InsightSofa

Customer Experience Management

WHITEPAPER · 2026

NPS and Retention

When the correlation holds, and when it's an illusion.

By the InsightSofa CX Strategy Team

For CX directors, insights leaders and CEOs

Contents

INTRODUCTION

The most over-cited correlation in CX

Few claims in customer experience have been repeated as confidently — and as often — as the one Fred Reichheld put into circulation in 2003: that the Net Promoter Score predicts customer loyalty and, through it, company growth. Two decades and several thousand keynote slides later, the claim has become near-axiomatic. Promoters stay, detractors churn, everything else is detail.

And in many businesses, that is roughly what the data shows. But not in all of them, and not for the reasons most CX leaders assume. A growing body of peer-reviewed research — and our own client work — finds that the correlation between NPS and retention is sometimes genuine, sometimes spurious, and most often somewhere in between. The difference matters enormously, because every CX dashboard, executive bonus and strategic decision built on the simpler version of the story risks being built on sand.

This paper is a careful look at when the link is real and when it is an illusion. It is not anti-NPS. NPS remains a useful, cheap, well-understood pulse-check on customer sentiment, and we use it heavily at InsightSofa. But sentiment is not behaviour, and behaviour is not always revenue. Treating those three as interchangeable is the single most consequential mistake we see CX programmes make.

What you will get out of this paper

By the end of fifteen pages, you should be able to:

- State precisely what Reichheld did — and did not — originally claim.
- Recognise the four mechanisms by which NPS genuinely can drive retention.
- Spot the six structural conditions under which the correlation is an illusion.
- Run a five-step diagnostic on your own data to find out which world you are in.
- Decide whether NPS alone is sufficient for your business, or whether it needs company.

Who this is for

CX directors, insights and analytics leaders, CEOs and product leaders whose company currently uses NPS as a primary north-star metric and who want to understand exactly how much weight it can carry — and where it stops carrying any at all.

A note on tone

Parts of this paper are critical of how NPS is commonly used. None of it is critical of NPS as a measurement instrument. The distinction matters: bad practice does not invalidate a good tool. Our goal is to make NPS more useful, not less.

PART I

What the NPS literature actually says

Before assessing whether the link between NPS and retention is real, it is worth re-reading what Fred Reichheld originally claimed in *Harvard Business Review* in December 2003. The article, *The One Number You Need to Grow*, makes a relatively narrow argument: that the question "how likely are you to recommend?" correlated, in his sample of fourteen industries, with company growth rates over the following years. From that observation, he proposed NPS as a single, simple management metric.

Reichheld did **not** claim, in that original article, that NPS is the best predictor of growth in every industry, that the link works in every business model, or that there is a universal causal mechanism. Those claims were added later, often by practitioners, consultancies and software vendors — sometimes with supporting evidence, sometimes without.

What the peer-reviewed literature has found

Since 2007, when Timothy Keiningham and co-authors published the first major peer-reviewed replication attempt in the *Journal of Marketing*, the academic picture has become considerably more nuanced. The headline findings are:

- NPS is **one of several** mindset metrics that correlate with retention — not consistently the best.
- Top-two-box satisfaction often correlates with two-year retention as strongly as NPS does.
- The single best predictor of retention is frequently the **length of the existing relationship**, not any score.
- Across industries, NPS explains a relatively small share of the variance in growth (typically < 5%).

Where this leaves us

The replication crisis around NPS is real, but it does not mean the metric is useless. It means the metric is *contingent*: it performs well in some conditions and poorly in others. Parts II and III of this paper map those conditions in detail.

Three numbers that frame the debate

The macro evidence on NPS-retention correlation is genuinely mixed. Three representative data points from the published research:

r ≈ 0.17

Reported correlation between NPS and two-year retention in Keiningham et al. (2007). Real, but weak.

< 5%

Share of variance in firm growth explained by NPS in Lee et al. (2021), across 600+ companies.

0.18 vs 0.17

Top-two-box satisfaction correlation vs. NPS correlation with retention — essentially identical.

These numbers are not an argument against using NPS. They are an argument against using NPS *alone*, and against treating a movement in NPS as a reliable proxy for a future movement in revenue. The rest of this paper builds the framework for using it well.

PART I (continued)

Why the question still matters

Three uses of NPS that depend on the link being real

If the correlation between NPS and retention is weak, ambiguous or absent for *your* business, three common uses of the metric become problematic — sometimes seriously so.

