



THE NEW MULTI-MARKET LAUNCH PLAYBOOK

How AI is redefining campaigns and what to do about it

CONTENTS

Why strategy is no longer the bottleneck

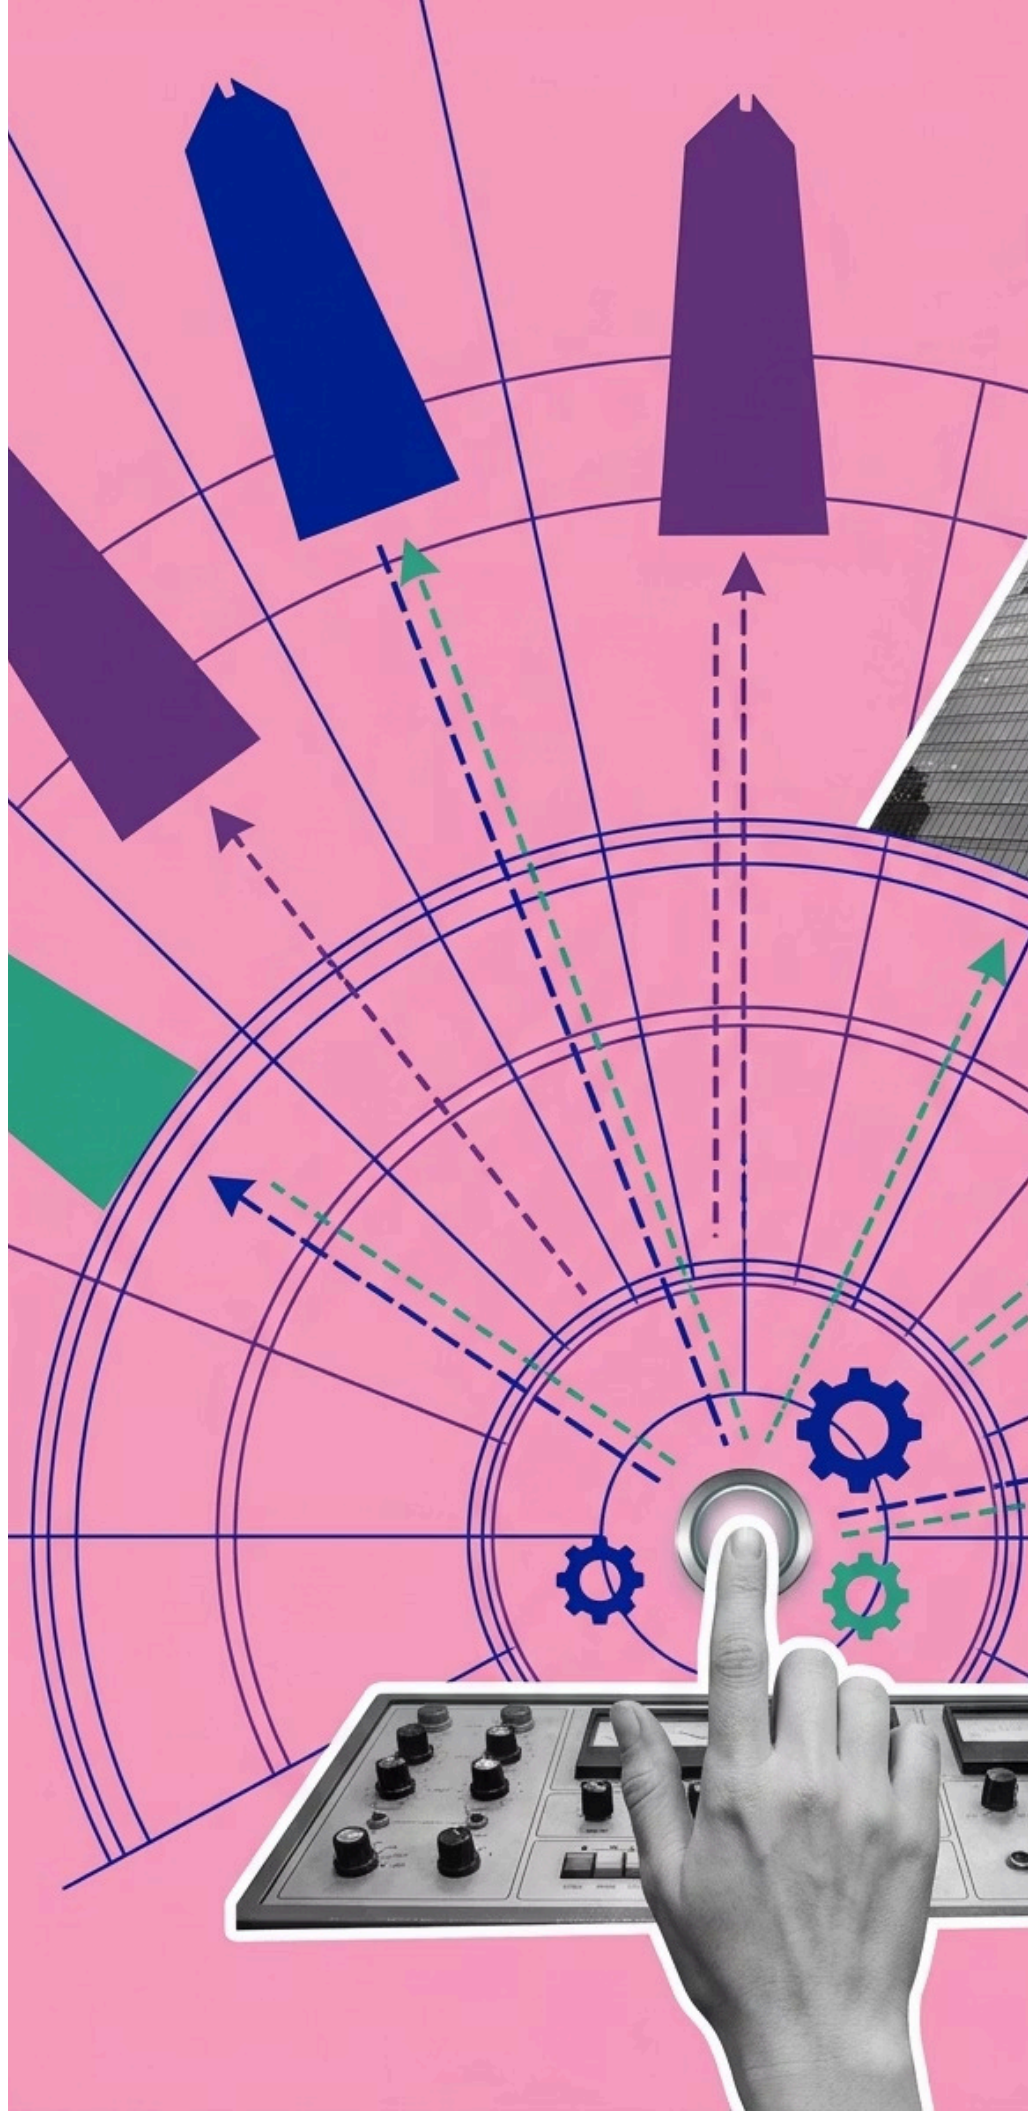
What a multi-market campaign looks like

The AI toolkit trap

How AI moves the needle

A new playbook:
Stage 1: Pre-launch
Stage 2: Launch
Stage 3: Post-launch

What's coming next (and why now is the time to act)



Part 1

WHY STRATEGY IS NO LONGER THE BOTTLENECK

For most of the last 20 years, the hardest part of marketing was thinking.

Figuring out who the audience really was, what to say to them, how to position against the incumbent, which channels actually worked, what the message hierarchy needed to be. Strategy was scarce. Production was cheap, or at least cheap enough. A good brief from a smart strategist would meaningfully outperform a bad one, every time, by margins that justified months of strategic work and a six-figure agency retainer.

