

AI MARKETING REWORK COST

The Hidden Cost of Fixing AI Marketing Work

A time and cost analysis of review, corrections, failed retries and release QA across 12 AI-first B2B drafts.

THE RESULT

73% of active time went to review, correction and failed retries - not first generation.

This is a secondary analysis of the same Topic 06 sample. It is not an independent replication.



The answer in four numbers

The visible output took minutes. Making it acceptable took most of the active workflow time.

46 min

INITIAL GENERATION

Reconstructed operator time for the first 12 outputs.

205 min

REWORK

167 review/correction minutes plus 38 minutes in failed retries.

73%

OF ACTIVE TIME

Rework share of the 280-minute end-to-end workflow.

6.1x

COST MULTIPLIER

All-in active labor versus counting first generation alone.

The practical lesson

Generation cost is a poor proxy for delivery cost. Use accepted work as the denominator, with review, retries and release QA included. Boundary: 167 review minutes come from Topic 06; generation, retry and final-QA fields are a row-level reconstruction.

Why the first-output metric can mislead

External evidence shows that speedups depend on task fit, quality thresholds and whether the whole workflow is measured.

19% slower

REAL-WORK TASKS CAN REVERSE THE RESULT

In a randomized trial, experienced open-source developers took 19% longer with early-2025 AI tools, despite believing they were faster. [2]

80% vs 37%

INDIVIDUAL PRODUCTIVITY VS EBIT

McKinsey's 2026 survey found a much wider reported effect on individual productivity than on organization-level EBIT. [3]

50%

QUALITY CONTROL MATTERS MORE

Microsoft's 2026 survey found quality control of AI output was the top human skill respondents said became more important. [4]

The measurement implication

Count the work between output and acceptance. That is where hidden labor, discarded attempts and business judgment appear.

How the cost analysis was built

This report recodes the same 12 briefs from Topic 06. It adds a stage-by-stage time reconstruction without treating the result as new independent evidence.

Stage	What is counted	Evidence status
Initial generation	Prompting, launch and first-output handling	Reconstructed per row
Review + correction	Reading, checking, editing and decision	Recorded in Topic 06 audit
Failed retries	Discarded or superseded generation loops	Reconstructed from audit trail
Final QA	Last acceptance and release-format check	Reconstructed per row

Cost rule




Labor cost = active minutes / 60 x loaded hourly rate. The base case uses \$90/hour; \$60 and \$120 scenarios show sensitivity. Tool subscription and media costs are excluded.

Acceptance remains the endpoint

Ten drafts passed the Topic 06 acceptance rule. Two stopped because evidence or decision authority was missing. Their time remains in the denominator because it was real workflow effort.

Where 280 active minutes went

The first output represented only 16% of active time. Review and failed retries together consumed 205 minutes - 73% of the workflow.






73% rework | 84% after the first output

Rework is review plus failed retries. All post-output work also includes the 29-minute final QA stage.

The first output is visible. The 234 minutes after it are where quality, evidence and approval are completed.

The \$69 draft became \$420 of delivery work

At a \$90 loaded hourly rate, counting only first-generation operator time understates active labor by a factor of 6.1.





\$420 all-in active labor

\$42 per accepted output across the 10 drafts that passed. The two rejected drafts remain part of the cost base.

Scenario, not invoice: the hourly rate is a transparent sensitivity assumption, not a claim about a universal market rate.

Evidence problems consumed the most review time

The 167 recorded review minutes were allocated across the primary correction reasons in the Topic 06 register.



Why evidence costs more
A tone edit can be local. A claim check may require reading the source, narrowing the sentence, changing the surrounding logic and reviewing the conclusion again.

Allocation is by primary trigger. Multi-purpose edits are counted once to avoid double counting.

Rework clustered in source-heavy tasks

Blog source plans and insight memos accounted for 74% of all rework minutes, even though they represented half of the sample.

Task type	Review	Retries	Rework total	Share	Cost @ \$90/h
Thought leadership	27	0	27	13%	\$40.50
Blog + source plan	55	17	72	35%	\$108.00
Campaign email	24	3	27	13%	\$40.50
Insight memo	61	18	79	39%	\$118.50



The risk signal

Evidence density and decision consequence were stronger cost signals than word count. The longest review belonged to an insight memo, not the longest piece of copy.

Cost includes review and retry time only. Final QA and initial generation are excluded from this task-type comparison.

Failed attempts were small individually - expensive together

Seven drafts triggered at least one discarded retry. Ten failed retry cycles added 38 minutes before the final decision.



38 min

TOTAL RETRY COST

Equivalent to \$57 at the \$90/hour base rate.

7 / 12

DRAFTS AFFECTED

Retries were concentrated in source plans and insight memos.

A retry is counted only when an output was discarded or superseded. Edits made directly in the accepted draft remain review time.

A four-minute draft can still become a 35-minute job

BL-2 illustrates the cost gap. The draft was generated quickly, then required source validation, two retry cycles and a final release check.