1. Setting executive bonuses on NPS movement

Many large organisations now tie a portion of executive compensation to NPS. If NPS does not actually predict revenue or retention in your business, you are paying bonuses for moving sentiment — possibly in ways that do not produce profit. We have seen companies improve NPS by 8 points over two years with no measurable effect on revenue, while paying out bonuses each year.

2. Building churn-prevention models on NPS alone

Predictive churn models that use NPS as the dominant signal will mis-target their interventions in any business where the correlation is weak. Real prevention requires *behavioural* signals — usage patterns, support contacts, payment friction, contract milestones — combined with sentiment, not sentiment alone.

3. Forecasting revenue from NPS targets

If you state in a board pack that "a 10-point NPS improvement will lift revenue by X%", you are making a quantitative claim. If the claim is wrong, leadership decisions based on it will be wrong too. The most expensive version of this mistake is over-investing in NPS-moving initiatives while under-investing in things that actually move retention.

How we got here

The reason the correlation became so widely accepted has less to do with the original evidence and more to do with three structural forces:

- **Simplicity** — one number, one question, one chart. Executive teams want metrics they can hold in their head.
- **Cross-company comparability** — NPS is standardised, which makes benchmarking easy and consultancy work scalable.
- **The vendor ecosystem** — every major CX platform sells NPS dashboards, and prefers to sell certainty rather than caveats.

None of these are wrong, exactly. NPS *is* simple, comparable and well-supported by tooling. The problem only appears when those operational virtues are confused with scientific validity.

A working position

For the rest of this paper, we will treat NPS as a **useful but limited indicator** of future retention. The interesting question is not whether the correlation "exists", but *under what conditions* it exists in your specific business, and what to do when it doesn't.

Worth re-reading

Fred Reichheld & Rob Markey, *The Ultimate Question 2.0* (Harvard Business Press, 2011); Timothy L. Keiningham et al., *A Longitudinal Examination of Net Promoter and Firm Revenue Growth*, Journal of Marketing (2007); Hyun Lee et al., *The use of NPS to predict sales growth*, Journal of the Academy of Marketing Science (2021).

PART II

Four mechanisms behind a genuine NPS-retention link

When the correlation between NPS and retention *is* real, it is almost never just one single mechanism doing the work. In our experience there are four distinct behavioural pathways through which higher NPS translates into customers actually staying. Each one operates on a different time horizon and is driven by a different psychological process. Knowing which mechanisms are active in your business is the difference between a CX programme that compounds and one that doesn't.

Mechanism	What it is	Strength of effect	Time horizon
1. Reduced search behaviour	Promoters spend less mental energy comparing alternatives. They renew or repurchase by default.	Strong — and fast	Months 0 – 6
2. Repeat-purchase decision	At the next renewal or purchase trigger, promoters are more likely to actively choose you again.	Moderate to strong	Months 6 – 18
3. Reduced price sensitivity	Promoters tolerate small price increases without leaving; detractors do not.	Moderate	Months 3 – 24
4. Word-of-mouth feedback loop	Promoters recommend, which improves their own reasons to stay (commitment bias, social proof).	Weak but compounding	Months 12 – 36+

Why mechanism 1 is the most under-rated

The largest single contribution of NPS to retention, in most subscription and low-friction businesses, is not active loyalty — it is the *absence* of active search. Promoters simply do not open competitor websites. They renew, they repurchase, they accept the auto-renewal. This is sometimes called "frictionless loyalty" and it is what gives high-NPS companies an enormous, quiet operating advantage.

The implication: NPS-retention correlation is strongest in markets where customers have *frequent, low-cost opportunities* to switch and choose not to. In markets where they couldn't switch easily even if they wanted to, this mechanism collapses (see Part III).

Why all four mechanisms must be present

Looking only at one mechanism — usually mechanism 2, the repeat purchase — leaves three quarters of the story uncovered. A useful diagnostic when reading any "NPS drives growth" claim is to ask which of the four mechanisms is being proposed. If the answer is "all of them, vaguely", treat the claim with caution. If the answer names one specific mechanism and the data behind it, take it seriously.

In Part V we work through a specific example of a company where mechanisms 1, 2 and 3 are all active (and the correlation is strong), and one where only mechanism 4 is active (and the correlation is essentially meaningless for short-term decisions).

