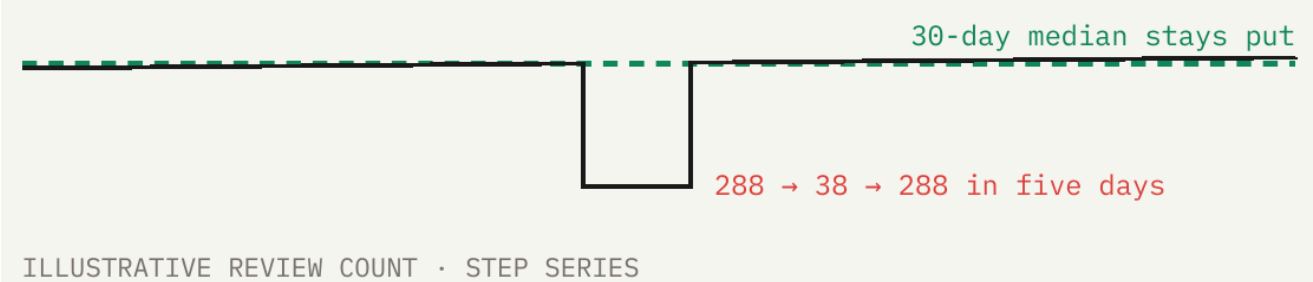
We count the reviews each child gained over the last **6 months** and divide by the whole family's gain for the same period. Buyers review what they bought, so the child collecting most of the new reviews is most likely the one selling. A share, rather than a raw count, lets the same scale work for a family of 5 and a family of 50.

SHARE OF THE FAMILY'S 6-MONTH REVIEW GAIN	POINTS
15%+ of family	25
8%+	18
3%+	10
1%+	5
Under 1%	0

Window: 180 days. Medians: 30 days. Reset: a drop of more than 50% and at least 20. Settle: 7 days.

Why medians, not single readings

Both ends of the window are 30-day medians. In February 2026 one child's count dropped from 288 to 38 and back to 288 over five days. A single reading would have scored that as a collapse. The median doesn't move.



Illustrative review count, step series. 288 to 38 to 288 in five days; the 30-day median stays put.

Why reset detection

Amazon sometimes splits pooled reviews out to each variation. On 2026-07-19, every child in a 27-child family dropped from about 21,300 to its own count (11 to 851) on the same day. Without a guard, every child would have scored "reviews falling." A drop of more than half, and at least 20 reviews, now counts as a reset. Only history after it is measured, starting 7 days later once the new counts settle.

Why all or nothing

Share is scored only when every child's review history was read. Otherwise one previously cached child could be "100% of the family" out of a total made up only of itself. A child that was read and simply has no review series gained 0.

POINTS FOR · SIGNAL 3 OF 3 · WORKS BELOW THE FLOOR

Offer turnover: up to +15

Over the last 90 days, we add up every change in the child's new-offer count, up or down, and divide by its average offer count. Sellers selling out and restocking is inventory that turns, and inventory that turns is inventory that sells.

Under 0.5x a child earns nothing here, in any family. Above that floor the points come from where the child sits in its own family's range, from 15 at the top down to 5 at the bottom. Children that tie share the middle, and a family's only mover is its top.

TURNOVER, AGAINST THE REST OF THE FAMILY	POINTS
Top of family	15
Upper range	12
Mid range	9
Lower range	6
Bottom	5
Under 0.5x	0

The floor is absolute. Everything above it is read against the rest of the family.

Why a ratio, and why against the family

The first version counted changes directly, but parents can carry 50 to 100 offers. A 70-offer child moving by 5 isn't turning over (5 ÷ 70 = 0.07x). A 1-offer child going 1 to 0 to 1 turned over completely (2 ÷ 0.5 = 4x). Dividing by the average puts both on one scale.

Fixed thresholds then failed for a different reason. On one 24-child family the turnover ran from 1.3x to 60x, so a top band opening at 1.0x paid the 60x mover what it paid the 1.8x one. On an eight-child tool listing the best child reached 3.9x. Position inside the family reads both correctly: 5x is a strong mover in a quiet family and an ordinary one in a churning family. A child that hits zero offers and recovers inside the window is noted as back in stock.

