REVENUE OPERATIONS & AI

The funnel bottleneck nobody's measuring.

The funnel you instrument is the one your CRM can see. The most expensive interval in it happens between two systems, and shows up in neither.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED

25 July 2026

READING TIME

9 min

WORDS

2,042

https://isoglu.com/journal/funnel-bottleneck-nobody-measures/

CONTENTS

MANAGEMENT SUMMARY

Win rate falls and the pipeline review goes straight to the closing stage, because that is where the CRM has fields. But the interval between a lead existing and a human acting on it sits between two systems and is recorded by neither. Smith, Gopalakrishna and Chatterjee found that high lead volumes can reduce downstream sales, through delay caused by limited capacity. Sabnis and colleagues found something more uncomfortable: managerial tracking of follow-up loses its effect as reps get more experienced and higher-performing, while the quality of prequalification gains it. The fix teams reach for weakens exactly where they need it most. Search data says the market has priced this correctly even though attention has not.

Keywords: Sales funnel measurement · Lead follow-up · Sales and marketing handoff · Pipeline management · Revenue operations

Win rate is down four points and the pipeline review has been on the closing stage for forty minutes. Someone thinks it is discounting. Someone thinks it is a competitor. Someone pulls up loss reasons, which say "price" — because "price" is the answer a buyer gives when they do not want to explain the real one.

Nobody is looking at the top of the funnel, because the top of the funnel is green. Leads are up. Marketing hit its number.

Leads up and win rate down are the same problem seen from two ends — and it happens somewhere nobody has a slide for.

The funnel you instrument is the one your CRM can see

Every stage in a pipeline report exists because someone built a field for it. Stages are objects in a database, and an object gets measured when a system owns it.

The interval between a lead existing and a human acting on it is owned by nobody. Marketing's system records that the lead was created and passed. Sales' system records that an opportunity eventually appeared. What happened in between — whether it was two hours or eleven days, whether anyone opened it at all — falls in the seam. Not because it is unimportant, but because there is no field where it would go.

So the funnel has a stage that nothing writes to, and a pipeline review can run for an hour without anyone mentioning it.

More leads, fewer sales

The mechanism is in the literature. Timothy Smith, Srinath Gopalakrishna and Rabikar Chatterjee modelled the marketing–sales interface as three sequential stages and published it in the Journal of Marketing Research. Their statement of the problem is the clearest I know:

> **high lead volumes generated through effective marketing campaigns could actually hurt downstream sales because of wasted effort on poorly qualified leads and/or delays in sales follow-up resulting from limited sales force capacity**

Read that as an operating equation rather than a sentence. Capacity is fixed in the short run. Volume is not. Push volume against fixed capacity and you get the same number of conversations, later, with more of the wrong people in the queue ahead of the right ones.

Marketing's number goes up. Sales' number goes down. Both teams are reporting accurately. Separate numbers do exactly this: each side hits its own while the result between them gets worse, which is the case for a number neither side can move alone.

The scope matters and I will state it rather than let it slide: this is a study of a large home-improvement retailer, which is a consumer setting. The mechanism — capacity constraint producing follow-up delay producing lost conversion — is what transfers. Anyone citing it as B2B evidence would be overreaching, including me.

The uncomfortable finding

Once a team accepts that follow-up is the problem, the fix it reaches for is the one I have watched teams reach for every time I have been in the room: track it. Put follow-up on a dashboard, set an SLA, review it weekly.

Gaurav Sabnis, Sharmila Chatterjee, Rajdeep Grewal and Gary Lilien examined what actually drives whether a rep works a marketing lead, across 461 reps at four firms, and published the result in the Journal of Marketing under a name that has since escaped into the language — the sales lead black hole. On their framing, it covers 70% of marketing-generated leads.

Their finding about the fix is the part worth sitting with:

> **as sales reps' experience increases, their responses to managerial tracking of lead follow-up and marketing lead volume decrease; responses to the quality of the lead prequalification process increase**

And, separately:

> **As sales reps' performance improves, their response to the volume of marketing leads increases, but their response to managerial tracking decreases**

Tracking works on your weakest, newest reps and loses its grip on your strongest, most experienced ones. Prequalification quality does the opposite: the better the rep, the more it moves them.

Which is an awkward result, because tracking is cheap and prequalification is expensive. You can add a dashboard on a Tuesday. Changing what "qualified" means requires marketing and sales to agree on something they have been arguing about for years — and it is the only one of the two levers that gets stronger on the people closing your largest deals.

