RESEARCH · AUGUST 2026

One Shot

What 39,000 listings and 25 years of MLS data say about pricing Arizona's luxury market

Matthew Walsh

Russ Lyon Sotheby's International Realty · Scottsdale, Arizona

martawalsh.com/research/one-shot/

THE HEADLINE FINDINGS

This study analyzed every residential listing ever entered in the Arizona Regional MLS with an asking price of $2 million or more — 39,539 listings spanning 1998 through August 2026, resolved into 29,322 distinct selling campaigns. It is, to my knowledge, the first published full-history, data-mined study of pricing outcomes in Arizona's upper market. Nine findings:

01 **Nearly half of all campaigns fail, and failure climbs with ambition.** Across 25 years, 48% of $2M+ selling campaigns ended without a sale. The gradient is smooth and relentless: 44% at $2–2.5M, 52% at $5–7M, 56% at $7–10M, and 65% above $10M. Even in the strongest seller's market in Arizona history, a quarter of campaigns died.

02 **"Priced right, it sells in days" is false at every price above $2 million.** Only 13% of genuine market sales went under contract within 10 days; the median was 80 days. Even in the most liquid tier studied — $2–2.5M — fewer than one sale in five contracted inside 10 days, and the median was 50 days. Above $5M the median runs 117 to 134.

03 **Days on market and price achieved are tightly, negatively linked.** Sales contracted within 10 days achieved a median 99% of original ask. The curve then steps down without exception: 96% by a month, 91% by ninety days, 87% by nine months, 80% past a year, 68% past two. Across 15,321 market sales the rank correlation between days to contract and share of ask achieved is −0.57 — one of the strongest relationships in the dataset.

04 **But the discount is taken in public, not at the negotiating table.** Sale price as a share of the *final* asking price — the number on the sign when the buyer actually appeared — is nearly flat at 93–95% whether the home had been on market five weeks or two years. The entire days-on-market penalty is extracted through visible price reductions before the offer is ever written. Sellers imagine defending their price against a hardball buyer; the record shows the market takes the discount from them directly, cut by cut. The mirror image holds for buyers: a stale listing carries no extra negotiating leverage — its discount is already in the ask.

05 **Price cuts do not rescue broken campaigns — at any size.** Campaigns that never reduced sold 55% of the time. Campaigns that reduced — by 5%, 15%, or 30% — all sold at a statistically flat 48–50%, with no dose-response whatsoever. Relists after failure are flatter still: cut nothing, 45% success; cut more than 20%, 43%. There is no reduction size that reliably repairs a campaign, because the reduction arrives after the launch audience is gone and the market-time meter has already repriced the listing. And when a campaign does break, the record is equally clear on the two ways to pay: conceding within the campaign reached the same total discount in roughly one-third the market time of cancelling, resetting, and relaunching — a reset restores the MLS clock, not the market's memory.

06 **Chasing the market costs eight times the calendar.** Properties that failed at least once before selling spent a median 397 cumulative days on market; properties that sold in one campaign spent 52. Within a single market era, the chase surcharge measured roughly three additional points of the original ask — on top of the time.

07 **A failed first campaign is a coin flip on ever selling — and recovery usually means dropping a tier.** Of properties whose first campaign failed, roughly half ever recorded an MLS sale afterward. And recovery routinely happened *below* the price band the seller started in: the chase does not merely trim asking prices; it demotes properties across tier boundaries.

08 **Sellers get roughly a five-month window in which success outruns failure — shorter at the top.** Market-wide, success stays ahead of failure through the first year. Above $4 million the race inverts: dead campaigns overtake successful ones between day 90 and day 180, and by one year the luxury ratio is 52% dead to 35% sold.

09 **The top of the market is increasingly a private market.** Roughly one closing in fifteen at $2M+ was pre-arranged — under contract within a day of listing, the signature of pocket sales and after-the-fact comp entries. The share rises monotonically with price: 6.5% at $2–2.5M, one in ten at $7–10M, one in seven at $10M+.

The synthesis is a different way of thinking about price than either the industry's gospel or the seller's instinct: **above $2 million, an asking price is not a prediction of value. It is a survival strategy.** The buyer for a distinctive property arrives on their own schedule — typically within two to four months, longer at the top. The asking price determines whether the listing is still credible when they do — and it is, on this record, functionally the only lever a seller controls.