That equation has flipped in recent years. Strategy was a competitive advantage when guessing was the only option. When you had three concepts, two markets, and a quarterly cycle, the smartest guess won. **Strategy was how you made the best guess possible before committing the budget.**

But guessing isn't how modern launches are won anymore. Testing is. And testing has a hard prerequisite that strategy alone can't satisfy: a lot of creative, in market, against live signal, fast enough to learn from. Ten variants beat one in any half-decent A/B framework. Fifty beats 10. Two hundred, properly structured across audiences and markets, beats every strategist in the room arguing about which headline is best. The team that ships more variants doesn't need to guess as well. They only need to read signals better, which is a different and more reliable skill.

That's the part of the equation that makes strategy commoditised. Every marketing leader now has access to the same data, the same competitor intelligence, the same audience tools, the same AI-assisted analysis that used to sit behind an agency wall. Two competent teams looking at the same market in 2026 will land on roughly similar strategic conclusions. The conclusions aren't the moat. The ability to test against them is. And that's a production problem.

The numbers speak for themselves. Adobe's **2025 creative workflow study** found that content demand has doubled for 96% of marketers, with nearly two-thirds reporting a fivefold increase. **Content Marketing Institute** now puts the share of marketers using AI for written content at nearly 90%. **BCG's research** has tracked cases where creative development time dropped by up to 75% once AI was briefing and testing the work. And **McKinsey** has documented teams personalising content 50 times faster than the manual approach. None of this is because teams want to ship more for the sake of it. It's because testing requires volume, and volume requires production capacity that current tools and workflows can't supply.

Your strategy is probably fine. The question is whether your production engine behind it can keep up with the testing required to prove it right. Production is where testing happens. And testing is where the wins are.



WHAT A MULTI-MARKET CAMPAIGN LOOKS LIKE

A single multi-market launch isn't one campaign. It's hundreds of assets across dozens of touchpoints. Here's a rough sense of what a five-country launch actually puts on the team.



STAGE 1 PRE-LAUNCH

*3 months
before launch*

Strategy

- 1 master positioning doc
- 5 country audience briefs
- 1 competitive landscape report
- 1 core narrative
- 1 messaging doc
- 1 campaign concept

Execution

- 15 message variants
- 5 translated copy decks
- 5 hero images
- 3 video masters
- 90 modular content blocks
- 30 country-level creative adaptations

148 ASSETS



STAGE 2 LAUNCH

*Launch week
to first month*

Strategy

- 1 launch narrative and calendar
- 5 country GTM plans
- 1 PR strategy and influencer strategy
- 5 paid media plans

Execution

- 15 localised landing pages
- 5 updated product pages
- 15 in-app announcements
- 15 launch blog posts
- 120 display ad variants
- 90 social ad variants
- 15 search ad copy sets
- 20 video cutdowns
- 15 announcement emails
- 25 nurture emails
- 5 press releases
- 60 organic social posts

400 ASSETS



STAGE 3 POST-LAUNCH

*Months 2 and 3
after launch*


Strategy

- 10 per-country performance readouts
- 1 channel mix reallocation plan
- 5 segment refinements
- 1 next-launch insights doc
- 1 brand lift study

Execution

- 60 A/B test variants
- 40 creative refreshes
- 5 early customer case studies
- 5 competitor comparison pages
- 5 thought leadership pieces
- 5 webinars

120 ASSETS



THE AI-POWERED TOOLKIT TRAP


Every tool in your stack now has some kind of AI feature. Microsoft Copilot is everywhere. Your email platform suggests subject lines. Your analytics dashboard pushes AI-generated insight cards. It all works. But none of it will move the needle on a multi-market launch.

Why? The production bottleneck isn't writing a better subject line. It's producing 50 of them, testing them against five audiences, and pulling the underperformers before they burn budget. It isn't generating a nicer image. It's shipping 200 on-brand creative cuts across formats, channels, and markets, ready to swap the moment performance dips.

The AI features baked into most marketing tools were designed to make individual people slightly faster at individual tasks. They weren't built to lift a team's total output by an order of magnitude—which is exactly what testing-led marketing needs.

That's the lens we have to apply to any AI capability before adding it to the stack. The question isn't whether the feature is AI-powered. Almost everything is. The question is whether it raises the team's total production capacity, or just makes one person's task slightly easier.

