



The Delayed Impacts of Upper-Funnel Marketing

An analysis of post-treatment window effects across marketing channels, business verticals, and consumer purchase behaviors.



Chandler Dutton
Measurement Strategy Lead



Tyler Horner
Head of Solutions Consulting



Emily Schwartz
Head of Content and Brand

Intro

You run an incrementality experiment against a marketing campaign. You evaluate the results. You take action based on the results. Business outcomes ensue. Pretty simple, right?

One catch: The *delayed effects* of a marketing campaign can be notoriously challenging to measure. If you cut or reallocate spend too soon based on initial results, you potentially miss out on not just additional lift, but real revenue.

Knowing how, when, and where to measure lagged lift in your marketing strategy is more important than ever, and this report looks at hundreds of incrementality tests across Haus customers – from enterprise powerhouses in fintech and retail to consumer brands in beauty and wellness – to dive into delayed impact patterns in the data.

Executive summary

1

The missing quarter

Across hundreds of upper-funnel experiments, roughly 25% of incremental lift arrives after the media spend window closes. If you're not measuring lagged effects of your marketing, you could be leaving a quarter of your upper-funnel ROI on the table and potentially making budget decisions based on an incomplete picture.

2

Brands selling on retail shelves experience lagged ROI most acutely

For brands selling through retail channels — like Target, Walmart, and Sephora, for example — the median share of lift arriving in the post-treatment window rises to 29%. For 25% of brands in this category, ROI doubles when factoring in latent impact. This makes intuitive sense: Consumers rarely see an ad and immediately drive to a store.

3

Lagged impacts differ by vertical

Beauty and skincare brands see the least latent effect with “only” 11.5% of lift arriving post-treatment – in other words, consumers are more likely to buy these products when they want them, without a longer consideration period. On the other end, CPG, home and furniture, fintech, and consumer apps all see roughly 27-28% of lift arriving late.

4

Bad ROAS reads lead to bad budget calls for upper-funnel channels

There are real business consequences to not considering lagged impacts for upper-funnel channels. An upper-funnel channel showing weak blended return on ad spend (ROAS) in month one may simply be doing exactly what it's supposed to do. Cutting that spend locks in the cost while forfeiting the return. Worse, the efficiency rebound in the following month — driven by latent effects of the spend you just cut — can create a false confirmation that cutting was the right call.

5

Delayed effects are real for nearly all brands across verticals

At the 75th percentile across all experiments, 41.5% of lift arrives post-treatment — meaning ROI is 71% higher when you factor in delayed effects. For any brand near that end of the spectrum, optimizing on real-time ROAS can be actively misleading.

6

Longer time horizons pay off

From the subset of data we've analyzed, upper-funnel tests beat the brand's median iROAS 65% of the time when run 7+ weeks — compared to just 44% success rate when run 4 weeks or less. That's quite an endorsement, and shows under-investment in upper-funnel channels that can be unlocked when fully measured over a longer time horizon.

Background

Upper-funnel media has long carried a measurement burden that lower-funnel channels don't. Click-through rates, last-touch attribution, and traditional MTA were built to measure direct response. When you apply those frameworks to awareness and consideration channels — YouTube, CTV, Linear TV, Podcast, Programmatic Audio — you're using the wrong ruler.

What these frameworks don't account for is timing. Upper-funnel media works by [building awareness](#): Getting your brand onto a consumer's mental shortlist before they're actively in-market. By definition, that means the conversion often happens later — sometimes days, sometimes months after the initial impression. If your measurement window closes when your campaign does, you're under-counting the value of the channels doing the heaviest brand-building work.

This report quantifies that under-count. Drawing on hundreds of Haus incrementality experiments conducted over the past year and a half on upper-funnel channels, we analyzed how much incremental lift arrives after the treatment window closes, how that varies by vertical and channel, and what the financial implications are for brands making budgeting and forecasting decisions today.

Important to note: This analysis covers medium-term (3-4 week) effects. For businesses thinking about reorienting upper-funnel paid media strategy, longer tests matter — big time. We'll come back to this later in the report.

