



Rand Fishkin

Cofounder of SparkToro, Alertmouse, & SnackBar Studio. Author of Lost & Founder. Feminist. I love underdogs, cooking, & helping people do better marketing

202k followers • [View Profile](#)

3

POSTS / WEEK

10 PM, UTC

MAIN POSTING TIME

330

AVG REACTIONS
55 comments


70%

CALL-TO-ACTION
FREQUENCY

#1 SUCCESS FACTOR

He wins on LinkedIn by making people feel both validated and smarter: he takes a clear side on a timely marketing/AI problem, explains it in a simple way, and backs it with real evidence or a repeatable pattern. That trust-first "truth-teller" position (founder + author + values-forward human voice) makes his ideas easy to argue about, easy to share at work, and easy to convert into downloads, trials, and event interest without sounding salesy.

FORMAT BREAKDOWN



Video	62%	310 avg reactions
Image	31%	330 avg reactions
Text Only	6%	470 avg reactions
Carousel	1%	260 avg reactions

7 Tactics You Can Steal

01 [Start with a spicy take, and name the idea you disagree with](#)

02 [Open with a gut feeling, then stack a short list of reasons](#)

03 Make your post a tool people can forward to leaders

04 When you share research, prove it's real with effort numbers

05 Turn "we don't know" into a group experiment with clear rules

06 Use a one-line hook that scares people, then calm them fast

07 Share someone else's idea with a huge endorsement and plain explanation

7 Tactics You Can Steal

1 Start with a spicy take, and name the idea you disagree with

He often opens with a strong opinion and spells out the “wrong idea” in plain words so people know exactly what to agree or argue with. That simple setup makes commenting almost automatic, because readers can react without doing mental work first. You can see how powerful this is in his highest performers: these kinds of openings show up in posts that hit 5.87x, 3.74x, and 2.9x relative engagement. To copy it: write the one-sentence belief you’re pushing against, then say your take in one sentence, then give 2–4 short reasons.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

2 Open with a gut feeling, then stack a short list of reasons

He starts by saying what everyone’s feeling (“this is unfair,” “this is frustrating”), then piles up a short list of specific reasons so readers feel seen. That combo pulls in both emotional comments and practical “yes, and…” replies. One of his top posts using this style reached 4.68x relative engagement, and it’s consistent with his best-performing ‘hard opinion’ content in general (the strongest contrarian+data subset averages 3.05x). To copy it: open with the feeling in one line, then list 3–6 real-world reasons people can recognize, then give one clear next step.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

3 Make your post a tool people can forward to leaders

He literally tells readers who to forward it to (boss, finance, leadership) and frames the post as ‘here’s the thing that will help you win that meeting.’ That turns readers into distributors because sharing helps them look smart at work. A top post with this “forward it to your execs” framing reached 3.44x relative engagement. To copy it: name the exact person they should send it to, promise one practical outcome, then give the key takeaway in a sentence before you link anything.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

4

When you share research, prove it's real with effort numbers

Instead of saying “new research,” he adds numbers that signal real work (months, surveys, thousands of runs) and shows just a taste of the charts. That makes people trust it and also makes them curious enough to click for the full story. His research announcement posts can land near 2.89x relative engagement when they combine ‘big claim’ with ‘here’s how much work we did.’ To copy it: state the question you tested, list 2–3 effort numbers, show one small visual or result, and save the full conclusion for the link.

[Example 1 →](#)[Example 2 →](#)

5

Turn “we don’t know” into a group experiment with clear rules

When there’s missing data, he doesn’t just complain—he asks the community to help run a simple experiment with clear steps, time cost, and a promised public result. That makes people comment and repost to recruit others, because participation is the point. A top post built around this “volunteers wanted” experiment format reached 2.85x relative engagement. To copy it: ask one exact question, tell people the steps and time needed, and promise what you’ll publish back to the group.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

6

Use a one-line hook that scares people, then calm them fast

He earns attention with a short first line that sounds like big news, then quickly explains what’s actually happening so readers don’t feel tricked. That’s a clean way to get opens without resorting to fake drama. A top launch/update post using this approach hit 3.26x relative engagement, and it also fits his personal-story pillar where the stronger posts average 1.47x. To copy it: write the one-line “wait, what?” opener, then immediately answer the obvious worry in the next sentence, then share the real point and a clear next step.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

