Measuring PR Impact Beyond EMV

Idea In Short

Earned media value puts a dollar figure on coverage a brand did not pay for, calculated by estimating what the same reach would have cost as an ad. Finance teams like it because it reads like a familiar unit, but the number rests on assumptions about audience size and ad rates that nobody can fully verify, and it says nothing about whether anyone acted on what they saw. It cannot show whether a reader visited the brand's site, searched for it afterward, or remembered it later. EMV descends from advertising value equivalency, a method the PR industry formally rejected in 2010 for the same reason: it prices exposure rather than proving impact. Keeping EMV in a report is reasonable as a rough comparison between campaigns, as long as the method is disclosed and the figure sits beside outcome measures like referral traffic, branded search, and AI citations.

EMV, short for earned media value, is the dollar amount a PR team puts on coverage it didn't pay for, based on what the same exposure would have cost as an ad. The math works backward from advertising. If the brand had bought that same reach as an ad, what would it have cost? That answer is the EMV.

It's popular because finance people can read it without any explanation. The problem is how it gets made. Someone has to guess how many people saw the article and what an ad next to it would cost, and then the tool multiplies those numbers. It can't tell you whether anyone clicked, searched for the brand afterward, or remembered it a week later.

So keep it in the report, but keep it small, and add a line on how it was calculated. Give room for traffic from the coverage, branded search and whether AI answers are quoting it. And if you're the one approving the budget, ask for the starting numbers from before the campaign began. Without them, nobody can say what changed.

Is EMV a financial or accounting metric?

No. It's a marketing number, and there's no standard formula behind it, so two tools can price the same article differently. It shouldn't be booked as revenue or treated like money the company earned. Marketers and finance leaders tend to speak different languages on this point, and the work of translating a marketing number into something a CFO can act on is exactly where EMV runs into trouble, since engagement and exposure figures alone rarely carry much weight in a budget conversation.1 Treating it as financial data invites the kind of scrutiny it was never built to survive.

How reliable is EMV as a measure of PR impact?

Good enough for rough comparisons, like this quarter against last quarter, as long as the method stays the same. Not good enough to prove the PR worked. It doesn't pick up whether the coverage was positive or whether the right people read it, and it says nothing about what they did next. Research comparing coverage volume against business results across several campaigns found that volume alone can track only loosely with outcomes, and that the correlation strengthens considerably once a specific key message actually appears in the coverage rather than coverage existing for its own sake.2 Volume by itself, in other words, is a weak stand-in for whether the coverage actually worked.

What should sit next to EMV in a PR report?

Start with what people did. That means visits that came from the coverage, a rise in searches for your brand name, and whether AI answers cite the pages that mentioned you. Then add how visible you were inside the coverage, like share of voice against competitors and sentiment.

A PR professional writing for the Open University's business school recalls being taught in 2011 that advertising value equivalency was the way to show a client what media coverage was worth. It spread because it came with a return-on-investment figure, a metric clients already trusted. The same professional later argued that it answers the wrong question.

The name is different now, but plenty of teams still use the number the same way. What follows explains how EMV is calculated, where it can be misleading, and what else belongs in the report.

Why leaders still ask for a dollar figure

Anyone who has to defend a budget ends up comparing channels. Paid media has a cost per thousand impressions, a cost per click, a cost per lead. Earned coverage shows up as a folder of links, and that folder is hard to put next to those numbers. A dollar figure fixes that quickly.

Muck Rack's guidance says as much. Executives and clients tend to find a dollar amount easier to read than share of voice or a sentiment score, and one number can show the size of the coverage at a glance. It's a fair point, and probably the main reason EMV is still around after years of criticism. The same guidance adds a condition, though. The figure works best as context, alongside outcome measures.

How earned media value is calculated

There is no single formula. Social platforms and analytics tools usually take impressions and multiply them by a cost per thousand impressions, or CPM, a figure borrowed directly from display advertising, where it is calculated as advertising cost divided by number of impressions and then scaled to a thousand.3 Some also add dollar weights for likes, comments and shares. Sprout Social's documentation describes it as turning impressions and engagement into dollars using the brand's own paid media costs, like CPM and cost per click.