FIRST OUTPUT | 4 MIN

FINAL QA | 3 MIN

Visible cost: \$6 -> all-in cost: \$52.50

The cost ratio is 8.8x for this draft at \$90/hour. The first output was cheap; the evidence work was not.

The result is not an argument to avoid AI. It is an argument to design the source pack, acceptance rule and reviewer role before generation.

The conclusion survives three labor-rate scenarios

The share of time does not change with the rate. Only the dollar translation does.

Loaded rate	First output	Rework	All-in labor	Per accepted output
\$60 / hour	\$46	\$205	\$280	\$28
\$90 / hour	\$69	\$308	\$420	\$42
\$120 / hour	\$92	\$410	\$560	\$56

What scales linearly


Review, retry and QA labor. At higher skill rates, preventing one unnecessary loop becomes more valuable.

The base-case denominator

$\$420 / 10 \text{ accepted outputs} = \$42 \text{ per accepted output}$. Dividing only by the 12 generated drafts would hide both rejection and completion cost.

Measure six events, not one generation timestamp

A lightweight event log is enough to reveal whether AI reduces accepted-work cost or merely shifts labor downstream.



How to lower rework without lowering the quality bar

The fastest fix is not weaker review. It is moving recurring corrections upstream into the brief, source pack and routing rules.

PRE-APPROVE SOURCES

Give the workflow an explicit evidence set and expiration date.

BAN UNSUPPORTED LEAPS

Define which claims require proof, qualification or escalation.

ROUTE BY RISK

Send source-heavy and decision-heavy work to deeper review.

LOG DISCARDED OUTPUTS

Retry time disappears unless every superseded attempt is recorded.

PRICE ACCEPTED WORK

Use cost per accepted output, not cost per generation.

FEED DEFECTS BACK

Turn repeated corrections into better instructions and reusable checks.

Do not optimize the wrong stage

Saving one minute on generation matters less than preventing an 18-minute evidence correction.

The 12-brief time register

The reconstruction is published row by row. Red rows are drafts that did not pass the Topic 06 acceptance rule.

ID	Task	Gen	Review	Retry	QA	Total	Result
TL-1	Thought leadership	4	7	0	2	13	Accept
TL-2	Thought leadership	4	9	0	2	15	Accept
TL-3	Thought leadership	4	11	0	2	17	Accept
BL-1	Blog + source plan	4	15	4	3	26	Accept
BL-2	Blog + source plan	4	18	5	3	30	Accept
BL-3	Blog + source plan	4	22	8	3	37	Reject
EM-1	Campaign email	3	6	0	2	11	Accept
EM-2	Campaign email	3	8	0	2	13	Accept
EM-3	Campaign email	3	10	3	2	18	Accept
IN-1	Insight memo	4	17	5	3	29	Accept
IN-2	Insight memo	4	21	6	2	33	Accept
IN-3	Insight memo	5	23	7	3	38	Reject

Limits

This is a secondary analysis of one small synthetic B2B sample. The review minutes are inherited from Topic 06; the other stages are a transparent reconstruction. Hourly rates are scenarios. Tool fees, waiting time and downstream performance are excluded. Replication upgrade: instrument live workflows automatically and compare matched AI-assisted and non-AI work.

Sources, scope and reuse notes

Reviewed in September 2026. External sources provide context; time and cost findings come from the Teamulate secondary analysis described in this report.

[1. Teamulate. Where Human Review Changes AI Marketing Quality. Research Topic 06 \(2026\).](#)
[2. METR. Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity \(2025\).](#)
[3. McKinsey. The State of AI: Global Survey 2026.](#)
[4. Microsoft. 2026 Work Trend Index: Agents, human agency, and opportunity.](#)

[5. NIST. Artificial Intelligence Risk Management Framework: Generative AI Profile. NIST AI 600-1 \(2024\).](#)
[6. Nov, S. and Zhang, W. Experimental Evidence on the Productivity Effects of Generative AI. Science \(2023\).](#)
[7. Dell'Acqua, F. et al. Navigating the Jagged Technological Frontier. SSRN 4573321 \(2023\).](#)
[8. Bruniolfsson, E., Li, D. and Raymond, L. Generative AI at Work. NBER 31161 \(2023\).](#)

Definitions

Rework: review/correction plus failed retries. Active time: hands-on operator or reviewer minutes; excludes waiting. All-in labor: initial generation + review/correction + failed retries + final QA. Accepted output: a draft meeting the Topic 06 release rule.

Interpretation

The analysis exposes a cost mechanism. It is not a market-wide estimate of AI rework.

Reuse notes

When quoting the 73% result, preserve the 12-draft sample, secondary-analysis status and reconstructed timing boundary. Cite original publishers for external evidence.

Suggested citation: Teamulate (2026). The Hidden Cost of Fixing AI Marketing Work. Evidence Report.