The ladder

29,322 CAMPAIGNS, 1998–2026, BY ORIGINAL ASKING PRICE

ORIGINAL ASK	CAMPAIGNS	FAILURE RATE	MEDIAN DAYS TO CONTRACT	UNDER CONTRACT ≤10 DAYS	PRE-ARRANGED CLOSINGS
$2.0–2.5M	8,123	44%	50	19%	6.5%
$2.5–3.0M	7,255	49%	84	12%	5.8%
$3–4M	6,823	47%	102	11%	6.1%
$4–5M	2,919	48%	99	11%	8.0%
$5–7M	2,430	52%	117	8%	7.4%
$7–10M	1,034	56%	126	7%	10.4%
$10M+	738	65%	134	7%	14.3%

One table, four stories: mortality rises with price; the wait lengthens with price; the fast sale becomes rare with price; and the private market grows with price. Everything else in this study is the detail behind these four gradients.

Why this study exists

Two pieces of conventional wisdom dominate pricing conversations, and they point in opposite directions.

The first is the agent's gospel: *price it right and it sells fast*; *sold comparables are the truth.* The second is the seller's instinct: *my home is unique*; *comparables can't capture it*; *records get set by people who asked more.* The seller's instinct contains a genuine insight — if sold comps were binding, prices could never rise, and every record-setting sale in this dataset was, by definition, an ask above anything that had sold before.

Both camps argue from anecdote. Arizona has excellent market analytics — the Cromford Report has tracked supply and demand here for nearly two decades — but no one had mined the full MLS record to answer the seller's actual questions: *What are my odds*? *What does time cost*? *How large a premium does the market actually ratify*? *What happens to the sellers who overreach*?

So I pulled the entire record and counted.

Data and methods

Source and universe. Arizona Regional MLS (ARMLS), all residential listings with an original or final list price of $2,000,000 or more, list dates 1998 through August 2026, retrieved in five price-banded exports and deduplicated by MLS number: 39,539 unique listings. The dataset is statewide; the large majority of listings sit in the Scottsdale–Paradise Valley–Phoenix core, with the remainder concentrated in Sedona, Flagstaff, and Prescott. Fields include list date, under-contract date, close date, cancellation date, original and final list price, sold price, and parcel number. Analysis was performed programmatically using AI-assisted tooling directed by the author; all classification rules and thresholds were fixed before results were computed on the cleaned data.

What was excluded, and why. Honest failure statistics require removing listings that were never really for sale in the ordinary sense:

Speculative and to-be-built listings (4,130 listings, 10.4%). Builder catalog entries, renderings, and pre-construction listings follow a different lifecycle — serial cancellation and relisting as construction progresses is normal, and many close off-MLS on the builder's own paper. Including them inflates failure rates for reasons unrelated to pricing behavior. They were identified through construction-status fields, remarks language, and year-built dates postdating the list date. (The most-listed parcel in the entire dataset, with 57 listings, is a builder machine of exactly this kind; it appears nowhere in the statistics or records below.)

Pre-arranged sales (1,024 closings, 6.7% of all closings). Listings that went under contract within one day of listing never faced the market: pocket sales, off-market deals, and comp entries recorded after the fact. They are excluded from time-to-contract statistics — counting them would falsely flatter the "sells in days" narrative — and reported instead as Finding Nine.

Data errors (60 listings). A recurring typo class — an extra digit producing, for example, a $44.9 million original ask on a home that listed and sold near $4.4 million — was quarantined by rule (original ask ≥ 5× the same listing's final price).

The campaign correction. ARMLS continues a property's cumulative days on market across any relist that begins within its reset window — currently 45 days off-market, and 90 days for most of this study's span — so a quick cancel-and-relist produces a fresh MLS number and fresh listing alerts, but not a reset clock. Linking listings by parcel number surfaced thousands of relist events; more than a quarter occurred within 30 days of the prior listing's end, over half by the same agent — a marketing refresh of a continuing effort, not a new decision to sell. Any relist within 30 days was therefore merged into a single *campaign* — 34,143 clean listings resolving into 29,322 campaigns — a rule consistent with, and stricter than, the MLS's own continuity convention. The headline failure rate is robust to the rule: merging at 45 days yields 47.1%, and at the historical 90-day window 44.2% — "nearly half" survives every defensible definition. The relist-gap distribution itself contains a finding: gaps cluster in visible bumps just past the reset thresholds — at 45–59 days and again at 90–104 days — the signature of sellers deliberately waiting out the cumulative-DOM window before relisting. Genuine clock-resetting exists in this market; it is done by waiting, not by quick relists, and it is one more measure of how much sellers believe visible market time costs them.

