

You are not competing with one seller. You are competing with *three*.

What 54,779 geographic price scans across four months revealed about how the Buy Box actually works, and why both standard pricing strategies leave money on the table.

96.8%

Of products showing geographic price variation, the share where **more than one seller** holds the Buy Box across different parts of the country at the same moment, on the same ASIN.

The question nobody asks

Ask an Amazon seller how they price and you will get one of two answers.

"I price at the Buy Box." You check the current Buy Box price, match it or shade under it, move on.

"My repricer handles it." You set a floor, a ceiling, a rule set, and let software chase the box all day.

Both answers share a hidden assumption so basic that almost nobody examines it: **that there is one Buy Box price**. One number. One competitor to beat. One market.

We spent four months measuring whether that is true. It is not, and it is not close.

Why we went looking

This started as a cost audit, not a study. We were spending more than expected on the infrastructure behind PATH Profit Zones' geographic scanning, and went in to find waste.

We found waste. We also found we were sitting on a dataset almost nobody else has, or at least that nobody has shared publicly: the same products, priced simultaneously, in 18 metropolitan markets across the United States, captured repeatedly over four months.

What we measured

WINDOW	TOTAL SCANS	DISTINCT ASINS	MARKETS / SCAN	MARKET PRICES
117 days	54,779	28,664	18	~986k

Each scan captures the **anonymous Buy Box**, meaning the offer a first time, not logged in shopper in that ZIP code would see, along with the winning offer's fulfillment type and delivery promise. All 18 markets are sampled inside the same scan, so we compare prices at the same moment rather than across time.

That last point matters. Most price variance analysis is really measuring *time* variance and attributing it to geography. We hold time constant and vary only location.

Every figure in this paper is an aggregate count of products and prices. No account, seller, or customer information is attached to any of it. We are counting how prices behave, not who looked at what.

Your repricer is only doing its job on about half your catalog

23,164 of 54,779 scans (42.3%) showed measurable price variation across markets. At the product level, 12,423 of 28,664 ASINs (43.3%) varied at least once.


Turn that around and it becomes uncomfortable.

A repricer that chases one national Buy Box number is potentially doing useful work on the 57% of products where one national number is real. On the other 43%, the number it is chasing is real in one place and irrelevant everywhere else.

Nobody sold you a repricer with that caveat attached. It is not that the software is broken on those products. It is that it cannot see the thing that matters, so it confidently does the wrong arithmetic.

FINDING • MAGNITUDE

Most of the time it is a material piece of your margin. Sometimes it is the entire margin.



Half of the products that vary do so by less than a dollar, which is easy to read as noise.

It is not noise. Seventy cents on a \$25 product at 12% net margin is roughly a **quarter of your profit on that unit**. That is not a rounding error. That is the difference between an item that carries its weight and one that does not, and it is invisible to you today.

Then there is the top of the distribution. A quarter of variable products move by **two dollars or more**, one in ten by **four dollars or more**, and the largest single spread we recorded was **\$222**.

Now put that next to something every seller already knows about their own catalog: roughly 20% of your ASINs generate roughly 80% of your profit. If any meaningful share of that 20% also sits in the top quarter for geographic spread, this stops being an optimization exercise. It becomes the margin that decides whether the P&L is black or red, and you are currently leaving the decision to a tool that cannot see the variable.

FINDING • COMPETITION

The finding that changes how you think about competition

This is the one we did not expect. For every scan showing price variation, we checked whether the *same* seller was holding the Buy Box in every market, or different ones.

In **96.8%** of cases, more than one seller held the Buy Box across the country at the same moment, on the same product.

The median product had **three different Buy Box holders** across the 18 markets simultaneously. The most we recorded on a single product was **eleven**. In 3.3% of cases, **Amazon itself** held the box in some regions while a third party seller held it in others.

Read that again, because it takes the floor out from under both standard strategies.

When you **price at the Buy Box**, you are pricing against one offer, in one location, and then applying that price nationally. In 96.8% of variable products, whoever you just matched is not your competitor across most of the country. You may have calibrated your entire national price against a seller who wins in one region and loses everywhere else.

When your **repricer chases the Buy Box**, it is chasing a single number that is perfectly real somewhere and misleading almost everywhere else. It reacts to one competitor and drags you toward the most aggressive seller in whichever region it happens to be observing from, including regions where you were already winning at a higher price.

Which raises a question for your repricer vendor: **where does it observe the Buy Box from?** We have not seen a repricer disclose that. If the answer is a single location, then on the 43% of your catalog that does vary, it is reporting a regional price as a national truth.

There is a fair objection to that, which is that a tool might sample several locations rather than one. It moves the question rather than settling it, and the follow up is the more revealing one: **how many locations, and what does it do when they disagree?**

Our own subset testing, described further down, says the sampling has to be wider than most people would guess. Fewer than ten markets systematically understates the spread, and a six market sample missed it entirely in 17% of cases. So a tool sampling two or three cities is not meaningfully closer to the truth than one sampling a single city, and it may be more dangerous, because it produces a number that looks reconciled.