Subject-line suggestions in your email app fall into the second category. A brand-trained AI platform producing hundreds of on-brand email variants, dynamically assembled per segment per market, falls into the first.



HOW AI MOVES THE NEEDLE

Let's take a closer look at the existing tools your team already uses, what it produces today and what AI capability—measured by output, not feature parity—actually boosts productivity.

Tool	Status quo	With AI	What gets better
Content management system	Every page is produced manually, one at a time, market by market.	Pages and variants are generated on-brand at scale.	More page variants per market, shipped in hours not weeks.
Design tools	A key visual is created, then adapted by a designer.	One concept fuels hundreds of on-brand generations.	Hundreds of on-brand creative cuts per concept, not a handful.
Email platform	Every email and nurture stream is built one at a time.	Emails are written at scale using the same campaign brief.	On-brand email campaigns without manually writing every email.
Localisation vendor / Translation management system	Two- to six-week turnaround per project.	On-brand transcreation in minutes, not weeks	Faster turnaround, lower spend, voice stays intact across markets.
Business analytics/ intelligence	Last week's numbers reported on Monday, acted on by Wednesday.	Cross-market signals are synthesised in real time.	Faster read-react cycles; learnings surfaced, not buried.

The question isn't whether the feature is AI-powered. Almost everything is. The question is whether it raises the team's total production capacity, or just makes one person's task slightly easier.

Part 2

A NEW PLAYBOOK




STAGE 1: PRE-LAUNCH

The biggest shift happens at planning, because you're no longer just building a campaign. You're building the campaign and a testing program to optimise it from day one.

The old pre-launch process gave you one approved version of everything—one positioning doc, one narrative, one creative set, one localised version per market. The new process means planning a variant library to match your testing volume, building a campaign kit to guide future creative generation, and getting clear on what to track so you know what's working.

Your team stops defending the initial strategy and starts collecting data and shipping improvements based on it. The strategy and brief don't change. The output does.



PRE-LAUNCH: BILL OF MATERIALS

Let's revisit our bill of materials for a five-market launch from earlier and see how the process differs with and without AI:

Output	Without AI	With AI	What changed
Master positioning doc	1 document	1 document	No change. Strategy work stays human.
Country audience briefs	5 briefs	5 briefs	Same count, but production time drops from weeks to days.
Competitive landscape report	1 report	1 report	No change. Refresh cadence shifts from one-off to continuous.
Core narrative	1 narrative	1 narrative	No change. Still a single approved narrative.
Campaign concept	1 concept	1 concept	No change. Concept work stays human.
Messaging	1 approved doc	15 message variants	15 times more angles, ready to test on day one of launch.
Translated copy decks	Handled at launch by vendors	5, ready before launch	Localisation moves earlier and in-house. Vendor saved for high-stakes copy.
Hero images	1, manually designed	5 hero images	1 per market instead of 1 for all 5.
Video masters	1, manually designed	3 masters	3 cuts ready for modular reuse, not 1 universal edit.
Modular content blocks	Assembled ad hoc post-launch	90 content blocks	Variant library built at pre-launch instead of being cobbled together later.
Country-level creative adaptations	Built at launch, market by market	30, ready before launch	30 market-specific cuts in hand before launch week starts.

PRE-LAUNCH: 3 THINGS TO GET RIGHT

Start with a testing plan

To plan for the right variant counts, refresh cadences, and everything else, you have to flip the order of planning. Instead of writing a bill of materials for the minimum viable launch, start with the testing plan. That way, the creatives needed to support your testing volume get built into the brief from the start.


Build a campaign kit

Brand guidelines are a pain for most teams to enforce because they live in a PDF that gets referenced occasionally and ignored at volume. With AI, the priority is a campaign kit that can guide production at scale. Build it pre-launch, not after launch week starts going sideways. It's probably the most important thing you'll do before launch.

Define clear roles

Name the AI-to-human handoff per market in the brief itself. What's AI-generated, what's human-edited, what's human-only, and where the review gates sit. Get the roles in the brief from day zero. Clarity here is what turns AI from a productivity claim into real throughput, which is what makes the testing program runnable across all five markets at once.