The censoring correction. An earlier version of this analysis, run on the $4M+ tier alone, found that only one-third of failed first campaigns ever sold. Extending the data to $2M revealed the artifact: sellers who failed in one tier and succeeded in a lower one had been invisible. The corrected figure — roughly half, where the ladder can see the full descent — replaces it throughout, and the same logic implies the $2–3M recovery figure remains understated for the identical reason (its descent exits the dataset at $2M). I flag this openly: truncated data flatters whatever story it accompanies.

What this study cannot see. Sales that never touched the MLS are invisible. Recent campaigns have had less time to resolve. Sub-$2M outcomes are outside the frame. Findings on price reductions and negotiation outcomes are observational: reduction depth correlates with initial mispricing, so those findings describe what happened, not a controlled experiment.

Finding one: failure is the base case

Of 29,322 campaigns reaching a terminal status, 47.9% ended in cancellation or expiration rather than a sale, rising along the ladder from 44% to 65%. The pattern holds within every market era, which matters, because it means this is not a story about bad markets — it is a story about the tier itself:

CAMPAIGN FAILURE RATE BY ERA OF LAUNCH

ERA OF LAUNCH	CAMPAIGNS	FAILURE RATE
1998–2007	4,874	64%
2008–2012 (crash)	3,566	72%
2013–2019	5,598	55%
2020–mid-2022 (surge)	4,289	26%
mid-2022–2026	10,995	38%

Read the surge row again. In the most frenzied seller's market Arizona has ever recorded — bidding wars, cash migration, record velocity — one campaign in four still failed, and nearly one in two at $10M+. The upper market's default state is that listing a home does not sell it.

Sellers do not experience it this way, and neither does the public. Failed listings disappear from the record consumers see; sold listings become comps, tour stops, and news stories. The upper market is a survivorship-bias machine, and every pricing conversation that begins from sold comps alone has already excluded close to half of all attempts.

Finding two: the ten-day myth

Among 14,249 genuine market sales, 13% went under contract within 10 days. About 28% contracted within 30 days; 53% within 90. The median was 80 days — and 50 days even in the fastest, most liquid tier studied.

This is the finding that cuts against our own industry's favorite sermon, and the reason is structural, not motivational. Below roughly $2 million, buyers are continuously present, so price is the only gate and a correct price clears quickly. Above it, the buyer population thins with every step up the ladder — the specific buyer for a specific estate *arrives*, at a frequency the seller does not control. A correctly priced home in this market is not slow because something is wrong. It is waiting, and waiting is the mechanism. The ladder's middle column measures it precisely: the wait lengthens from 50 days to 134 as the ask climbs, because the arrival rate falls.

What a launch does control is the one demand spike a campaign ever gets. A new listing aggregates every buyer who accumulated during the preceding months into its first weeks of exposure. The launch is not when the home should sell. It is when the home meets the largest audience it will ever have, and the asking price decides what that audience does.

Findings three and four: time prices — but not where sellers think

The market does not merely delay overpriced homes. It charges them, on a schedule, and the full curve deserves to be seen whole. Among 15,321 market sales:

PRICE ACHIEVED BY DAYS TO CONTRACT

DAYS TO CONTRACT	SALES	SOLD PRICE AS % OF ORIGINAL ASK	SOLD PRICE AS % OF FINAL ASK
2–10	2,200	99.1%	99.0%
11–30	2,594	96.0%	96.2%
31–60	2,533	93.3%	94.8%
61–90	1,661	91.4%	94.0%
91–120	1,332	90.0%	93.5%
121–180	1,914	88.5%	93.3%
181–270	1,456	86.8%	93.1%
271–365	768	83.4%	92.1%
1–2 years	735	80.4%	91.8%
2+ years	128	68.4%	88.9%

The left ratio falls relentlessly — the rank correlation between market time and share of original ask is −0.57. The right ratio barely moves: from the second month onward, buyers pay 92–95% of whatever the current asking price is, whether the home has been listed five weeks or two years.