PART III

Six illusions that fake the correlation

An illusion in this context means a pattern in the data that looks like "high NPS leads to retention" but is actually produced by something else entirely. The trouble with illusions is that they look identical to genuine correlations on a dashboard. The only way to tell them apart is to ask, mechanistically, *why* a high-NPS customer would stay — and to be honest about the answer.

STEP 1	STEP 2	STEP 3	STEP 4	STEP 5	STEP 6
Survivor bias	**Contractual lock-in**	**Market structure**	**Habitual buying**	**Aggregation paradox**	**Response bias**

Each illusion is most prevalent in a particular kind of business. The table below gives a first-pass diagnosis: industries and business models in which each illusion is especially common, and a simple symptom you can look for in your own data.

#	Illusion	Most common in	First symptom to look for
1	Survivor bias	Subscription, SaaS, telecom, any panel-based survey	Average tenure of NPS respondents is > tenure of full base
2	Contractual lock-in	Telecom, B2B SaaS with annual contracts, insurance	Churn spikes at renewal date, flat between renewals
3	Market structure	Utilities, regulated banking, local monopolies	Retention is high but NPS is mediocre or negative
4	Habitual buying	Groceries, FMCG, low-involvement retail	Detractors keep buying anyway, at similar frequency
5	Aggregation paradox	Multi-segment businesses, multi-product portfolios	Overall correlation flips when you split by segment
6	Response bias	Any programme with response rates below ~25%	Respondents are not demographically representative

STEP 1

1 Survivor bias

Most companies measure NPS by emailing customers who are still active. By definition, the customers who have already churned are not in the sample. The customers who give detractor scores today are therefore disproportionately the detractors who haven't left *yet*. Customers who would have given the lowest scores have already gone — and they are invisible.

The visible pattern

Detractors in your active base churn at, say, 25%. Promoters churn at 5%. The data appears to show a strong NPS-retention link. But the customers giving the worst scores left long before you could survey them, so the "detractor" group you see is actually the milder half of the true detractor population.

Why it matters

- The measured correlation systematically overstates how much NPS predicts churn in the broader population.
- Movements in NPS over time can look better than they are, because the worst customers keep dropping out of the sample.
- Any predictive churn model trained only on surviving customers will under-predict churn at the segment where it matters most.

STEP 2

2 Contractual lock-in

In businesses with annual or multi-year contracts (B2B SaaS, telecom, insurance, utilities), customers stay because they have to, not because they want to. The NPS score they give in month four has very little to do with whether they renew in month twelve — unless the score is catastrophically low *and* the renewal happens to fall in the right window.

The visible pattern

Retention looks high across all NPS bands during the locked-in period. The correlation only appears, faintly, around the renewal window. Outside of renewal events, NPS is uncorrelated with what customers actually do. This makes it dangerous to use NPS as an early-warning signal for churn at any other point in the lifecycle.

Where to look

- Plot churn-by-month-since-contract-start. If you see flat-then-spike at renewal, lock-in is doing most of the retention work.
- Compare NPS at month 3 vs. month 11 for the same customer. If month 11 is dramatically lower, dissatisfaction has been suppressed by contract.
- Look at "silent churn" — customers who downgrade or stop using but technically remain on contract. They will inflate retention metrics.

A useful test

Strip out contractually-locked-in customers from the retention calculation entirely. If the NPS-retention correlation among the *renewable* base remains strong, the relationship is real. If it disappears, you have been measuring contract law, not customer experience.

STEP 3

3 Market structure (no real alternative)

In categories where the realistic number of alternatives is very small — local utilities, regulated banking in concentrated markets, niche B2B suppliers without competitors — customers may stay regardless of how they feel. The NPS score reflects sentiment honestly, but sentiment doesn't translate into behaviour because there is nowhere meaningful to go.

The visible pattern

Retention rates are very high (90%+), NPS is mediocre (often barely positive or negative), and the two are essentially uncorrelated. Worse, the relationship can be inverted: customers who give high scores do so because they're not paying attention, while engaged customers — who would be the most likely to leave if they could — give the lowest scores.

Diagnostic question

Ask three customers, in their own words, what they would do if you doubled your prices tomorrow. If the realistic answer is "complain, but stay", market structure is the dominant driver of retention. NPS movements may still be worth tracking for reputation and brand reasons, but not as a leading indicator of revenue.