PR monitoring tools go by readership instead. Take Muck Rack's published method for its AVE estimate. It starts with an outlet's monthly unique visitors, multiplies by 0.025 (the share assumed to view any given article), then multiplies by $0.37 per visitor. Run that for an outlet with 2 million monthly visitors and one article, and it comes out at $18,500. The number is only an illustration, but look at what it rests on. Two constants somebody picked. Change either one and the value moves, even though the coverage hasn't changed.

EMV, AVE and the Barcelona Principles

EMV has a family history. AVE came first. Muck Rack describes it as a print-era method that priced coverage at the outlet's advertising rate, so a quarter-page article was valued like a quarter-page ad.

In June 2010, several industry bodies, including the Institute for Public Relations and PRSA, presented the Barcelona Declaration of Measurement Principles. Its fifth principle states that

AVEs do not measure the value of public relations, and it rejects the concept. AMEC, the measurement association, later ran a Say No to AVEs campaign, according to a G2 article on proving PR value.

EMV arrived as the digital-era cousin. Britopian, a PR measurement firm, calls it conceptually and mathematically close to AVE, with advertising rates as the shared reference point. Muck Rack notes that AMEC and the Barcelona framework still discourage relying on AVE as a primary metric. A new name has not removed the original objection.

Where the number breaks down

Britopian lists three problems. EMV counts outputs and ignores outcomes, it treats all coverage as equal, and it is subjective because no standard formula exists. The second one hurts most in practice. Muck Rack points out that a negative article can produce the same AVE as a positive one of the same size. It also notes that a niche trade title read by your buyers can score lower than a general magazine they never open.

Stability is another issue. A PRSA commentary observed that tools reporting in real time can shift baselines, so the same coverage may carry a different value from one report to the next. The author's fix was to measure the same values in every report, so each report compares like with like.

Then there is inflation. Influence Online argues that EMV overstates commercial impact and pushes clients to chase surface-level success. Muck Rack warns against blanket multipliers added for editorial credibility, because they vary across the industry and have no standard justification.

What to measure besides EMV

Outcome measures ask what people did after they saw the coverage. Three signals work for most digital PR campaigns.

Referral traffic comes first. Open your web analytics referral report, find visits from pages that covered you, and see what those visitors did on your site. A placement with a modest EMV that sends engaged readers can be worth more than a large one that sends none. The same logic behind measuring paid media performance applies here: web analytics can trace

which categories of earned coverage actually drive visits, registrations or sales, not just which ones generated the most mentions.4 Muck Rack's guidance lists website traffic from coverage among the metrics that supply the substance AVE lacks.

Branded search comes next. When coverage lands, more people look for the brand by name. Use Google Search Console's Performance report to review query-level clicks, impressions and CTR, while remembering that branded-search changes alone do not prove PR caused the increase. Compare branded search volume in the weeks after publication with the weeks before, and note any paid campaign or seasonal swing that could distort the comparison.

AI citations are the newest signal. In its May 2026 report, Muck Rack, a PR software company, analyzed more than 25 million links cited by ChatGPT, Claude and Gemini across 17 industries. Earned media accounted for 84 percent of citations, journalism for 27 percent, and paid and advertorial content for 0.3 percent. Earned media in that study covers more than press coverage, since it also includes academic, government, encyclopedic and third-party corporate sources. The report says citation behavior varies by platform, so a practical monthly check runs the questions your buyers ask through each assistant and logs which sources appear.

Agencies that report on digital PR for clients have settled on a similar order. According to Vatsal from viaENS:

> When we report on digital PR campaigns to our clients, EMV is one line in the report, not the headline. We also track referral traffic, growth in branded search, and whether the coverage gets cited in AI answers. Those show whether the coverage actually changed anything.

Exposure estimates come last in that order, and evidence of change comes first.

Better complementary metrics that show the true visibility of performance

Behavior signals show what happened after the coverage. A second group of metrics shows how visible the brand was inside the coverage, and these sit closer to the PR work itself.