POINTS AGAINST · FIVE MARKS

The patterns of a fading child

Points for only show what a child does right. Five patterns subtract points, and each needs enough history to back it up. A short or missing series never triggers a mark: unknown is not zero.

How a child has no offer: an ASIN isn't deleted when its last seller leaves. Amazon keeps it in the family as "Currently unavailable," and Keepa still lists it: a color you can click but can't buy. Confirmed by eye on King Crab, Big Wave Blue and Sharptail Taupe: reviews flat for 12+ months and offers gone or not moving.

MARK	WHAT TRIGGERS IT	POINTS	SAFEGUARD
No listed offer	No seller has offered the child for 30+ days	−40	A shorter stockout isn't a mark; once it fires, the tier is Very Low regardless
No recent reviews	Zero reviews gained in 6 months	−15	Fires only when history covers the full 6 months since any reset
Reviews falling	30-day median below the median 6 months ago	−10	Needs 60+ days of history; a dip that recovers in 30 days never counts
Stale offers	Turnover under 0.15x in 90 days	−10	Skipped for a child at zero offers right now, which is a stockout
Never took off	Tracked 24+ months, under 10 lifetime reviews	−10	An unknown review count never counts as zero

THE VERDICT

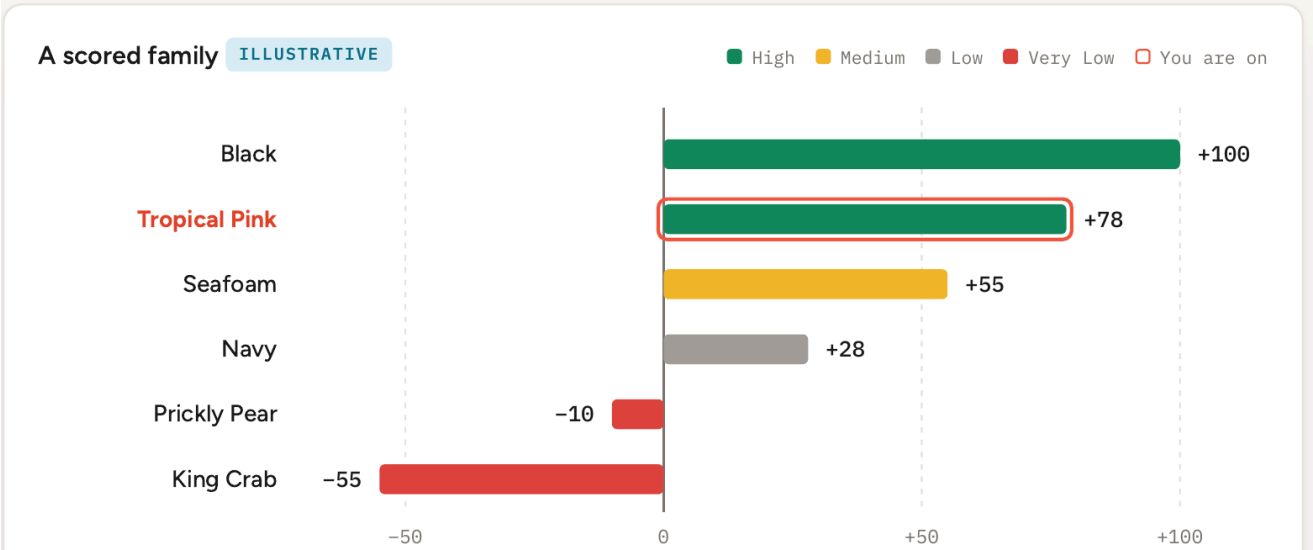
Netting to a tier, and why the leader is always green

Points for minus points against gives the net score, which can go negative. The net sets the tier.