STEP 4

4 Habitual purchasing

In high-frequency, low-involvement categories — supermarket purchases, fuel, subscription media, FMCG — customers buy on autopilot. They are not making active comparison decisions every week, and the question "how likely are you to recommend?" activates a cognitive process that has nothing to do with the behavioural process driving their actual purchases.

The visible pattern

Detractors and promoters show only modestly different repurchase rates. Promoters say nice things in surveys and buy at a similar rate as everyone else; detractors say critical things and also buy at a similar rate. The data *looks* like a weak positive correlation, but it is largely noise.

What to do instead

- Measure declared next-purchase intent at moments of decision, not after the fact.
- Use behavioural cohort analysis — share-of-wallet, frequency, basket size — alongside NPS.
- Treat NPS as a brand-health metric in habitual categories, not as a retention predictor.

An accidentally useful finding

If your habitual-purchase customers show *any* behavioural difference by NPS band, that is itself interesting — it usually means a sub-segment within them is making active choices, which is exactly the segment you want to identify and protect.

STEP 5

5 Aggregation paradox

Sometimes the correlation between NPS and retention is real *within* each segment of your customer base, but the aggregate looks completely different — or vice versa. This is the well-known Simpson's paradox, dressed up in CX clothes. It is one of the most common and most underappreciated illusions we encounter in client data.

A concrete example

A B2B software company finds that, in aggregate, NPS and retention are uncorrelated. Splitting by customer size reveals the truth: among enterprise customers, NPS strongly predicts renewal. Among SMB customers, NPS doesn't predict anything — those customers are price-driven. The aggregate hid both signals. Acting on the aggregate would have led to the wrong investment decisions in both segments.

How to find it

- Re-run the NPS-retention correlation separately by segment (size, tenure, channel, geography, product).
- Look especially at moments-of-truth segments: new customers, renewing customers, recent complaint cases.
- Watch for paradoxical signs — the aggregate correlation pointing one way while every segment points the other.

STEP 6

6 Response bias and self-selection

NPS programmes with low response rates have a fundamental representativeness problem. The customers who respond are systematically different from those who don't — typically more engaged, more loyal, and more polarised (either strongly satisfied or strongly dissatisfied). The metric ends up measuring the views of a self-selected minority.

The visible pattern

Response rates below 25–30% should be treated with caution. Below 10–15% they are essentially unrepresentative. Any correlation between NPS and retention you observe in this sample tells you about behaviour among engaged customers, not your customer base as a whole.

Two practical mitigations

- **Boost response rate** — alternative channels (SMS, WhatsApp, in-app, terminals), shorter surveys, better timing. We routinely see response rates double when transactional channels are used at the right moment.
- **Weight or model** — apply a weighting scheme that adjusts NPS for the demographic and behavioural profile of non-respondents.

What "good" looks like

If your response rates are above 35%, your sample is demographically close to the customer base, and the NPS-retention correlation holds when you slice by segment, you almost certainly have a real signal. If any of those three conditions fails, treat the headline correlation as circumstantial rather than predictive.

PART IV

A diagnostic test for your own data

The point of the previous sections is not to argue that NPS is useless — far from it. The point is that whether NPS predicts retention in *your* business is an empirical question with a knowable answer. The five-step diagnostic below is the test we run internally with InsightSofa clients when this question matters. Each step takes between a few hours and a few days of analyst time. Combined, they produce a defensible verdict on whether to treat NPS as predictive, suggestive, or essentially decorative.

Step	What to do	What you are looking for
1. Calibrate the sample	Compare the demographic and behavioural profile of NPS respondents to the full base.	Match on tenure, segment, ARPA. Large drifts mean response bias is contaminating your data.
2. Lock-in adjustment	Remove customers who cannot churn in the next 12 months (active contracts, etc.) from the analysis.	If the correlation only exists in the locked-in cohort, you have contractual lock-in, not loyalty.
3. Behavioural cross-check	For each NPS band, compute repeat-purchase rate, share-of-wallet and time-to-next-purchase.	All three should move in the same direction as NPS. If they don't, NPS is not behavioural.
4. Segment-level test	Re-run the NPS-retention correlation separately for each major segment (size, tenure, channel).	If the segment-level correlations contradict the aggregate, you have an aggregation paradox.
5. Time-lag check	Lag NPS by 3, 6, 12 months and re-correlate with retention in each subsequent window.	A genuine leading indicator should peak at 6–12 month lag. No correlation at any lag = no signal.

