

Solve & Signal: the paper

A practical B2B growth system for the AI age. The premise, five practices, research, measurement and limits in one publication.

EDITION 2.0 / 13 SEPTEMBER 2026 / KILL THE DRAGON GMBH

Abstract

Solve & Signal connects what a business can deliver with how buyers recognise, evaluate and purchase it. Its premise is simple: **solve a buying problem that matters, make the offer recognisable, and give people the evidence and access to act.** AI can help prepare the work behind that decision. The supplier remains responsible for the promise.

The system has five connected practices: buying moments, outcome offers, distinctive signals, access and evidence. It also asks what survives after a conversation: a useful specification, an approved supplier record, an implementation plan or a verified case. These materials can preserve knowledge when people move on.

This is a practitioner framework developed by Kill the Dragon. It combines established concepts with operating proposals. The complete system has not been validated as a causal model of growth. Its value should be tested against buyer behaviour and commercial results.

1. The decision is the unit of work

A manufacturer needs to replace equipment during a planned shutdown. Engineering needs compatibility data. Operations needs a workable installation plan. Procurement needs terms and a supplier it can approve. Finance needs to understand the cost. One enthusiastic contact cannot answer all these questions alone.

The marketing task is to help the right people recognise the supplier and understand its relevance. The commercial task extends further: make the offer assessable, supply the necessary documents and deliver what was agreed. A campaign, a sales conversation and a technical pack can all contribute.

Solve means defining a useful outcome in a particular situation, with scope, constraints and evidence. Signal means giving that offer a recognisable identity and a consistent presence where buying happens. Both depend on actual delivery. A clear promise attached to a poor product makes the problem easier to see.

Our working hypothesis is that reducing avoidable work in these decisions improves the chance of consideration and purchase. This is a hypothesis to test, not a guarantee. Price, fit, capacity, competition and procurement rules still matter.

2. What the research supports

Buying situations matter. Ehrenberg-Bass work on category entry points connects brand recall with the circumstances in which people consider buying. Our application is to organise selected offers and communications around recognisable triggers, such as a contract renewal or a production deadline. [1]

Future buyers matter. John Dawes's 95:5 argument highlights the long intervals between many B2B purchases. The ratio illustrates buying frequency; it is not a measured constant for every category or week. Use the principle to plan beyond today's enquiries and estimate timing in your own market. [2]

Buying requires access. Ehrenberg-Bass research on B2B physical availability examines how brands become easier to find and buy. We apply this to distribution, specification, enquiry routing and purchasing requirements. Visibility alone does not establish that a buyer can order. [3]

Different participants assess different risks. LinkedIn and Bain's enterprise buyer research describes the influence of procurement, finance, legal and operations on supplier evaluation. We use this as a reason to prepare material for the people who approve a decision as well as those who use the product. Its findings should not be read as universal percentages for every B2B purchase. [4]

Context helps define an offer. Jobs to Be Done provides a lens for understanding the progress people seek in a particular circumstance. We use it to investigate desired outcomes, anxieties and trade-offs in buyer interviews. It complements category entry point research; the two concepts answer different questions. [5]

These sources support parts of the framework. They do not establish that credibility compounds at a predictable rate, that persuasion is obsolete, or that a supplier can reliably become the default.

3. Five practices, one connected loop

01. Buying moments

Ask customers about the last actual purchase. What changed? What did they do first? Who became involved? What almost stopped it? Include lost deals and decisions to do nothing. Sales interviews alone can miss buyers who never reached sales.

Start with three priority moments for a pilot. This is a workload choice, not a scientific optimum. Record the buyer's wording, the underlying job, the people involved, existing alternatives and evidence of demand. Expand coverage when the first set is useful.

Output: a ranked moment map. **Decision:** which situation will this offer address first? **Owner:** marketing with sales and customer research.

02. Outcome offers

Translate capability into a decision that can be assessed. Describe who the offer is for, what changes, what is included, what is excluded, what the buyer must provide and how completion will be judged. Provide a price, range or clear quoting basis when commercially possible.

Separate a target from a commitment. A supplier can commit to an assessment deliverable while treating future savings as an estimate. Only guarantee results within your control, under stated conditions. Document assumptions in any value calculation, including implementation costs and the buyer's own effort.