The takeaway from pre-launch is simple. Strategy work looks the same as it always did, but everything downstream of it has shifted. The teams winning this stage treat pre-launch as the moment they build the testing program. Get your testing plan, campaign kit, and roles right, and your odds of a successful campaign go way up.



LAUNCH: BILL OF MATERIALS

Let's compare the BOM again. Little to no compromise now—every asset gets personalised for each market and optimised for conversion. For brevity, we've left out the strategy items, which remains human-owned.

Output	Without AI	With AI	What changed
Localised landing pages	1 hero per market, locked at launch	15 pages	3 pages per market instead of 1, variants swapped on live performance.
Updated product pages	1 hero product page	5 pages	1 per market, refreshed against in-market signals instead of a single universal edit.
In-app announcements	Templated, English-first	15 announcements	3 per market, market-native, ready on day one.
Launch blog posts	1 announcement post, manually translated	15 blogs	3 per market, written off the launch narrative.
Display ad variants	A handful per market, manually created	120 variants	24 times more variants in rotation, signal-triggered swapped.
Social ad variants	A handful per market, manually created	90 variants	18 times more variants, generated on-brand against the launch narrative.
Search ad copy sets	Written per market	15 sets	3 sets per market, ready before paid spend goes live.
Video cutdowns	1 hero edit, weeks to localise	20 cutdowns	Cutdowns from a single master, market native voiceover and captions same day.
Announcement emails	Templated send, segments configured manually	15 emails	3 per market, assembled live per segment instead of 1 universal blast.
Nurture emails	Pre-built flow, light personalisation	25 emails	5 per market, content assembled live per signal, not templated drip.
Press releases	1 release, manually translated	5 press releases	1 per market, market-native voice and angle from the launch narrative.
Organic social posts	Single calendar, manually translated per market	60 posts	12 per market, pulled from a shared editorial guide featuring market-native copy.



LAUNCH: 3 THINGS TO GET RIGHT

Launch with a variant library

The classic mistake is treating day one as the finish line—pre-bake every variant, lock the creative, lock the sequencing. That caps the program at whatever the team could imagine pre-launch, which is always less than what live data would surface. Again, flip it. Day one is a start state. All the ad variants in the BOM are the inventory for the launch.

Set performance-based triggers for variant swaps

Most teams default to a calendar refresh: Monday reports, Wednesday swaps, repeat. That caps the loop at whatever pace the team can manually keep up with. Replace the calendar with a signal-triggered rule. Underperform by day two? Swap it. Overperform? Double down in that direction the same day.

Measure success with the same metrics across all markets

Five markets running five measurement frameworks is why insight stays siloed. A breakout creative in Singapore never reaches Australia because they're reading different metrics on different cadences. Make it a point to use the same metrics, same definitions, same refresh cadence across all five markets.

Launch week is the only stage where the variable is time. Get the swap cadence and signal frame right, and put your campaign on the path to success for the next 90 days.

STAGE 3: POST-LAUNCH

Most launches lose 60% to 80% of their momentum within four weeks. The campaign's still running, but the launch team has moved on to the next one. And because no testing is happening, nothing changes unless performance tanks unexpectedly.

It's a shame, because post-launch is where the biggest return on a multi-market launch actually gets realized. Launch week proves what works. Post-launch is where you scale the winners, refine segments, build case studies, and stack comparison pages, thought leadership, and webinars on top—turning a launch into a 90-day demand engine instead of a two-week moment.

The biggest challenge is, again, under-resourcing. Sustained testing across five markets needs roughly twice the throughput of launch week, and most teams never budget for it. So the next launch starts from zero instead of building on the last one's insights.



STAGE 3: POST-LAUNCH

Here's what an ideal post-launch bill of materials looks like, compared to the status quo.