How to score the diagnostic

Treat each of the five steps as a pass/fail check. A step "passes" when the data behaves the way a real NPS-retention link should make it behave. The rubric below maps the number of passes to what you can credibly do with NPS in your business.

Diagnostic outcome	If 4–5 checks pass	If 2–3 checks pass	If 0–1 checks pass
Treat NPS as…	A genuine leading indicator of retention	A directional indicator, not a forecast	A brand-health metric, not a retention predictor
Use it for…	Forecasting, early-warning, exec bonuses	Trend monitoring, sentiment tracking	External comparability, reputation tracking
Pair it with…	Behavioural cohort dashboards (light pairing)	Repeat-purchase and complaint data (medium)	Full behavioural model — NPS is supporting evidence only

Why each step matters

The five checks were chosen to cover the six illusions of Part III with as little analytical effort as possible. Each step rules out a specific failure mode:

Step 1 controls for response bias (Illusion 6)

If respondents don't look like the customer base, no headline correlation tells you anything about the base. This step should always be done first; if it fails, the others can be misleading.

Steps 2 + 4 control for spurious patterns (Illusions 2 + 5)

Lock-in and aggregation hide segment-level truths. Looking at the renewable cohort and splitting by segment usually exposes them.

Step 3 distinguishes sentiment from behaviour (Illusion 4)

Habitual buyers behave the same regardless of NPS. A behavioural cross-check — repeat purchase, share-of-wallet, time-to-next — separates real change from polite responses to a survey.

Step 5 confirms NPS is leading, not lagging (Illusion 1)

If NPS only correlates with retention at zero lag, it is being measured after the customer has already decided. A real leading indicator peaks at 6–12 months of lag — that is the fingerprint of a sentiment that *then* drives behaviour.

Final caveat

Even a diagnostic that passes all five checks does not prove that NPS *causes* retention. It only proves the correlation isn't spurious. To prove causation, you need a controlled intervention — change the experience for one cohort, hold another as control, and watch the gap open.

PART V

A tale of two companies

To make the diagnostic concrete, consider two fictitious but realistic companies with *identical* headline NPS — both score +28 — but very different underlying realities. The first, **AlphaCloud**, is a mid-market B2B SaaS company. The second, **OrbitMobile**, is a regional mobile carrier. Running the five-step diagnostic on each produces opposite verdicts, and consequently very different recommended uses of NPS.

Diagnostic step	AlphaCloud · B2B SaaS, mid-market	OrbitMobile · Regional telecom	Score
1. Sample calibration	Respondent profile matches base on tenure, size and ARR.	Respondents skew towards higher-tenure post-paid; pre-paid under-represented.	**Pass / Fail**
2. Lock-in adjustment	No contracts > 12 months; correlation holds in renewable cohort.	78% on 24-month contracts; correlation collapses once locked-in cohort is removed.	**Pass / Fail**
3. Behavioural cross-check	Share-of-wallet and renewal rate both move with NPS bands.	Detractors and promoters consume the network at similar volumes.	**Pass / Fail**
4. Segment-level test	Correlation present in enterprise, mid-market and SMB segments.	Correlation present only in pre-paid; absent in post-paid.	**Pass / Mixed**
5. Time-lag check	NPS at month 0 predicts retention at month 9 (peak lag 9–12 months).	No meaningful correlation at any lag beyond month 0.	**Pass / Fail**
Total score	**5 / 5 checks pass**	**0.5 / 5 checks pass**	—
Verdict	**NPS is a genuine leading indicator. Safe to use for forecasting and incentives.**	**NPS is a sentiment metric, not a retention predictor. Do not tie executive bonuses to it.**	—

5 / 5

AlphaCloud — every check passes. Use NPS as a leading retention indicator.

0.5 / 5

OrbitMobile — only a mixed pass on segment test. Brand-health metric at best.

Same NPS

Both score +28 NPS. The headline number, in isolation, means very little.

PART VI

Five things to do this week

1. Stop pretending the correlation is universal

If your communications, dashboards or board packs still describe "NPS drives retention" as a general rule, soften the language now. Reframe it as "NPS is correlated with retention in the parts of our business where customers actively choose us each cycle". That is both scientifically accurate and operationally useful.