Output: an offer sheet with acceptance criteria and supporting evidence. **Decision:** can the buyer evaluate fit without reconstructing the offer? **Owner:** product or delivery with commercial leadership.

03. Distinctive signals

Choose a small set of recognisable visual and verbal cues. Use them across advertising, offer pages, presentations, partner material and delivery documents. Test whether buyers associate the cues with your brand; internal enthusiasm is not evidence of recognition.

Recognisability and meaningful differences do separate work. Keep a relevant technical advantage, service model or price advantage explicit. Consistent design helps identify who is speaking. It cannot establish that the supplier is suitable.

Output: a compact identity kit with examples and fixed rules. **Decision:** can a buyer recognise the source and understand the offer? **Owner:** brand and design.

04. Access

Trace a real purchase from first enquiry through approval and order. Look for missing specifications, unclear coverage, slow quotes, unanswered questions and hand-offs without an owner. Account for distributors, installers and procurement portals where relevant.

Publish stable pages for offers and evidence. Provide downloadable material that a buyer can share internally. Keep a named contact and a clear next step. On the public website, ensure the important content is present in HTML and reachable through ordinary links.

Output: a tested buying path and a decision pack. **Decision:** can the next participant move the purchase forward? **Owner:** sales operations with marketing, IT and channel partners.

05. Evidence

Capture the baseline before delivery. Record the result, measurement period, scope and conditions. Ask the customer to verify any attribution or quotation and obtain permission before publishing.

Separate measured results from projections and worked examples.

Each completed project can improve the offer and its supporting material. A failed implementation is also evidence: it may reveal an unsuitable segment, a missing dependency or an unrealistic promise. Keep an owner and review date on every proof asset.

Output: verified cases and a change log. **Decision:** what should be retained, corrected or retired?

Owner: delivery with marketing and the customer.

4. Credibility that can travel

A useful insight in the earlier Solve & Signal model was organisational memory. Knowledge can remain in procurement records, specifications, implementation guides and shared case libraries after an individual leaves. This gives teams something concrete to maintain.

Call these items anchors if the term is useful. An anchor needs an owner, a current version and a reason to be used. A document sitting unread in a folder is not evidence of influence. Supplier approval can expire; a specification can become obsolete; a poor delivery can undermine years of familiarity.

Treat accumulation as a management intention. Reuse verified work, keep it current and make it accessible. Do not draw an exponential credibility curve without data. Prefer observable questions: Is the approved supplier record still valid? Did the new stakeholder receive the implementation pack? Was the case used in an actual evaluation?

Default status can sometimes result from familiarity, proven performance and low switching incentives. It is not the operating objective for every market. Competitive tenders, new requirements and buyer choice remain legitimate. Aim to earn consideration and make the decision easier to justify.

5. Made for the AI age

The division of work is straightforward: **keep the buying surface simple; use AI to help manage the detail behind it.** A visitor should quickly find the problem addressed, offer, evidence, limits and next step. Technical depth should remain available when needed.

Useful internal applications include retrieving approved evidence, drafting a specification pack from current product data, identifying contradictions between offer versions and preparing translations for review. Give source documents an owner, version, permitted use and review date. Require a person to approve commitments, commercial terms and customer claims.

For discovery, publish descriptive headings, accessible text, stable URLs and clear source links. Google's guidance for its AI search features continues to emphasise useful, crawlable content and established SEO. No special AI file guarantees inclusion or ranking. [6] A small llms.txt index can serve as a convenience for readers and tools; it is not evidence of visibility.

Test the actual experience. Can a person, or a tool using the public pages, accurately describe the offer and find its exclusions? Record errors against the source. Use such checks to improve clarity, not

as a substitute for customer research. AI-assisted research does not mean all buyers delegate selection to AI.

6. Measure decisions and delivery

Use a small set of defined measures. Establish the baseline and observation period before choosing targets. Segment by offer and buying situation where the sample supports it.

Question	Measure	Caution
Are relevant buyers considering us?	Shortlist inclusion in recorded eligible opportunities	CRM data excludes unseen demand
Can a buyer move forward?	Median and 90th percentile time from complete enquiry to usable quote	Define a usable quote and track missing inputs
Does the offer help decisions?	Wins, losses and no-decisions among comparable opportunities	Small samples and changing deal mix distort rates
Are we delivering the promise?	Share of completed projects meeting agreed acceptance criteria	Include failures and exceptions
Is the evidence current?	Priority offers with verified, in-date supporting material	Document counts do not establish effectiveness

Use recall and recognition research when the budget and sample permit. Search and referral data can help diagnose discovery, but branded search is influenced by existing demand and other activity. It is not a direct measure of trust or a causal growth result.