Share of voice compares your coverage with your competitors' over the same period, a comparison that tends to correlate with market share and revenue in a way that raw exposure numbers do not.5 In a Muck Rack webinar recap, Claire Walsh, who leads public relations at ZipRecruiter, named it her gold-star metric. She also tracks story volume, messaging, story types and where the company appears in a story, for example, in the introduction or in the middle of a list, and compares each against competitors. The placement detail matters. A name in the opening paragraph and a name in a list of ten are different kinds of visibility, and a formula built on outlet traffic cannot tell them apart.

Sentiment and message pull-through show whether the coverage said what you wanted it to say. Muck Rack's PR KPI guide points teams toward sentiment, share of voice and conversions in place of AVE.

Search impact adds a third view. Agility PR recommends tracking SEO metrics to see how coverage affects rankings, and Muck Rack's guide notes that tools can track backlinks. A live link from a respected publication keeps sending readers long after the news cycle has moved on.

How to set a baseline and attribute change honestly

Walsh's advice in the same recap mentioned earlier was to set a baseline before a new campaign goes out. Skip that step and any movement in traffic or search is guesswork. So before launch, write down your branded search volume, referral traffic from the sites you're targeting, your share of voice against two or three competitors, and what AI assistants currently say when someone asks about your category.

After that, keep a dated log. Record the day each piece of coverage went live. Next to it, note anything else happening in those weeks, like a paid push, a product launch or a seasonal spike. Coverage rarely works alone. A chart with dated notes lets you say something modest, such as branded search rose in the three weeks after the feature. That's a claim you can defend. Saying the feature caused the rise usually isn't.

How to build a report that survives a budget meeting

Start with the goal. Britopian notes that output metrics like EMV can suit a program aimed at reach, because more placements put the message in front of more people. They fail a program aimed at lifting sales among existing customers. Pick measures that match what the campaign set out to do.

Then order the page. Open with what changed in behavior, follow with the evidence for it, add coverage visibility, and put the EMV figure last with its method written beside it. Muck Rack's guidance gives a checklist for the estimate itself. Label it clearly as an estimate, record the formula and data sources, add the date the data was collected, and show sentiment and traffic next to it.

Where EMV still earns a place

It does earn one. It gives finance a familiar unit. It lets a team compare second-quarter coverage with first-quarter coverage, or one campaign with another, as long as the method stays fixed. Muck Rack's guidance says it can help executives see the scale of coverage in terms they readily understand. Treated as a directional indicator with its method disclosed, it does a useful job. The trouble starts when someone promotes it to proof of impact.

EMV is a price tag on exposure. It tells you what the attention might have cost if you'd bought it, and it doesn't tell you what the attention did. The numbers behind it are assumptions; each tool does the math a little differently, and the PR industry formally rejected its older cousin, AVE, back in 2010. That said, finance teams get it right away, and it's useful for lining up one campaign against another as long as you use the same method every time and say what it is.

The report should open with what happened after the coverage went live. Did visits from those articles go up? Did searches for the brand move against the baseline, are AI answers citing the pieces? After that, show how visible the brand was inside the coverage, using share of voice, sentiment and where it was mentioned in each story. Put EMV below all of it. A leader reading the page then gets a straight answer on whether the PR did anything.

- 1From metrics to money: how to win your CFO's confidence
- 2Tip sheet: does pumping up the volume pump up business?
- 3Cost per impression
- 4Using web analytics

- 5What is share of voice?

Summary

A PR report earns credibility by leading with what changed rather than what coverage might have cost as an ad. Referral traffic, branded search movement, AI citation, share of voice, and sentiment all show whether anyone acted on the coverage, while EMV only estimates its size. None of that makes EMV worthless, since it still gives finance a familiar unit for comparing one campaign against another, but only when the method is disclosed and the comparison uses the same formula every time. The fix is mostly about order and honesty: set a baseline before a campaign launches, log what else was happening at the same time, and put EMV last in the report rather than first. Treated as a directional, clearly labeled estimate, it still earns a place in the report. Promoted to proof of impact, it misleads the people reading it.