Shift90

The problem hiding in plain sight in investor-backed portfolios

One in every 238 private-equity owned companies has a homepage a buyer can act on.

607	476	72	2
HOMEPAGES READ	PE-OWNED	FUNDS	REACH STRONG

The Commercial Transfer Index · European business-to-business commercial surfaces · figures current at 20 August 2026

WHAT THIS MEASURES

Private equity buys businesses that already work. Revenue, customers, a pipeline. The thesis is usually that it can grow faster, or wider, or through more people than are doing it today. That assumes the reason customers buy is available to those people.

It generally sits with a founder and one or two sellers who have had the conversation enough times to know what moves a decision. **Commercial transfer is whether that reason has made it out of those people and onto the surfaces a buyer meets:** the website, the deck, the proposal, the whitepaper, the sales conversation.

The website is the only one of those a stranger can read without being invited. That is what this index holds: 607 European business-to-business homepages, read the way a buyer reads them.

The question the index answers

Can a stranger understand why this business wins before speaking to anyone?

Each homepage is scored out of thirty, six conditions at nought to five.

CONDITION	WHAT IT ASKS
The Buyer's Situation	Does the page name the moment that sent this buyer looking?
The Buyer's Words	Is it written in language a customer would use, or in the company's own?
Who It Is For	Can a right-fit buyer place themselves, and a wrong-fit buyer rule themselves out?
Proof It Works	Is there a result with a number, a period and a customer's name against it?
The Way In	Is there a route for a buyer who is interested and not yet ready to talk?
The Difference	Could a competitor run the same page with their logo on it?

Bands: Fragile nought to eleven. Partial twelve to twenty-one. Strong twenty-two and above. Strong is a score across all six conditions together. It is not a judgement about a company's positioning, its messaging, or the idea the business is built on. It says a buyer arriving cold can recognise themselves, place themselves, believe the evidence, find a way in, and say why this company rather than another.

THE POPULATION

POPULATION	N	MEAN	STRONG
European private-equity owned	476	12.6	2 · one in 238
European index	607	12.7	3 · one in 202
Whole calibrated index	624	12.7	3 · one in 208
Fragile, 0–11	182	38%	
Partial, 12–21	292	61%	
Strong, 22–30	2	0.4%	

> **On provenance.** Eighty-three per cent of the index was assembled from published fund portfolio lists, with no reference to whether the company had a commercial problem. Three per cent approached us. A dataset built from companies who came asking would describe companies who knew they had a problem. This one does not.

WHERE THE VARIATION SITS

Nineteen funds hold eight or more companies in the index. Their averages run from 10.0 to 15.8. The typical spread **inside a single fund** is 10.9 points, and twelve of the nineteen span more than eight.

Eighty-three per cent of the variation sits inside portfolios rather than between funds. The portfolio average is usually set by two or three companies.

FUND	N	MEAN	RANGE	SPREAD
Fund A	9	15.8	13–20	**7**
Fund B	18	15.3	11–22	**11**
Fund C	13	15.1	10–19	**9**
Fund D	9	14.8	10–18	**8**
Fund E	10	14.6	8–21	**13**
Fund F	18	13.7	9–22	**13**
Fund G	9	13.3	9–17	**8**
Fund H	15	13.3	8–20	**12**
Fund I	23	13.2	5–20	**15**
Fund J	8	13.1	9–21	**12**
Fund K	13	13.1	9–16	**7**
Fund L	11	12.9	5–19	**14**
Fund M	10	12.7	8–21	**13**
Fund N	12	12.7	3–21	**18**
Fund O	9	12.6	9–17	**8**
Fund P	12	12.5	9–17	**8**
Fund Q	24	12.3	7–20	**13**
Fund R	21	11.1	7–18	**11**
Fund S	17	10.0	6–14	**8**

Funds holding eight or more companies in the index, ordered by average and identified by letter. Fund averages are a property of this sample: composition, sector mix and sample size differ between houses, and eleven per cent of these rows have been certified by hand. We do not publish them as a ranking and they should not be read as one. The spread column is the finding, and it holds regardless of which fund is which. Partners who would like to know which letter is theirs can ask.

> **What a spread of that size means in practice.** Inside a single investment strategy, some portfolio companies read to a buyer like category leaders and others like commodity suppliers. That is an implication of the spread rather than a measurement of it, and it is the reason the portfolio average is the least useful number on this page.

Our hypothesis is that this leads the number by two or three quarters. A homepage that fails to help a buyer recognise themselves leaks opportunity from the day it goes live. The loss reaches a board pack a quarter or two later, by which point it presents as a market problem. We have not tested it. The index holds no pipeline data, so it cannot.

WHAT IS MISSING

CONDITION	MEAN	SCORE 0–2	SCORE 4–5
The Way In	2.54	48%	17%
Who It Is For	2.45	49%	7%
Proof It Works	2.20	60%	12%
The Difference	2.17	69%	3%
The Buyer's Words	1.76	88%	0.4%
The Buyer's Situation	1.46	89%	3%

The six do not fail evenly. Who it is for, the way in, and why this company average **2.39**. The buyer's situation, the buyer's words, and proof it works average **1.81**.