Read together, the two columns dismantle the negotiation folklore of both sides. Sellers picture the risk as a hardball buyer beating them down at the table; in fact the table is where prices are *most* stable. The entire cost of overpricing is extracted earlier, in public, through the visible reductions that connect an unratified original ask to a final ask the market accepts — and the eventual buyer simply pays the standard haircut off wherever the cuts landed. Buyers, meanwhile, picture a stale listing as leverage; in fact 300 days of market time buys an offer no deeper, relative to the current ask, than 45 days does. The discount they imagine extracting was already conceded, in price cuts, before they arrived.

It is worth being precise about causation, because the left column is routinely abused in listing presentations. Slow sales did not close lower *because* they were slow; both the slowness and the discount flow from the same source — an original ask the market declined to ratify. The homes that sold near ask in 10 days were not lucky or better marketed. They were priced where the arriving buyer, whenever they arrived, could say yes.

Finding five: price cuts don’t rescue campaigns

If overpricing is the disease, the industry’s prescribed cure is the price reduction. The record says the cure barely works.

Campaigns that never reduced their price sold 55% of the time. Campaigns that reduced sold 48–50% of the time — *regardless of how much they cut*. Five percent, fifteen, thirty: the success rate is statistically flat across every reduction depth. Relists after a failed campaign show the same dead line: relist at the same price or higher, 45% success; relist more than 20% lower, 43%.

The observational caveat is real — deeper cuts correlate with worse initial mispricing, so this is not a controlled experiment. But the complete absence of any dose-response is the finding. If reductions repaired campaigns, bigger reductions would repair more of them; they don’t. The mechanism is the one Finding Two identified: the launch audience is the largest a listing ever meets, and it has already rendered its verdict by the time the first cut arrives. The reduction is broadcast to a smaller audience, alongside a market-time meter that Finding Three prices, and the combination cancels most of what the concession was meant to buy.

The operational conclusion is uncomfortable and unavoidable: the mental model “we’ll adjust if we have to” assumes the adjustment works. It mostly doesn’t. The launch ask is not the first move in a flexible sequence. It is, functionally, the decision.

The head-to-head: cutting in place versus resetting and relaunching

Every seller in a struggling campaign eventually faces the same fork: reduce the price where you stand, or cancel, let the days-on-market clock reset, and relaunch at a new number. The record can compare the two paths fairly by conditioning on the total concession the market ultimately demanded — among properties that sold, those that got there through in-campaign reductions in a single campaign (2,265) versus those whose first campaign died and that relaunched after a genuine reset gap (2,611):

SAME TOTAL CONCESSION, TWO PATHS

TOTAL HAIRCUT FROM ORIGINAL ASK	CUT IN PLACE: MEDIAN MARKET DAYS	RESET & RELAUNCH: MEDIAN MARKET DAYS	RESET & RELAUNCH: CALENDAR, FIRST LIST TO CONTRACT
5–10%	101	346	726
10–15%	116	308	597
15–25%	174	384	740
25–50%	316	710	1,819

At every level of concession, the in-place cut reached the closing in roughly a third of the market time — and the calendar gap is wider still, because the reset itself is months of dead air. The reset path's nominal haircut looks slightly gentler, but its median journey ran 4.7 years; the market appreciated underneath those sellers, so in real dollars — and in carrying cost, and in what the proceeds would have earned elsewhere — they conceded more.

Nor does the reset buy what sellers believe it buys. Fresh first campaigns went under contract within 30 days 31% of the time, median 71 days; genuine relaunches managed 25% and a median of 90. A reset restores the MLS clock, but only about four-fifths of a fresh launch's attention: agents pull the listing history, and the neighborhood remembers. The usual observational caveat applies — sellers who cancel rather than cut may be a more stubborn breed on any path — but the gap is monotonic across every band and too large to be selection alone.

One reset, and only one, earns the record's blessing: the relaunch priced at the number that is finally real. The dataset's cleanest recoveries are properties that failed at fantasy prices, went dark, and returned at market — several selling at full price in days. A reset works exactly once, when the new number is correct; as a repeated tactic, it is how properties end up in the record book.

Finding six: the cost of the chase

Linking campaigns by parcel produces the study's clearest natural experiment: what does starting wrong cost, relative to starting right?

Properties that sold after at least one failed campaign spent a median 397 cumulative days on market. Properties that sold in a single campaign spent 52. Eight times the calendar — and at upper-market carrying costs (debt service, taxes, insurance, staff, and the estate-grade maintenance a showing-ready home requires), calendar is money.