7

Share someone else’s idea with a huge endorsement and plain explanation

He sometimes boosts another person’s idea, but he doesn’t just repost—he adds a bold recommendation and explains, in simple words, what the reader will learn. That makes the post feel like a trusted friend saying “read this, it matters,” which people take seriously and share. One of these endorsement-style posts reached 4.13x relative engagement, showing curation can perform like original content when the take is strong. To copy it: name the person, make one clear claim about why it matters, then explain the main idea in 2–3 short lines.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

Profile Breakdown

HEADLINE FORMULA

Cofounder of SparkToro, Alertmouse, & SnackBar Studio. Author of Lost & Founder. Feminist. I love underdogs, cooking, & helping people do better marketing

[Concrete authority roles with specific brands] + [one flagship proof asset] + [one clear value signal] + [one human detail / interest] + [who he helps]. It works because the brands + book create instant trust, the values word ("Feminist"/underdogs) attracts the right people and repels the wrong ones, and the human detail makes him memorable and approachable instead of "generic marketing expert."

BANNER STRATEGY



Candid behind-the-scenes studio/interview-style photo (people + production gear) with no big text overlay. It works because it signals "I share ideas publicly and do real work with real humans," which matches the honest, conversational voice in the posts and keeps the profile from feeling like an ad.

Content Strategy

PILLAR	%	WHAT CRUSHES	EXAMPLES
AI/Search/Zero-click + big marketing/platform takes (thought leadership)	53%	This pillar performs best when it's a hard, clear point of view backed by data or a simple model. In the 90-post sample, this pillar is 48 posts (53.3%) averaging 1.69x relative engagement; the strongest subset is the "contrarian + data" style (17 posts at $\geq 2.0x$) averaging 3.05x. It flops when it's a softer share/link without a punchy opening: 9 posts in this pillar are $\leq 0.60x$ (average 0.41x), often reading like lightweight amplification rather than a decisive takeaway.	Post 1 Post 2 Post 3
Product updates & launches (SparkToro + Alertmouse)	23%	Product posts underperform unless they're framed as a big release or they pull the audience into the decision. Across 21 product posts (23.3% of the sample), average performance is 0.84x relative engagement, but the "launch/input" subset (5 posts at $\geq 1.0x$) averages 1.91x. Small incremental feature/changelog-style posts are the danger zone: 6 posts at $\leq 0.35x$ average only 0.28x—usually because the 'why this matters' isn't strong enough.	Post 1 Post 2 Post 3
Events, webinars, hiring, and community asks	13%	These posts work when the event is secondary and the hook is a story, datapoint, or speaker insight; the "hook-first" subset averages 1.32x (3 posts). Pure logistics/reminder posts are consistently weak: the reminder-style subset averages 0.41x (3 posts) because it gives people no reason to stop scrolling unless they already planned to attend. The most reliable approach is 'teach something first, then invite.'	Post 1 Post 2 Post 3

Personal stories, values, and behind-the-scenes**10%**

Personal posts do best when they're specific, a little vulnerable, and give readers something to respond to (5 of 9 posts in this pillar are $\geq 1.12x$; average 1.47x). They flop when they feel like a utilitarian update with no story arc: 3 of 9 are $\leq 0.67x$ (average 0.52x). The sweet spot is 'real-life moment \rightarrow lesson \rightarrow small prompt' rather than 'here's what I did today.'

[Post 1](#)[Post 2](#)[Post 3](#)

Conversion Strategy

Their Offer(s)

[SparkToro \(audience research SaaS\)](#)[Alertmouse \(brand/mention tracking\)](#)[SparkTogether \(event\)](#)[Lost & Founder \(book\)](#)

How They Promote in Posts

He uses CTAs frequently, usually after the value/story as a short last line with "👉" (try/register/download). Across the 90-post sample, posts with CTAs (63 posts) averaged 1.26 relative engagement and about 400 likes per post, while posts without CTAs (27 posts) averaged 1.48 and about 466 likes—so the "ask" slightly depresses engagement unless the hook is very strong. One pass over the same time window labeled CTA usage higher (81.1%), which suggests the difference is how strictly a CTA is defined; either way, the safest pattern is value-first and low-friction. Promotions also stay on-brand by selling useful tools/assets ("free trial," "slide-ready graphs") instead of hype.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

How They Promote in Profile

The profile sells trust first: the headline stacks founder roles + the book, and the About section reads like a clear credibility inventory (current ventures + prior Moz leadership + speaking). The banner reinforces "public thinker / creator energy" but doesn't push an offer. The Featured section spotlights the book as the main click path, which is strong for authority but slightly mismatched with how often posts promote SparkToro/Alertmouse; adding a Featured item for the primary SaaS or a best-performing report would tighten the path from post → profile → action.