The market has priced this correctly. Attention is elsewhere.

There is a way to see where a profession's attention actually sits, as opposed to where it says it sits: what it searches for, and what it will pay to be found for.

Figure 1 What the funnel searches for, and what it pays for

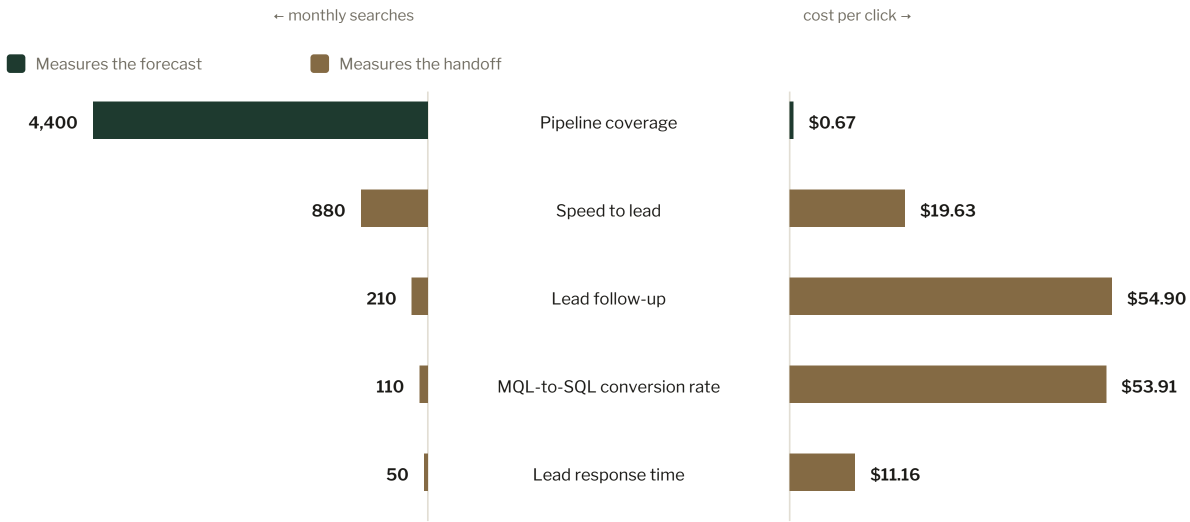

Author's analysis of DataForSEO Google Ads data, 12-month average to June 2026, pulled 25 July 2026. Volume is reported in rounded bands; cost per click is the reported average. US market, English.

A cost per click is not an opinion. It is what a company with a budget will pay to reach one person with that problem, set by an auction against everyone else who wants the same person. The seam costs fifty-five dollars a click and draws two hundred searches a month. The forecast costs sixty-seven cents and draws four and a half thousand.

Two readings are available and both are useful. The generous one: whoever searches for "lead follow-up" is very close to buying something, which is what a high cost per click means. The blunter one: an entire category of software has worked out that this is where the money is, and is bidding accordingly, while the people with the problem are searching, in far greater numbers, for something else.

The usual caveats hold and are in the source line: Google reports volume in rounded bands, one market and one language, and search is one channel among several. I would not build a forecast on this. I would notice the shape.

Where the instrumentation stops

Table 1 **Which parts of the funnel a system writes to**

The interval	Which system owns it	What gets recorded
Campaign to lead	Marketing automation	Volume, source, cost
Lead to first human action	Neither	Nothing, in the stacks I have worked in
First action to qualification	Sales, informally	An outcome, not a duration
Qualification to opportunity	CRM	Stage, date, owner
Opportunity to close	CRM	Everything, in detail
Close to renewal	Customer success	Increasingly, everything

Author's own classification. Not a finding of any cited paper.

Harri Terho, Anna Salonen and Meri Yrjänen looked directly at this function and reported in the Journal of Business & Industrial Marketing that handover is a distinct process with its own mechanics, varying across four different lead contexts. Their own note on the state of the field is the useful part, and it is carefully hedged: "To the best of the authors' knowledge, this study is the first to focus purely on the sales development function as a form of inside sales."

The first. On a function that stands between every marketing euro and every closed deal.

That is also the connection back to why some commercial work accumulates and some resets. A handoff standard is a stock: built once, it keeps working, and it makes every subsequent lead cheaper to convert. A weekly follow-up chase is a flow. It works while you are running it and stops the week you stop.