The dollar cost of the chase requires honesty that most tellings of this statistic skip. Measured naively across the full 25 years, long chasers often eventually matched or beat their original ask — but only because a decade of market appreciation rose underneath them while they waited; in real terms they paid heavily for the privilege of being nominally whole. Measured within a single market era, where appreciation cannot mask it, the chase surcharge ran roughly three additional points of the original ask on top of the time. And at the extreme the chains are extraordinary: estates that listed five, seven, nine times across a decade, surrendering 40–73% of their original asks before finally trading.

And there is a third cost the MLS cannot record: what the money would have done elsewhere. Years ago I advised the owner of a home in a north Scottsdale golf community — already long on market — to take the $2 million offer in hand and put it in an index fund. He held out. The home eventually sold about a decade later, for roughly $2.5 million — a "win" of half a million dollars. The $2 million, compounding in the S&P 500 over that same decade, would have grown to something on the order of $10–12 million. He was right that the house would fetch more eventually. The waiting still cost him a fortune, because the alternative to selling is never zero — it is whatever the proceeds would have earned. Carrying cost is the visible price of the chase; foregone compounding is the invisible one, and over a decade it dwarfs everything else in this study.

Finding seven: one shot — a coin flip, and a demotion

Of properties whose first campaign failed, roughly half ever recorded an MLS sale afterward — 50% at $4M+, 47% at $3–4M, where the data can observe the full descent.

Half. The mental model most sellers carry — "we'll test it high, and re-run it if we have to" — assumes a durable option to repeat the campaign. The record says the option is a coin flip, and exercising it usually means something the seller never planned: recovery routinely occurred in a *lower price tier* than the one the campaign began in. The chase does not merely trim the ask; it demotes the property. Whatever the reason a seller lists — divorce, estate, relocation, retirement, consolidation — the circumstances that made selling possible tend to be temporary, and the first campaign usually happens inside that window or not at all.

One more species deserves honest separation, because it protects the integrity of the failure statistics: some chains were never sales attempts at all. The dataset includes properties listed seven, eight, even fourteen times across years whose asking prices *rose* through the failures — a $6 million first ask becoming a $12 million fourteenth ask with no sale ever recorded. These owners are not chasing the market down; they are advertising, testing, or enjoying the attention. No statistical rule separates them cleanly, so they were not removed — but their presence means the true failure rate for *sincere* sellers is somewhat lower than the headline, while the one-shot conclusion for sincere sellers is, if anything, stronger.

Finding eight: the seller's odds table

Combining both risks — sale and death — into one view produces the table every seller in this market should see before choosing a number. Across all tiers, among 28,112 campaigns:

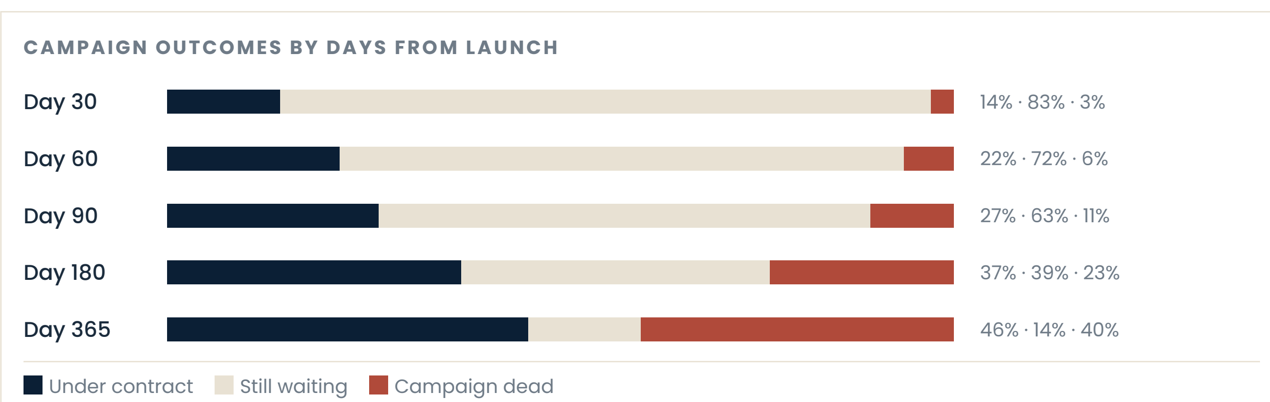

Two lessons live here. The first is patience: at day 60, nearly three-quarters of campaigns are still waiting, and waiting is normal — a seller sixty days in with real showing activity has no cause for panic and no reason to slash. The second is the clock: success holds its lead over failure through the first year across the market as a whole, but the lead shrinks with every band up the ladder — and above $4 million it inverts. In the luxury tiers, dead campaigns overtake successful ones between day 90 and day 180; by one year the ratio is 52% dead to 35% sold. The higher the ask, the shorter the window in which a campaign is more likely to succeed than to have already failed — and everything about a campaign, from the launch ask to the preparation before it to the pre-committed review points, should be designed around that window.