[View Profile →](#)

Top 10 Posts

- 1 Unpopular Opinion: in 90%+ of cases where leadership demands marketing hit a certain growth target, then fires the team for not hitting it, the problem has nothing to do with marketing.**

He argues that when leadership sets aggressive growth targets and then fires marketing for missing them, the root problem is usually not marketing. He lays out a repeated pattern of leadership cycling through marketing leaders instead of confronting product or strategy issues, and he urges marketers to have an early, direct conversation about what growth is actually realistic.

♡ 1869 💬 362 ↻ 96

- 2 It's not frickin' fair. How are marketers supposed to get budget for anything in a world where:**



He explains why it's getting harder for marketers to justify budgets as tracking and attribution break down. After listing the modern forces making measurement unreliable, he proposes a practical alternative: focus on lift-based experiments, and he points to a video that teaches what to measure and how.

♡ 1564 💬 159 ↻ 131

- 3 I'm calling it. Britney Muller's post here yesterday is the most important short piece on AI you'll read this year. Why? 🙌🙌🙌**



He highlights another creator's AI explanation and makes a strong case for why it matters right now. He frames the value as correcting a common misunderstanding about how large language models work, and positions the takeaway as something marketers should internalize when they talk about AI outputs.

♡ 1244 💬 217 ↻ 177


- 4 There are lots of erroneous assumptions people make about what AI adoption has changed. But NONE ARE MORE WRONG THAN:**



He challenges the idea that AI search tools are replacing Google searches and argues the opposite is happening. He cites a large clickstream-based dataset from named sources and points readers to methodology, framing the claim as something that can be tested, not just debated.

♡ 1177 💬 170 ↻ 135

- 5 Here's the video you show your CMO, CFO, VP-who's-obsessed-with-AI-marketing.**



He offers a plug-and-play resource that readers can send to leadership to set realistic expectations about AI search volume versus Google. He emphasizes updated numbers using newly shared data and calls out the difference between "prompts" and "searches," aiming to prevent sloppy comparisons inside companies.

♡ 1088 💬 132 ↻ 143

6



I have a new job.

He opens with a surprising personal update, then quickly clarifies he isn't leaving his existing companies. He uses that attention to announce a new business (Alertmouse) with a clear, simple use case: free tracking for names, products, and brands, positioned as a better alternative to basic alerts.

♡ 1060 💬 202 ↻ 31

7



In my humble opinion, this is the most useful new Google feature I've seen in years.

He highlights a small Google feature he finds surprisingly useful and compares it against bigger, more hyped product updates. The example shows a practical "copy/paste" style workflow win that saves time for anyone who posts or writes online.

♡ 1049 💬 123 ↻ 11

8



AI responses aren't lies; they're rarely dead wrong. But they're also nowhere near consistent.

He reframes the "AI lies" debate into a more practical warning: inconsistency. He explains why this matters for people using AI tools for recommendations or trying to track brand presence, and he offers a simpler video recap of deeper research for people who won't read charts.

♡ 859 💬 166 ↻ 125

9



NEW research: https://lnkd.in/gwyQz_wg

He announces original research on how consistent AI tools are when asked to list brands or products. He raises the stakes by suggesting that if outputs aren't stable, many AI brand-tracking products may be misleading, and he shares only a small sample of the charts while linking to the full findings.

♡ 814 💬 178 ↻ 153

10



If 100 people ask ChatGPT 5, Claude, or CoPilot the exact same question, how many different answers will they get? 🤔

He asks a very specific question about how much AI answers vary when many people ask the same thing. Instead of speculating, he recruits volunteers to run a simple test, shares the estimated time cost, and promises to aggregate and publish the results.

♡ 860 💬 204 ↻ 68