REBEL.REBEL LETTERS · PAPER

SECOND OF THREE · SEPTEMBER 2026

The Translation Subsidy

For twenty-five years, customers did the company's understanding for free. The machine just took the job, and sent the company the bill.

BY AUSTIN RIPMASTER · PUBLICATION PAPERS

THISISREBELREBEL.COM · RR-P-2026-007

Every search you have ever typed was a small unit of unpaid work.

Not the reading of results. The typing of the query. The moment before the typing, when a complicated circumstance was compressed into three or four nouns a machine could index. A pain. A child. A deadline. A fear.

That compression was labor. Somebody had to understand the situation well enough to reduce it. Somebody had to know the category's vocabulary well enough to translate into it. Somebody had to strip out everything the index couldn't use, which was almost everything that mattered, and keep only the keywords.

The customer did it. Billions of times a day, for twenty-five years, for free.

This paper is about who was collecting the value of that work, what happened when it stopped, and why a word this house has always used, *authorship*, just changed from a conviction into a cost of doing business.

I. The Subsidy

An economist would recognize the arrangement instantly. The cost of a transaction includes the cost of finding and understanding what to transact. For a quarter century, the understanding half of commercial discovery was performed almost entirely on the customer's side of the counter.

Consider what the search bar actually demanded. To use it well, a person had to hold the category in her head before the category had done anything for her. The back that stopped cooperating, the daughter falling behind in the pool: she had to know that her situation belonged to a market, know what that market called itself, and phrase her life in its terms. *Mattress, firm, back pain. Swim lessons, kids, near me.*

The company, meanwhile, had to do remarkably little understanding in return. It published the category's shared vocabulary: the same premium, the same world-class, the same personalized as everyone else. The index matched code to code. The averaged page worked not because it meant anything but because the customer had already done the meaning. She arrived pre-translated. The company's language could be empty because her labor had filled the gap.

This is the arrangement to name precisely: **the translation subsidy**. The customer's unpaid interpretive work subsidized the company's ability to remain unauthored. Vague language was affordable because someone else was paying the cost of specificity, at the rate of a few seconds of compression per query, across billions of queries, forever.

A Letter in this house has already named the arrangement for what it was: a labor transfer, the burden of converting a life into a market vocabulary, placed on the person living the life. That Letter found it in the archive. This paper prices it.

No line item ever recorded it. Subsidies you don't see are the ones you build on.

Entire disciplines were built on it. Keyword research is the study of how customers perform the translation, so the company can wait at the destination. Search optimization is the practice of decorating the destination. An industry that describes itself as understanding customers was, structurally, an industry that let customers do the understanding and charged for standing nearby.

None of this is an accusation. It was the rational response to the interface that existed. The point is only that the arrangement was never free. It was paid for continuously, invisibly, by the wrong party. And anything paid for invisibly is mistaken, eventually, for a law of nature.

II. The Withdrawal

The subsidy is being withdrawn now, and the withdrawal is measurable.

The average Google search in the United States runs about three and a half words. The average message to a conversational agent runs longer by every published measure, and the published measures spread widely, from roughly nine words to a hundred and three, depending on which population is counted and when. What they agree on is the shape of the thing being typed: full circumstances, pronouns intact, constraints and feelings and irrelevancies included. The customer is not composing queries anymore. She is describing conditions. The previous paper in this series documented the behavior; the numbers behind it, and the disagreement among them, sit in the receipts below.

Read those word counts as an economist rather than a marketer and they say something specific: *the compression stopped.* The labor of reduction is no longer performed at the point of query. The understanding of the situation, the mapping to a category, the stripping of the self: none of it happens there now. The customer has stopped translating because, for the first time, she doesn't have to. The machine in front of her accepts the situation raw and performs the translation itself, at its own expense, in both directions.

The function did not disappear. Commercial discovery still requires that a messy human circumstance be connected to a specific offering. Somebody still has to do the understanding.

The question is only who.

III. The Invoice

Here is the fact the new arrangement turns on: a translator must translate *from* something.

When the customer carried the translation, meaning entered the transaction on her side. Her three nouns were the distillate of her whole situation; the company's averaged page could be empty because her query was full. But when the machine performs the translation, it starts from her raw condition and goes looking for material on the other side: something specific enough, checkable enough, *meant* enough to connect her circumstance to an answer.

Meaning must now exist on the company's side of the exchange, because it has stopped arriving from the customer's.

If it isn't there, the machine does not fail politely. It substitutes. It reaches for the category's shared vocabulary, the tanker language every competitor also rents, and blends the company into the mean, because the mean is all an empty record can support. The company doesn't lose the answer. It dissolves into it.

And the dissolution exposes what the subsidy had been concealing: the unauthored company was never strong. It was subsidized. Its adequacy rested on interpretive labor performed continuously by people it never paid and could not retain, and a moat made of other people's work drains the day they stop digging. They have stopped.

So the cost that was externalized for twenty-five years has been repatriated. The understanding that customers performed for free is now an input the company must produce, in advance, in writing, or forfeit its position in the answer. The subsidy didn't end quietly. It converted into an invoice.

The translation crossed to the machine. The bill crossed to the company.

And this is the point at which a word this house has used from the beginning stops being a preference. *Authorship* was arguable as taste: the practice of recognizing which facts carry meaning and entering them in the record, specifically, checkably, in the company's own hand. Some houses wrote well; most rented words; the market forgave both, because the customer's labor covered the difference.