A note on scope:

- This analysis includes data from hundreds of Haus incrementality tests over the past 18 months, pulled in Q1 2026
- Tests span upper-funnel channels including YouTube, CTV, Linear TV, Podcast, Programmatic Streaming Audio, Radio, and Google Demand Gen
- All categories included a minimum of 20 unique experiments
- Average test duration: 28 days
- Average post-treatment window (PTW) duration: 16 days

This report intentionally excludes BFCM and Holiday period experiments, where PTWs overlap with peak promotional periods and create a distinct (and separately analyzed) latency dynamic.

For this analysis, see Haus' [Cyber Week Incrementality Report](#).

The big finding: You're only seeing 75% of the picture without the post-treatment window

Across all experiments in our dataset, the median share of incremental lift arriving in the post-treatment window is ~25%. Put differently: Three quarters of the value of upper-funnel media shows up during the campaign flight.

The remaining quarter shows up after you've stopped spending — and in many measurement setups, after you've stopped looking.

The practical implication: Total ROI is 33% higher when you factor in delayed effects. For a channel you might be evaluating at a 2x iROAS, the true number is closer to 2.67x.

But the median only tells part of the story. At the 75th percentile, 41.5% of lift arrives post-treatment. For brands in that range, factoring in lagged effects doesn't just adjust the ROI — it nearly doubles it. A 2x iROAS becomes closer to 3.4x.

When looking specifically at primary KPIs — the core conversion metric each experiment was designed to move — the median share of post-treatment lift is 23.4%, with a 75th percentile of 38.4%. The consistency between primary and secondary KPIs suggests this isn't a measurement artifact. Latent impact is real, broad, and proportionally distributed across the outcomes upper-funnel media is meant to drive.

Where delayed lag hits hardest: Retail sales

If there's one finding in this analysis that should change how retail brands think about measurement, it's this: **For experiments measuring retail outcomes (what Haus calls “retail KPIs”) — including lift at Target, Walmart, Sephora, and similar retailers — the median share of lift arriving in the post-treatment window is 29%.**

That means the typical retail brand is only capturing 71% of the value of their upper-funnel media during the campaign window. ROI is 41% higher when you account for what comes after.


At the 75th percentile, the story gets even more dramatic: Half of all incremental retail lift arrives after the campaign ends. ROI doubles.

The behavioral logic here makes sense: Brick-and-mortar and omnichannel retail purchases don't happen in real time. Consumers don't listen to a podcast ad and immediately add something to their Target cart. Upper-funnel media for retail is doing something more subtle: It's building the brand salience that earns a spot on the mental shopping list. The conversion happens on the weekend, during the next grocery run, during the next browse session. That's not a failure of paid media — it's exactly how life works! Your measurement and decisioning just have to be built to capture it.

Amazon and DTC: Consistent with the overall picture

For brands measuring Amazon and DTC KPIs, the latency profile looks more like the overall dataset median. The post-treatment window still captures meaningful incremental value — again, roughly 25% of total lift — but the extreme lag seen in retail omnichannel settings is less pronounced.

This likely reflects the nature of the purchase environment. Amazon and DTC channels offer lower friction between intent and transaction. A consumer who's been primed by an upper-funnel impression may convert within days rather than weeks when the path to purchase is a single click. That doesn't eliminate lag, it just compresses it.



The takeaway for Amazon and DTC brands: Lagged lift is still real and still worth measuring. You're still leaving roughly one quarter of your upper-funnel ROI uncounted if your measurement window closes with your campaign. The urgency is slightly lower than in retail, but the principle holds.

Does marketing channel matter? Less than you'd think.

One of the more counterintuitive findings in this dataset: **Across all upper-funnel channels analyzed, the median share of post-treatment lift was within a few percentage points of one another.** YouTube, CTV, Linear TV, Podcast, Programmatic Streaming Audio, Radio — no single channel stood out as dramatically more or less prone to latent effects.