NET SCORE	TIER
60 and up	High
35 to 59	Medium
15 to 34	Low
Under 15	Very Low

The family leader rule

The tier cutoffs are fixed numbers, so a family with thin review data can end up with nothing in green. That misses the point: a listing with any sales has at least one child driving them, and this feature exists to show which one. Any child **within 5 points** of the family's best net, above zero and still offered by at least one seller, is **High** no matter where its number lands, and its row says so. A family with no evidence at all reads Inconclusive, and nothing is promoted.



A scored family, illustrative. One bar per child, from zero to its net score, ranked highest first. An earlier design stacked points for and against on each bar, which gave every row two endpoints and made a +55 bar look like it ran past its own score, so we replaced it. The breakdown lives in the hover and in the table below.

Read it like a seller. Black and Tropical Pink are carrying this family. Seafoam sells with a badge but is a supporting color. Navy has no badge but still turns over, so it's a candidate for a small test buy. Prickly Pear would be a storage bill, and King Crab has no seller offering it at all.

RECEIPTS

Every score explained, as the product writes it

CHILD	NET	TIER	WHY
Black	+100	High	+60 Amazon 200+ bought +25 reviews 31.0% of family, last 6 months +15 offers turned over 1.3x in 90 days
Tropical Pink You are on	+78	High	+50 Amazon 100+ bought +18 reviews 12.4% of family, last 6 months +10 offers turned over 0.7x in 90 days
Seafoam	+55	Medium	+40 Amazon 50+ bought +10 reviews 4.2% of family, last 6 months +5 offers turned over 0.2x in 90 days

CHILD	NET	TIER	WHY
Navy	+28	Low	no badge: under 50 a month, not a mark against +18 reviews 9.1% of family, last 6 months +10 offers turned over 0.6x in 90 days
Prickly Pear	−10	Very Low	no badge: under 50 a month, not a mark against −10 offers stale, turnover 0.1x in 90 days
King Crab	−55	Very Low	no badge: under 50 a month, not a mark against −40 no listed offer for 30+ days −15 no reviews gained in 6 months

COST · KEEPA TOKENS

What a scan costs, and why it asks first

Scoring a family needs history for every child, and Keepa charges per product. A 24-color listing is a 24-token lookup. Variation Confidence shows the estimate, including how many children are already cached, and waits for your OK. Nothing is spent until you approve, and scoring uses your own Keepa API key.

1 token

Per child for the family fetch: offer history, badge and tracking date.

2 tokens

Per child with review history on, since Keepa bills review data as a second token.

0 tokens

For cached children, and for reopening the family from any child page for 24 hours.

1. **Estimate.**
   "Score 24 variations? 18 tokens, 6 cached"
2. **Approve.**
   You confirm before any token is spent.
3. **Score.**
   Tiers, chart, review ring and receipts.
4. **Reuse.**
   Free on every child page for 24 hours.

The cache stores the raw inputs, not the scores, so if a rule changes, reopening a cached family re-scores it under the new rule at no cost.

What the score can't tell you

Every model has edges. These are ours, stated plainly. Variation Confidence answers one question well: which children in a family are doing the selling. It deliberately doesn't try to answer the questions around it.

Demand, not profit

A High child can still be a bad buy at the wrong cost. Profit Zones and your break-even math still decide that.

Less to go on below the floor

Children under 50 a month score 0 on the strongest signal, so reviews and turnover do all the work, and both need history.

Reviews stand in for purchases

They follow sales with a lag and at different rates by category. That's why the share is worth 25 points and the badge is worth 60.

Only as deep as Keepa's history

A newly tracked child, or one whose review series Keepa never recorded, scores on what exists and takes no penalty for what doesn't.

Source of truth. Every number in this paper matches the extension's own rules table as shipped in version 1.6.2. The same table generates the Help page, and tests fail if Help and code disagree. Example families are illustrative and computed with the published rules.

> **This is our own scoring system, kept simple on purpose.** If it works for you, use it. If it doesn't, tell us what you think works better. Custom weights are planned for a later release. Until then the rules stay fixed, fully published, and one click away in Help. Variation Confidence is free in Profit Zones: open any listing in a family of two or more and click the Variations pill.