It no longer does. Authorship is now the factor input the new intermediary consumes. Not voice. Not craft. Not brand as decoration. The production of interpretable, verifiable meaning, demanded not by any doctrine but by the structure of the machine that sits between every company and every question. The doctrine merely arrived early.

IV. What the Machine Pays For

If authorship is an input, it has a specification. Three properties decide whether a sentence in a company's record functions as raw material or as filler.

It must be specific enough to match a condition. The machine is connecting a circumstance, not a keyword. "Programs for all ages" answers no one. The engineer who refused the cheaper material, the boot back for its third resole, the fact with a number and a date in it: these are the units that can meet a human situation halfway, because they describe something that happened rather than something the category says.

It must be checkable enough to survive suspicion. The machine has been lied to more than any reader in history, and the systems around it increasingly weight what can be traced. A specific with no source, no date, no record behind it is vivid boilerplate, and blends like boilerplate. Who recorded the fact, when, against what: that is not metadata. That is whether the sentence is evidence at all.

It must be written where the machine reads. The specifics almost always exist: in the production log, the repair ledger, the founder's head, the patter the best people say out loud on tours. Existing is not entering the record. An asset the reader cannot reach is inventory that was never shelved.

Notice what is absent from the specification: reach, frequency, adjacency, tone of voice. The inputs the last era priced are not the inputs this intermediary consumes. A company can dominate every metric the dashboard shows and still present the machine with nothing to translate from.

The diagnostic travels, and it takes three questions.

Strip from your public record every sentence that could appear, unchanged, on a competitor's site. What survives? That remainder is your authored substrate, the only part of your record that was ever yours.

Of what survives, how much can a suspicious reader check? A claim, a source, a date. That fraction is your evidence. The rest is assertion wearing specificity's clothes.

Who in your company is paid to notice what happened and write it down? Not to produce content. To enter the record. If the answer is no one, the answer is the audit.

Run the three questions and most companies discover the same balance: decades of spend on the subsidized side of the ledger, and almost nothing on the side the invoice now names. That is not a moral failure. It is a book kept accurately for an arrangement that has ended.

V. The Clock

A claim this large owes a way to be wrong.

The claim is that the machine's answers increasingly favor the authored record, the specific and checkable, over the category mean. If that is false, the invoice is imaginary and this paper is a preference in economic costume.

The test is public and dated. Retrieval and fidelity are measurable: ask the machines the category questions, score who appears and in whose words. Within two years of this paper's date, published measurement, this house's own included, will show either that records built from specific, traceable fact outperform averaged category language in machine answers, or that they don't. If blend beats author, if the tanker language holds the answer against the checkable specific, then authorship is a taste after all, this paper is wrong, and the correction will run here, in its own text, where the claim was made.

The previous paper in this series staked its clock on the customer's side of the counter. This one stakes it on the company's. The third will have its own.

VI. The Law

For a quarter century, the deal underneath commercial discovery was never written down, so almost no one noticed what it was: the customer understands; the company gets found.

That deal is over. The customer has stopped performing the understanding, because the machine took it from her hands. And the machine, unlike the search bar, demands real material to work with. The company that spent twenty-five years being paraphrased by its customers must now say something in its own hand, or be paraphrased by the category forever.

The subsidy has been withdrawn. The invoice has been issued. The work has not disappeared; it has changed sides.

Authorship crossed the counter.
The customer put it down.
The company has to pick it up.

THE RECEIPTS

The query-length figures are published measurement, not this house's, and they do not agree with one another. Average U.S. Google search: approximately 3.4 words, per Semrush citing Growth Memo, 2025. Semrush's own AI Mode analysis, covering 1 May to 5 July 2025, puts traditional Google search at 4.0 words and Google AI Mode at 7.22. For conversational agents: 23 words for prompts submitted without search enabled, in Semrush's study of more than 80 million clickstream records from the second half of 2024, published 3 February 2025; 8.7 words for search-enabled prompts and 13.5 for non-search prompts in Semrush's February 2026 clickstream update; and 103 words for the opening message, in WebFX's analysis of 13,252 publicly shared conversations, published 2 September 2025. The studies count different populations by different methods over different periods, and the spread is wide enough that it is named here rather than averaged away. The paper rests on the one thing they agree on, which is that the message to the machine is longer and less compressed than the query to the index, not on any single figure. If that agreement reverses in later measurement, Section V is where this paper said what to do about it.

The economic frame is old. That the cost of search falls on the searcher was formalized by George Stigler in "The Economics of Information" (Journal of Political Economy, 1961). The labor-transfer finding is the house's own: L-009, "The Archaeology of the Unguarded," Section V, which located it in the archive this paper prices. This paper's contribution is only to notice which side of the counter the cost just moved to.

"The translation subsidy" is this paper's coinage. In publishing, the phrase has an unrelated existing sense, grants paid to translate books between languages, noted here so no one mistakes the borrowing for ignorance of it.

The previous paper in this series is *The Question Stopped Dressing Up* (P-006, 9 August 2026), free with its own receipts at letters.thisisrebelrebel.com/p/the-question-stopped-dressing-up. This one was drafted, like that one, in conversation with a machine. The argument does not exempt its author.

Austin Ripmaster is the founder of Rebel.Rebel, an independent research and editorial practice for the retrieval age.