2. Run the five-step diagnostic — once, properly

Pick a single business unit, take the last 24 months of NPS and retention data, and run the diagnostic from Part IV. The exercise typically costs a week of analyst time. It produces the single most valuable artefact a CX programme can have: a defensible statement about what its headline metric actually means.

3. Pair NPS with one behavioural metric per journey

For each major customer journey (onboarding, support, renewal, growth) define one behavioural metric that moves when the experience genuinely improves: time-to-first-value, complaint resolution rate, expansion deal velocity, etc. Report the two side by side. If they diverge, you have a problem worth diagnosing.

4. Decouple bonuses from headline NPS

If executives are compensated on NPS, move the incentive to a composite that includes at least one behavioural metric. The total compensation stays the same; the incentive simply stops pointing at sentiment alone. This change is usually politically easier than it sounds, because finance teams welcome it.

5. Audit your response rate before doing anything else

Response bias is the single most common source of false correlation we encounter. If your response rates are below 25%, fix that first — through better channels, shorter surveys, and better timing — before drawing any conclusions about whether NPS predicts retention.

A short checklist

Before relying on the NPS-retention link in any major decision, run it through this checklist. If you cannot answer "yes" to all ten points, the headline correlation is circumstantial and should be treated as such.

- Response rate is above 25%, ideally above 35%.
- Respondent profile matches the customer base on at least three demographic axes.
- Correlation has been tested with the locked-in cohort removed.
- Correlation holds within each major customer segment, not just in aggregate.
- At least one behavioural metric moves alongside NPS (repeat purchase, share-of-wallet, etc.).
- NPS at month 0 correlates with retention measured 6–12 months later.
- The same diagnostic has been run for each major business unit independently.
- Executive incentives are tied to a composite metric, not NPS alone.
- Closed-loop response rate on detractors is tracked and above 50%.
- The diagnostic is repeated at least annually.

CLOSING THOUGHTS

What to do tomorrow morning

The point of this paper is not to retire NPS, and certainly not to replace it with something more elaborate. NPS is simple, cheap and well understood — it would be foolish to throw those properties away. The point is to use NPS with the same scientific honesty we would use for any other indicator: knowing when it works, knowing when it doesn't, and being willing to say so.

Three principles that tie the whole argument together

Sentiment is not behaviour.

The moment we conflate the two, we lose the ability to tell whether a CX programme is working. Always carry one behavioural metric per journey alongside the sentiment one.

Correlation is not causation — and not always even correlation.

A correlation observed in your data may be produced by a hidden third factor: contractual lock-in, market structure, survivor bias. The diagnostic in Part IV is how to find out.

A useful metric is one you can be wrong about.

If your dashboard never tells you that NPS has stopped predicting retention, the dashboard is decorative, not analytical. Build the diagnostic into the annual programme rhythm.

How InsightSofa can help

InsightSofa is a European customer experience management platform built around precisely this kind of analytical discipline. We capture NPS, CSAT, CES and our own OES across every channel — but we treat the score as the start of the investigation, not the end. Where it matters, we run the diagnostic from Part IV for our clients on their own data.

Where teams typically start with us:

- **NPS health check** — a single-week diagnostic of whether NPS predicts behaviour in your business.
- **Response-rate uplift** — multi-channel measurement (email, SMS, WhatsApp, in-app, terminals).
- **Behavioural pairing** — connecting NPS to renewal, share-of-wallet and complaint volume.
- **Segment dashboards** — Simpson's-paradox-proof reporting at executive level.
- **Closed-loop programme** — converting detractor feedback into operational actions.

Start with a one-week diagnostic

If you would like us to run the five-step diagnostic on your own NPS and retention data, we can typically do it within one calendar week. No obligation, no presentation — only the analysis and the answer. Reach us at **+420 777 661 368** or **info@insightsofa.com**.

Sources referenced in this paper. Fred Reichheld, *The One Number You Need to Grow*, Harvard Business Review (December 2003). Fred Reichheld & Rob Markey, *The Ultimate Question 2.0*, Harvard Business Press (2011). Timothy L. Keiningham, Bruce Cooil, Tor Andreassen & Lerzan Aksoy, *A Longitudinal Examination of Net Promoter and Firm Revenue Growth*, Journal of Marketing (2007). Hyun Lee et al., *The use of Net Promoter Score (NPS) to predict sales growth*, Journal of the Academy of Marketing Science (2021). Both AlphaCloud and OrbitMobile in Part V are illustrative and not representative of any specific InsightSofa client.