The objection

Someone reasonable will say the seam is already covered, because the CRM stamps a created date and a first-activity date, and the difference between them is computable.

Computable is not measured. A number nobody reviews, nobody is judged on and nobody has ever defended in a meeting is just a column. Some version of the data usually sits in the tables. The test is whether anyone has ever changed a decision because of it.

There is a third objection and it is the strongest, because it says the interval has already been measured. It has. In 2011, Oldroyd, McElheran and Elkington submitted test enquiries to 2,241 US companies and reported the results in Harvard Business Review: among the firms that answered inside thirty days, the average response took forty-two hours, and 23% never answered at all.

Read what that study had to do. To get a number for this interval, outside researchers had to pose as customers and time the replies, because the companies themselves could not produce it. A quantity that takes a covert field experiment to observe is a quantity nobody is managing. Fifteen years on, an entire software category sells against those forty-two hours — which is a market pricing the problem, not a team measuring it.

The stronger form of the objection is that response time is a vanity metric: fast contact with a bad lead is worse than slow contact with a good one. That is right, and it is why speed is the wrong lever to pull first. Sabnis and colleagues point at the other lever — prequalification quality — and it is the one that holds up on senior reps. Speed matters within a qualified set. Speed applied to an unqualified set just burns the set faster.

What to do about it this week

Take your last twenty closed-won deals and your last twenty closed-lost, and compute one number for each: the hours between the lead being created and the first human action against it.

Not the first automated email. Not the first task created. The first time a person did something a buyer could perceive.

Then plot the two distributions on top of each other. Twenty deals will not give you a correlation, and anyone who claims otherwise is selling you something. You are looking for whether the two distributions are the same shape. If the won deals were reached materially faster than the lost ones, you have found where four points of win rate went — a long way upstream of the closing stage.

And if the distributions are identical, that is worth just as much: you have eliminated the seam, honestly, with your own data, and the pipeline review can go back to the closing stage knowing why.

Figure 2 The response-time audit

THE DEAL Name it. Won or lost.	LEAD CREATED Date and time, from the record	FIRST HUMAN ACTION Not an automated email. Not a task.	HOURS BETWEEN The number nobody owns

Twenty won and twenty lost. Then lay the two distributions over each other and look at the shape, not the average.

Author's own worksheet.

Either way you now have a number for the interval nobody owns — which is the first condition for anybody owning it.

This stays invisible because every function reporting into the review is reporting truthfully about the part it can see, and no part of the organisation is accountable for the gap between two true reports. Marketing delivered the leads. Sales worked the pipeline. Both slides are correct.

The deals that fell into the seam do not appear on either one, because a deal that was never worked never became a deal. It became nothing, and reports have no row for nothing.

Oldroyd, J. B., McElheran, K., & Elkington, D. (2011). The short life of online sales leads. Harvard Business Review, 89(3), 34.

Sabnis, G., Chatterjee, S. C., Grewal, R., & Lilien, G. L. (2013). The sales lead black hole: On sales reps' follow-up of marketing leads. Journal of Marketing, 77(1), 52–67. https://doi.org/10.1509/jm.10.0047

Smith, T. M., Gopalakrishna, S., & Chatterjee, R. (2006). A three-stage model of integrated marketing communications at the marketing-sales interface. Journal of Marketing Research, 43(4), 564–579. https://doi.org/10.1509/jmkr.43.4.564

Terho, H., Salonen, A., & Yrjänen, M. (2023). Toward a contextualized understanding of inside sales: The role of sales development in effective lead funnel management. Journal of Business & Industrial Marketing, 38(2), 337–352. https://doi.org/10.1108/JBIM-12-2021-0596

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, July 25). The funnel bottleneck nobody's measuring. Sinan Isoglu. https://isoglu.com/journal/funnel-bottleneck-nobody-measures/

KEEP READING

Revenue operations & AI
The one number a commercial team should share.
https://isoglu.com/journal/one-number-commercial-team/

Revenue operations & AI
What AI actually changes in revenue operations
https://isoglu.com/journal/what-ai-changes-in-revenue-operations/

Growth that compounds
Why growth compounds – and activity doesn't.
https://isoglu.com/journal/why-growth-compounds/

Sinan Isoglu, MBA (Quantic)
Commercial growth leader, lecturer and doctoral researcher

Online version: https://isoglu.com/journal/funnel-bottleneck-nobody-measures/