What does this mean? This finding cuts against a common assumption in media planning: that certain channels are inherently “faster” or “slower” to convert. At the median, the data doesn't support meaningful differentiation.


This doesn't mean all upper-funnel tests should be run the same way. Channel still matters enormously for questions of reach, frequency, creative format, and minimum effective spend. A podcast test that doesn't generate sufficient reach won't generate sufficient lift — lagged or otherwise. The absence of channel-level heterogeneity in latency is a finding about timing, not about whether the media is working.

No matter what channels you're testing, we recommend running tests for a comparable duration and doing your best to evaluate tests on similar standards in terms of reporting delay or attribution window to make sure you're comparing apples to apples as best you can.

That said, Haus data makes one thing clear: The most notable variable when it comes to delayed effects doesn't appear to be what paid channel you're running, but rather the vertical you're selling in.

Business verticals are the real differentiator

While marketing channels show relatively similar latency profiles, business verticals tell a much more varied story. Purchase behavior, price points, and the nature of consumer decision-making all shape how quickly (or slowly) upper-funnel impressions convert into measurable outcomes.



The fastest-converting verticals

Beauty and skincare sits at one end of the spectrum. For these brands, the in-flight measurement window captured 88.5% of total ROI — meaning only 11.5% of lift arrived post-treatment. This aligns with what we know about beauty purchase behavior: High impulse, relatively low consideration, frequent repurchase. Upper-funnel media here works quickly or not at all.

Apparel and fashion is similar, with 17% of lift arriving in the post-treatment window. Fast fashion especially tends toward impulse-driven purchase cycles, and the data reflects it.

The slowest-converting verticals

Consumer packaged goods (CPG) shows a median of 27.4% of lift arriving post-treatment — a 38% ROI uplift when lag is factored in. This is largely driven by the omnichannel nature of CPG measurement: Most experiments in this category are measuring for in-store and retailer lift, which carries the same behavioral dynamics described in the retail section above.

Home and furniture showed similar latency, which is easier to explain: Higher price points mean longer consideration cycles. A consumer who sees a YouTube ad for a new sofa isn't buying that afternoon (well, at least, we're not).

Fintech and consumer apps are the slightly less intuitive cases. With a median of 28% of lift arriving post-treatment, these categories rival CPG for lag. Two explanations seem plausible: Upper-funnel ads are reaching consumers who aren't yet in-market for these products, or consumers in these categories do meaningful research before committing. If a subscription is involved, this may also impact consideration windows. Either way, the implication is the same: Real-time ROAS is particularly misleading for fintech and consumer app brands running awareness campaigns.

Extreme cases: When lag defines the entire story

A few outliers to note in the data we analyzed:

In one business' case, nearly 98% of incremental lift arrived after the treatment window closed. The in-flight measurement captured just 2% of the total ROI. In another test run by a different brand, the treatment window captured only 12% of total incremental value — the overwhelming majority arrived later.

These aren't edge cases to be dismissed. Rather, they show what's possible when the measurement window is too narrow for the channel and category combination being tested.

How marketing measurement problems become business problems

Everything detailed so far is a measurement problem. This section is about what happens when the measurement problem becomes a business problem (stick with us).

Month one will look worse than it should. If you start spending on a new upper-funnel channel on the first of the month, and your typical campaign flight is four weeks, you should expect to see roughly two thirds of the total ROAS within that first month. The remaining third arrives in month two — sometimes into month three. This isn't a channel underperforming. It's the channel doing its job on a timeline that doesn't fit neatly into a monthly reporting cadence.

To put numbers on it: Assume a 3x iROAS target and \$100,000 in upper funnel spend. In month one, you'll likely see roughly \$200,000 in incremental revenue — a 2x iROAS. The remaining \$100,000 arrives the following month, with no additional spend required. If you're evaluating that channel at the end of month one and seeing a 2x iROAS against a 3x target, you have an incomplete read — not an underperforming channel.