Output	Without AI	With AI	What changed
Per-country performance readouts	1 quarterly report, market-by-market slide	10 reports	2 per market across months 2 and 3, assembled continuously instead of quarterly.
Channel mix reallocation plan	Quarterly review, decision delayed by reporting lag	1 reallocation	Same count, but built off live results and decided in days instead of months.
Segment refinements	Manual lookalike build per market, quarterly	5 segments	1 per market, driven by continuous testing across all 5.
Next-launch insights doc	Hand-built recap deck, often skipped	1 insights doc	Synthesized weekly off live performance instead of assembled once at the end.
Brand lift study	1, late-stage, slow to land	1 study	No change in count. Tracked continuously instead of one post-hoc read.
A/B test variants	A handful per channel, monthly cadence	60 variants	10 times more variants in rotation with signal-triggered swaps.
Creative refreshes	Hand-cut, monthly at best	40 creatives	Generated on-brand with an AI platform, weekly per market instead of monthly.
Early customer case studies	1-2 flagship stories, weeks to produce	5 case studies	1 per market, drafted off interview transcripts with human edits on top.
Competitor comparison pages	1-2 long-form pages, manually researched	5 comparison pages	1 per market, refreshed continuously off competitor monitoring.
Thought leadership pieces	1-2 long-form assets, manually researched	5 articles	1 per market, produced in market-native voice.
Webinars	1 marquee webinar, heavy lift to localize	5 webinars	1 per market with AI-assisted scripting, captions, and market-native promotion.

POST-LAUNCH: 3 THINGS TO GET RIGHT

Replicate winners across borders

The default post-launch behaviour is for each market to run its own program, which is probably the biggest source of wasted budget in months two and three. Set a weekly cross-market review with one rule: every breakout variant gets ported to at least two other markets within seven days. Otherwise, you're running five small programs instead of one that compounds.

Turn paid winners into owned assets

Don't treat paid and owned as separate workstreams. A breakout ad hitting 3x CTR shouldn't stay trapped in a display campaign—it should fuel more content such as case studies and webinars. Make it a point to repurpose any paid creative that beats benchmarks for two consecutive reporting cycles within 10 days. Every dollar of paid spend should also fund the owned library that runs after the budget winds down.

Treat the first 90 days as a structured program

Post-launch isn't a maintenance window. It's when you amplify what works—and it deserves its own brief, budget, and resources. The bill of materials should reflect what you learned in the first month: readouts, A/B variants, case studies, and webinars. Done right, this stage drives more ROAS than any other part of the campaign by compounding what's already working.

Your next launch shouldn't start from zero. Getting post-launch right is how you find winners worth replicating in future campaigns, calibrate your bill of materials, and identify the right signals to act on. Each launch makes the next one easier, faster, and more likely to hit. That's the whole point of testing-led marketing.



Part 3

WHAT'S COMING NEXT (AND WHY NOW'S THE TIME TO ACT)

The interesting question isn't whether the next 18 months bring more AI features. They will. The interesting question is what marketing teams start measuring once production stops being the bottleneck. Once shipping isn't the constraint, the rest of the operating model changes shape—and the metrics that defined a good launch start to look like the wrong things to optimise for.

Here's what we expect will happen:

The campaign becomes a continuous program.

The line between one launch ending and the next beginning gets blurry. Teams stop talking about Q3 launches and start talking about the Q3 cohort of the testing program. The 90 days after a launch and the 90 days before the next one merge into one rolling system. Budget moves from per-launch buckets to always-on testing with launch spikes on top.

Localisation stops being a layer and becomes a setting. Transcreation as a vendor relationship fades for routine work. High-stakes work stays human; the rest gets absorbed into the production stack. A six-market launch and a 12-market launch start to cost about the same to produce. The interesting decision shifts from which markets to enter to which ones to enter first.

The org chart reshapes around the testing program.

Roles built around production—campaign managers, creative producers, localisation leads—fold into smaller, sharper teams whose job is to spec experiments, read signal, and decide what the engine produces next. Headcount doesn't necessarily shrink, but titles will change.

None of this is hypothetical. The teams furthest along are already a launch or two ahead, and the gap compounds with every cycle.

RUN MORE CAMPAIGNS DO MORE TESTS LAUNCH ON TIME

See Protaigé in action



o protaigē