Sidebar: when the buyer arrives

If the upper market is a waiting game, when does the wait end? Contract dates answer precisely, and the answer is a season, not a year. A third of the entire year's $2M+ contracts are written February through April; the July–September trough runs barely 6% a month. The town-level texture matters more than the average. Carefree is the most seasonal luxury market in Arizona — 44% of its contracts land in January–March, nearly half the year in one quarter. Fountain Hills posts the highest single month of any Valley town (13.9% in February). Sedona runs a month behind the Valley with a fall echo. Flagstaff inverts the state entirely: a June–August peak, a mountain market that contracts while the desert sleeps.

The launch implication is mechanical: a listing's largest audience is its first weeks, so a launch timed to enter the arrival season fresh — rather than to enter it carrying ninety days of visible market time — aligns the campaign's one demand spike with the market's.

Sidebar: what the crash looked like at the top

The 2008–2012 rows in this study are not abstractions. Matching sale pairs on the same parcel: 167 properties bought in 2004–2008 were resold at a loss in 2008–2013, at a median loss of 26% — sixty-four of them lost more than 30%, and the worst surrendered two-thirds of its price in five years. Distress told on itself in the listings: the share of $2M+ closings carrying short-sale, lender, or foreclosure language ran half a percent in 2006, hit 15–17% in 2009–2011 — one luxury closing in six visibly distressed, a floor, since not every distressed sale says so — and returned to baseline by 2014. The crash era's 72% campaign failure rate is the other half of the same picture: for four years, listing a luxury home in Arizona failed nearly three times in four.

Sidebar: the frenzy, measured

The 2020–2022 surge is the opposite pole, and the record quantifies exactly how far the market departed from its own physics. Bidding wars at $2M+: 2.7% of closings over ask in 2019, 14.2% in 2021, 19.3% at the 2022 peak — one luxury sale in five over asking. Median days to contract collapsed from 90 to 26. Same-parcel flips bought 2017–mid-2020 and sold into the surge (224 pairs) gained a median 28%, with tails beyond +200% in under three years.

And then it ended — into something new. 2026's texture is fully normalized: 4.2% of closings over ask, median sale at 95.7% of list, 56 days to contract — statistically indistinguishable from 2019 discipline. The volume is not: closings are tracking at roughly four times 2019's count. Part of that is mechanical — a fixed $2 million threshold plus five years of appreciation pushes homes across the line — but not four-fold's worth. The frenzy ended; the market it built didn't leave.

Finding nine: the private market at the top

Roughly one closing in fifteen at $2M+ was pre-arranged — under contract within a day of listing, the signature of pocket sales, off-market deals, and comp entries recorded after the fact. The share rises monotonically with price: 6.5% at $2–2.5M, one in ten at $7–10M, one in seven above $10M.

Two implications. For anyone reading comps at the top of the market: the public record materially understates ultra-luxury liquidity, and a meaningful share of the "sales" it does show were never available to buy. For buyers: at the top tiers, a seventh of the market transacts before it exists publicly — which is a structural argument that representation at this level is partly about access to the flow that never reaches a portal.

The record book

Twenty-five years of data produces a record book, and the records are cautionary monuments. All properties below are real, identified only by shape; each survived every data filter in the methods section.

The longest wait: a Sedona estate spent 4,040 consecutive days — more than eleven years — on a single uninterrupted listing at $2.9 million, and never sold. Another Sedona property logged two listings of 3,510 and 3,441 days — roughly nineteen years of combined market time on one house, no sale.

The most attempts: two properties are tied at fourteen listings each, and they tell opposite stories. One, in North Scottsdale, tried for nineteen years, 5,755 cumulative days, and finished asking *less* than it started — $2.95 million down to $2.5 million, never sold. The other, in Paradise Valley, raised its ask through every failure — $6.25 million to $11.9 million — and stands as the purest specimen of the listing-as-advertisement. A third Paradise Valley property has logged twelve listings since 2007 and is, as of this writing, still trying.