Mid-month starts make this worse. If spend begins mid-month, the first one to two weeks often show minimal lift as reach and frequency build. That means a channel started on the 15th may contribute very little to that month's blended efficiency numbers — dragging down overall ROAS in a way that has nothing to do with whether the channel is working.

Cutting spend can manufacture a false positive. This is the most dangerous failure mode. When blended ROAS dips in the month you start a new upper-funnel channel, the instinct is often to cut. If you do, the following month will look better — because the latent effects of the spend you cut will show up in the revenue line without any corresponding spend. Through pure correlation, it can appear that cutting the channel helped performance. It didn't. It just delayed the return while eliminating future spend that would have generated its own lagged benefit.

Scaling spend creates a temporary efficiency gap. Even for brands that have established a healthy upper funnel flywheel, increasing spend introduces a new version of the same problem. If you've been spending \$100K/month for a year, your blended ROAS already reflects the mature latency profile of that spend level. When you step up to \$200K, the incremental \$100K starts its own lag cycle. Your blended ROAS will look worse for a month or two — not because the new spend isn't working, but because its returns haven't materialized yet. Finance teams modeling the ROI of a spend increase need to account for this explicitly, or they risk consistently undervaluing the upper funnel at the moment it matters most.

How brands can take action based on these findings

Build lag into your measurement windows by design. The most direct fix: Don't close your measurement window when your campaign ends. Build post-treatment tracking periods into every upper funnel test from the start. The length may vary by vertical — for example, beauty brands need less runway than CPG or fintech — but the habit should be universal.

Model return timelines by channel tier, not just channel. Upper- and lower-funnel media have fundamentally different return curves. If your financial forecasting treats them the same, your monthly efficiency numbers will never make sense. A simple adjustment — modeling upper funnel contribution across two to three months rather than one — will dramatically reduce the noise in your blended ROAS reporting.

Use your vertical as a calibration tool. The vertical benchmarks in this report give you a reasonable prior for how much lag to expect. If you're a CPG brand measuring omnichannel lift, plan for roughly 27-30% of your ROI to arrive post-treatment. If you're in beauty, your in-flight numbers are closer to complete. Knowing where you sit on the spectrum helps you set appropriate expectations with finance and avoid premature optimization.

Treat mid-month starts as a forecasting input, not a performance signal.


If you're launching a new upper-funnel channel mid-month, flag it in your reporting. The blended ROAS drag in that first partial month is expected and explainable — but only if you've modeled for it in advance.

Be especially careful at spend inflection points. New channel launches and significant spend increases both trigger a lag reset for the incremental dollars. Brief your finance and leadership teams before the spend goes up, not after the efficiency numbers come in.

Closing thoughts on measuring delayed business impact of marketing campaigns

We mentioned at the top that the tests represented in this analysis cover medium-length incrementality experiments conducted over 3-4 weeks. We're still studying the data to better understand long-term effects beyond 3-4 week tests, and the Haus team is actively developing solutions for adjusting lift estimates based on long-term effects. But, to do that, a larger set of longer tests needs to be analyzed.

From the subset of data we've analyzed, upper-funnel tests beat the brand's median iROAS 65% of the time when run 7+ weeks – compared to just 44% success rate when run 4 weeks or less. That's quite an endorsement, and shows under-investment in upper-funnel channels that can be unlocked when fully measured over a longer time horizon.



The measurement gap documented here is a common challenge in how many organizations measure their upper-funnel media. The channels are generally working — generating salience, building consideration, and driving conversions that show up days and weeks after the impression. But are measurement infrastructures built to see these patterns and support decision-making based on lagged lift?

Across the hundreds of Haus experiments analyzed for this report, the answer for most brands is: Not fully.

The good news is that this is a solvable problem. Extending measurement windows, modeling return timelines by funnel tier, and calibrating expectations by industry vertical are all tractable changes. The brands that make them will have a clearer picture of what their media is actually doing — and a stronger foundation for the budget conversations that follow.

With that said, we're off to window shop. See you for the next report.