Our reading of that split is that the first three can be assembled inside the building, from a roadmap, a client list and a competitor's website, and the second three cannot. The index records what appears on a homepage. It does not record how the company arrived at it, so the explanation is ours and the gap is the measurement.

Eighty-nine per cent score two or below on whether the homepage names a situation the buyer recognises. Eleven per cent score zero. Only twelve per cent produce a single proof point carrying a metric, a timeframe and a named client.

What that costs. The homepage is not doing the work of recognition. Whatever persuasion happens, happens somewhere else, usually in a conversation with a founder or with one of the one or two sellers who know the answer. If that is right, the ceiling on growth is how many conversations those individuals can hold in a week, and new sellers will not reproduce the result, because what they were given was a product to describe rather than a reason a customer buys. That is the implication we are testing.

THE MATERIAL IS USUALLY ALREADY WRITTEN

We expected to find companies that had not done the work. What we found, again and again, was companies that had done it and put it in the wrong place.

A data infrastructure business runs a second version of its homepage at a URL nobody lands on, carrying a sharper opening line than the live one. A process orchestration company names four situations its buyers arrive with, five screens below a metaphor. A location intelligence firm places four buyer situations in a closing block under the footer navigation. A public sector software company makes an exclusivity claim in white text on a black panel, two screens down, while the hero says amplify your impact.

That material is already written and signed off. It is simply not where a buyer arrives.

A company that has never articulated why customers choose it has a research problem lasting months. A company that has articulated it and buried it has a decision to make on a Tuesday afternoon.

WHY A SCAN UNDERSTATES THE BEST SURFACES

Seventeen page constructions defeat an automated read. They are concentrated on the strongest surfaces, because a company with real evidence tends to present it in a form a machine cannot see.

CONSTRUCTION	WHAT IT HIDES
Animated counters	Figures initialise at zero and animate on scroll. Every text reader sees a zero where the number should be. Roughly one surface in five
Carousels	Twenty testimonials read as one or none. Logo bars, use cases and customer quotes all rotate
Analyst graphics	A quadrant showing a company top right is a picture. Zero text value
Proof mosaics	Capability claims, brands and outcomes assembled as one high-resolution image
Pinned frames	One panel stays fixed while another scrolls. The reading order a visitor experiences is not the order in the markup
Animated demonstrations	A full product walkthrough, informative to a person, invisible to a reader
No headline element	No `h1` on the page. The read falls back to the top of the document and drags in navigation
Bot protection	403, TLS failure, Cloudflare challenge. A surface unreadable to automated readers is unreadable to more than us

Eight of seventeen. This is why every score above two on the buyer-facing conditions is confirmed by a person before it enters a published figure, and why an automated pass alone is treated as triage.

WHAT A PARTNER CAN ESTABLISH IN AN AFTERNOON

1. **Open the homepage and read the first screen only.** Does it name a situation a customer would recognise, or does it describe the company? Eighty-nine per cent of this population does the second.
2. **Look for one proof point carrying all three parts.** A number, a period, and a customer named. Twelve per cent have one. A logo bar is not proof.
3. **Count the routes in.** If every button asks for a demo or a call, the page serves buyers who have already decided and nobody else.
4. **Run the substitution test on the differentiation.** Put a competitor's name on the sentence. If it still works, it is not differentiation.
5. **Scroll to the bottom and look for the good material.** In the majority of surfaces we read, the strongest line on the page is below the fold.
6. **Ask when a customer was last interviewed about why they chose.** The three lowest-scoring conditions are the three that cannot be answered without it.

WHAT THIS DOES NOT ESTABLISH

It measures what the commercial surface says. It does not measure the business, and a company can run well and score low.

It predicts commercial risk **where a company depends on its surface to win**. Anthropic's homepage scores 7 out of 30. Nobody arrives there to find out whether their problem is solvable, so the number tells you nothing worth acting on. A mid-market portfolio company consolidating a fragmented space is not a name its buyers already know. Nobody arrives already convinced. The homepage has to do the work.

We hold no evidence that a higher score causes better financial performance. We have run no controlled comparison and observed no company without the work. The corpus records no acquisition history, no product count and no service count, so it cannot establish that a company's capability has outgrown its explanation. That is an evidence gap and it stays one.

Of the 607 European surfaces, 86 are certified or human-read and the remainder are machine-scored. Every figure in this document carries that composition. Machine reads are corrected downward on review and a dimension scoring four or above is held out of published figures until a person confirms it.

THE TEST

If the hypothesis holds, the test is a scored portfolio held against pipeline data across two quarters. We would like to run that with a fund.

Pick one company. We read the five surfaces a buyer meets, the way a buyer meets them, and tell you within forty-eight hours whether it reads as one business or a list of parts. That is where it starts.

Craig Vintcent · craig@shift90.partners · +44 7706 349983
Shift90 Partners · Commercial Transfer Index · August 2026