The most market time ever endured by a home that actually sold: a North Scottsdale estate accumulated 7,124 days — nearly twenty years — across nine listings, opening at $12 million and closing at $5 million. Runner-up: a Cave Creek property, 6,599 days across fourteen listings and twenty-two years, $7.6 million to $5.5 million.

And the exception that proves the rule: one DC Ranch Country Club home raised its ask from $6.19 million to $11.9 million mid-frenzy — and caught it, closing at $11.2 million in late 2022. Chasing a price upward worked exactly once in the conditions where it can: when the market itself is rising faster than the seller's ambition. The record book's other entries show what the same move produces in every other market.

What the record says about setting the number

Records are real, and they are small. The seller's instinct is correct that sold comps are a floor, not a ceiling — every record in this dataset was an ask above every prior sale. But the record-setting premiums the market actually ratified were increments, not multiples: successful above-comp sales cleared modestly above the best prior print, while asks 40, 60, 80 percent above the established range populate the failure statistics and the record book almost exclusively. A genuinely superior property has earned an aggressive ask. It has not earned an unratifiable one, and the difference between the two is the difference between setting a record and becoming a cautionary chain.

The launch ask is a survival decision — and functionally the only decision. Since the buyer arrives on their own schedule, the number's job is to keep the listing credible — to buyers, to their agents, and to the appraisal and lending process behind them — for the two to five months a normal upper-market search requires. A number chosen to "leave room" negotiates with no one (Finding Four shows the table is not where discounts happen); it simply removes the listing from the arriving buyer's consideration, starts the meter that Finding Three prices, and stakes the campaign on a rescue — the price cut — that Finding Five shows rarely arrives in time to matter.

Reviews should be scheduled before launch, not provoked by panic. The odds table gives the calendar its structure: day-30 and day-60 reviews fall while waiting is still the likeliest state and adjustment is cheapest; the day-90 review is the serious one, because in the upper tiers the crossover is approaching and the launch audience is gone. Sellers who commit to evidence-based checkpoints before the campaign begins never have to negotiate with their own anchor in public, mid-campaign, from behind.

And some campaigns should not launch at all. The most expensive outcome in this dataset is not the 20% discount — it is the near-half of campaigns that spent months of exposure and carrying cost to produce nothing, the coin flip that follows failure, and the tier demotion that usually accompanies the second act. A campaign whose ask cannot survive its own window was, on the record presented here, better not begun.

Methodology appendix

Status codes. ARMLS terminal statuses were mapped as: C = closed (sold), L = cancelled, E = expired. Active, pending, contingent, hold, and withdrawn records (roughly 1,100) were excluded from outcome statistics as unresolved.

Speculative-listing rule. A listing was classified speculative if any of: the ARMLS construction-status field indicated to-be-built; the public or private remarks matched construction-marketing language (to-be-built, pre-construction, under construction, rendering, completion date, starting price, build-to-suit, proposed build, permit-ready); or the recorded year built postdated the list year.

Pre-arranged rule. A closing was classified pre-arranged if it went under contract within one day of its list date. These are reported as Finding Nine and excluded from all time-to-contract and price-achievement statistics.

Typo rule. Listings whose original ask was at least five times their own final list price were excluded as data-entry errors (60 listings).

Campaign merge. Listings on the same parcel were merged into one campaign when a relist began within 30 days of the prior listing's end. Statistics are campaign-level throughout.

Chains and pairs. Multi-campaign histories and repeat-sale pairs were linked by assessor parcel number, never by street address; unit and lot numbers are preserved to prevent false matches on shared addresses.

Era boundaries. 2008 (crash onset), 2013 (recovery), mid-2022 (rate shock) — set from Arizona market history before results were computed.

Censoring. Three censoring effects are acknowledged: recent campaigns have had insufficient time to resolve; off-MLS outcomes are invisible; and tier-boundary truncation, which produced and then corrected the recovery-rate artifact documented in the methods section, still understates recovery in the lowest band studied.

Based on information from the Arizona Regional Multiple Listing Service for the period January 1, 1998 through August 19, 2026. Statistics compiled and analyzed by the author; ARMLS did not produce and does not endorse this analysis. The statistics describe historical listing outcomes in aggregate and are not a prediction for any individual property. Study conducted August 2026.