Review operational faults weekly. Review experiments after their planned observation window, with a primary outcome and guardrails such as margin or delivery quality. Do not stop a test at the first favourable result. Brand effects and long sales cycles require longer observation than a weekly production meeting.

7. A pilot with a defined finish

A six-week setup is a planning example, not a promise of revenue within six weeks.

1. **Week 1:** interview buyers and review recent won, lost and stalled decisions. Pick one buying moment and document the baseline.

2. **Week 2:** write one outcome offer with scope, exclusions, pricing basis and acceptance criteria. Get delivery approval.
3. **Week 3:** assemble existing verified evidence. Mark gaps and estimates explicitly. Prepare the material each approving role needs.
4. **Week 4:** publish an offer page and decision pack using the existing brand codes. Test the enquiry and quote path.
5. **Weeks 5 and 6:** put the offer into relevant conversations and channels. Record objections, missing inputs and operational delays. Fix clear defects.
6. **After the appropriate sales-cycle window:** compare results with the baseline or a credible comparison group. Decide whether to expand, revise or stop.

The pilot needs one accountable lead, delivery participation and a record of decisions. The working canvas on this website provides a shared starting point.

8. Worked example: a planned factory shutdown

This is an illustrative scenario, not a client case or a performance claim.

A supplier has a broad catalogue of industrial components. A plant has a scheduled shutdown and needs replacement parts ready before work starts. A catalogue describes capability but leaves the buyer to assemble the decision.

The first offer is a shutdown readiness assessment: review a defined asset list, confirm compatibility where evidence permits, identify unresolved dependencies and produce a parts and scheduling plan. Installation and savings are outside scope unless separately agreed. The buyer provides asset records and site constraints.

Engineering receives the compatibility matrix. Operations receives the plan and unresolved items. Procurement receives availability, quotation terms and supplier information. All materials share a clear name, recognisable design and version date.

Before claiming an improvement, the supplier records quote turnaround, missing specifications, approval delays and later delivery against the agreed plan. A future case can report measured results with the customer's permission. Until then, the page describes the method and its limits.

9. Boundaries and failure conditions

Use the system where complex offers, several decision participants or repeated buying situations create avoidable work. A simple transactional purchase may need a product page, price and checkout, with little additional material.

Simplification must preserve material detail. A short page that hides exclusions creates risk. A bundle nobody needs is still a weak offer. Distinctive design cannot compensate for unavailable stock, poor service or an uncompetitive proposition.

The system is failing if it increases internal document production while buyers remain confused, if quoted commitments exceed delivery capacity, or if case claims cannot be checked. Return to an observed buying decision and fix the specific obstruction.

References

1. Ehrenberg-Bass Institute. Category Entry Points in a Business-to-Business (B2B) world.
2. John Dawes / LinkedIn B2B Institute. Why You Should Follow the 95-5 Rule.
3. Magda Nenycz-Thiel and Jenni Romaniuk. Easy to Find: Being Where B2B Buying Happens. Ehrenberg-Bass Institute, November 2025.
4. LinkedIn B2B Institute with Bain & Company. The Hidden Buyer Gap.
5. Christensen Institute. Jobs to Be Done Theory.
6. Google Search Central. Google's guide to optimizing for generative AI features on Google Search.

Sources reviewed on 13 September 2026. References support the specific concepts discussed, not the complete Solve & Signal system.

Publication and attribution

Published by Kill the Dragon GmbH, Vienna. Edition 2.0, 13 September 2026. Suggested citation: Kill the Dragon GmbH (2026), *Solve & Signal: A practical B2B growth system for the AI age*, edition 2.0, <https://www.solveandsignal.com/paper.html>.

© 2026 Kill the Dragon GmbH. All rights reserved. The framework's wording, design and publication are attributed to Kill the Dragon. The underlying research and third-party concepts remain attributed to their respective authors. The canvas may be used for internal workshops with attribution. Republishing the complete paper or selling adapted materials requires permission: ask@killthedragon.ai.