Then comes the part no amount of sampling solves. Whatever it observes, a repricer has to collapse it into one price to act on. When regions genuinely disagree, and on the 43% they do, there is no single national number that is correct in more than one of them. Sampling more locations makes the measurement better and the underlying problem more visible. It does not make one price fit the country.

The software may not be bad. It is **limited**, and far more limited than most sellers realize, because the limitation is invisible from inside the tool.

Regional structure is real, and it does not tell you what you need to know

Amazon operates a small number of US delivery regions. Our data agrees: **99.9% of scans resolved to eight or fewer distinct prices** across all 18 markets, and the markets that agree with each other most are exactly the clusters you would predict.



Seven of the 153 market pairs we track, measured across the full 117 days. The top six are the closest agreeing pairs. Charlotte and San Francisco, shown in amber, are the least agreeing pair of all 153.

So regional structure is visible in pricing. Here is the trap, and it is the reason we sample many markets instead of one per region:

No market pair agreed more than 82.2% of the time. Not one of the 153 pairs agreed above 99%. At the other end, Charlotte and San Francisco agreed on the Buy Box price only **24.3%** of the time. Three times out of four, those two cities are simply different markets.

Inside a single delivery region, the Buy Box holder still changes roughly **one time in five**. Two cities that Amazon serves from the same regional network, priced at the same instant, disagree about who wins the sale on one product in five.

We want to be careful about why, because we measured that it happens and not what causes it. The candidates are fulfillment center placement, meaning whose units physically sit closest to that ZIP, plus stock depth and seller performance metrics. Delivery speed is a known Buy Box input, so the likeliest

story is that within a region several sellers can promise similar delivery and something else breaks the tie. **We capture the delivery promise alongside every price and every winner, which makes this measurable rather than speculative, and it is the obvious next question to put to the data.**

What we can say now is the practical part: **you cannot infer a region's Buy Box from one ZIP code inside it.** As far as we can tell, nobody in this industry is looking at that, so we are clearly stating it.

What this means for how you price

Everything above is measurement. What follows is reasoning from it, meaning hypotheses to test against your own catalog rather than guarantees.

1. **Stop treating "the Buy Box price" as one number.** On any product with real volume, the useful question is not "what is the Buy Box price" but "what is it where my inventory actually ships from and to."
2. **Inventory placement is a pricing input, not a logistics detail.** If your units sit in fulfillment centers serving regions where you win at a higher price, a single national price may leave margin on the table exactly where you are strongest.
3. **A competitor who beats you in one region may be irrelevant in another.** Before you drop price to beat somebody, check whether they win anywhere you actually sell. With a median of three different holders per variable product, there is a real chance they do not.
4. **The 57% that do not vary deserve less of your attention.** That is where a conventional repricer already works. Spend your time acting on the 43% where it cannot.
5. **Watch the top quarter.** A regional strategy has a real cost, and it is attention: pricing these by hand instead of letting a repricer do it for you. On products moving two dollars or more across the country, that attention is cheaper than

the margin an automated repricer gives away chasing one national number it cannot see correctly. Below that tier, let the repricer run.

Every one of those requires seeing the spread in the first place. That is what the Buy Box Spread feature in **PATH Profit Zones Premium** was built to show you, and as far as we know it is the only place you can currently see it.

How we know this is reliable

What we are confident in

All figures come from direct observation of live Amazon pages, not modeling or extrapolation. All 18 markets are captured inside the same scan, so geography is the only variable. The findings are also **stable over time**: measured across the full 117 days, the share of scans showing variation moved by two percentage points and the median spread by five cents against a 30 day sample. This is not a strange month.

What we sample

Eighteen metros is a sample, not a census of American ZIP codes. We tested how small that sample could get before it lied to us, by exhaustively evaluating every possible smaller subset. Fewer than ten markets systematically *understates* variance. A six market sample missed the spread entirely in 17% of cases, including one product where a \$222 spread was completely invisible. That is why we do not cut corners on coverage.

Where our evidence stops

We measure the anonymous Buy Box, which can differ from what a logged in shopper with purchase history sees.

Whether acting on this makes you a specific amount of money remains an open question. Any margin figure we quoted today would be a guess on our part, and we would rather wait until it is not. We are measuring exactly that now, using a

round robin pricing strategy that prices for the best available outcome in each market, and we will publish the result the same way we published this.

What surprised us

We expected the variance to concentrate in a few volatile markets. It does not. Every one of the 18 markets was the high or the low price in **83% to 86%** of variable scans. The variance is spread evenly across the country, which is both why it is hard to detect cheaply and why most sellers never notice it at all.

The short version

Roughly 43% of Amazon products carry a different price depending on where the customer lives. When they do, **96.8% of the time more than one seller is holding the Buy Box in different parts of the country simultaneously**, a median of three, sometimes as many as eleven.

Pricing at "the" Buy Box, or letting a repricer chase it, both assume a single national competitor that usually is not there.

The sellers who understand their regional position, meaning where they win, where they do not, and where their inventory actually serves, are playing a different game.

Scans captured between 9 April and 4 August 2026 across 18 US metropolitan markets. Figures are aggregate counts of products and prices, reproducible from the underlying scan record. Full market list, the complete 153 pair agreement matrix, per market extreme frequencies, and the delivery promise analysis are available separately.