

Statistical validity of CX data

How many responses do you
really need?

Contents

Introduction	3
The three forces that set your minimum sample	4
The sample-size reference table	6
Confidence vs precision — the trade-off	8
The response-rate funnel	10
Slicing the data — segment sample sizes	12
Reading movements — when is a change real?	14
Six statistical mistakes CX teams make	16
Self-assessment scorecard	18
Closing thoughts	19

Introduction

Two failure modes dominate the way CX teams treat their data, and they are opposites. One team reports a 0.4-point NPS movement to the executive committee on a sample of 38 responses and recommends action. Another team, with 1,200 responses, refuses to act because the result "isn't significant enough." Both are wrong, and both are common. Statistical literacy in CX is not advanced statistics — it is knowing what minimum sample you need, knowing how to read a movement, and knowing when to stop counting.

This paper covers the working knowledge that solves both failure modes. Three forces determine your minimum sample — population, confidence and margin of error. A single reference table answers "how many do I need?" for almost every common CX scenario. Two further chapters cover response-rate planning and segment slicing — the places where most CX programmes accidentally under-sample. The last technical chapter covers how to read a movement: when a change is real and when it is noise.

None of this requires a statistics degree. The aim of the paper is to give a CX manager enough working knowledge to defend a sample-size decision to a CFO — and enough scepticism to push back when an executive demands action on a sample of forty.

What you'll get:

- The three forces that determine any minimum sample
- A reference table covering populations from 100 to 1,000,000
- A formula for working backwards from response rate to invitations sent
- Adjusted sample sizes for segment-level analysis
- A rule for reading period-to-period movements without false alarms
- Six common statistical mistakes — and how to avoid each

Who this is for:

- CX directors defending sample sizes to executives
- VoC programme leads designing transactional surveys
- Analysts reading dashboards and making the call on what's signal
- Anyone who has been asked "is this significant?" and froze



Chapter 1

Three forces

Three knobs — and one common misconception

Almost every "how big should the sample be?" question reduces to three settings. None of them is your population size — which is the most common misconception in CX. The required sample barely changes once your population is above a few thousand.

1. Confidence level. How sure do you want to be that your result reflects reality? The CX standard is 95% — meaning if you ran the survey 100 times, the true value would be inside your interval in 95 of them. 90% is fine for early-stage scanning; 99% is overkill for almost everything.

2. Margin of error. How precise do you need to be? ± 5 percentage points is the CX default and works for tracking. ± 3 points is needed when the decision


3. Population size. Counter-intuitively, this matters less than people think. A population of 10,000 needs almost the same sample as a population of 1,000,000. Below $\sim 1,000$ it does matter; above $\sim 10,000$ you can usually ignore it.

A hidden fourth variable: expected distribution

Formulas assume a worst-case distribution (50/50). If you know from prior data that your result is closer to 80/20 (e.g., 80% promoters), you can get away with smaller samples. In practice CX teams use the 50/50 assumption — it is safer and the maths is cleaner. Adjust only when you have stable historical data to lean on.

The working defaults

For 95% confidence and ± 5 percentage points on a population above 10,000, you need 385 responses. For ± 3 points at the same confidence, you need roughly 1,070. Memorise those two numbers — they cover the majority of CX use cases.



Chapter 2

Sample-size table


Required sample size by population and precision

The table below is the headline artifact of this paper. Read the row matching your population, the column matching your required precision, and the number in the cell is the minimum number of completed responses — not invitations sent — to defend a 95% confidence interval.

Population (N)	±5 % · 90% conf	±5 % · 95% conf	±3 % · 95% conf	±3 % · 99% conf	Notes
100	74	80	92	99	<i>Almost everyone</i>
500	176	218	341	446	<i>Most B2B accounts</i>
1,000	214	279	517	747	<i>Boutique B2C</i>
5,000	258	358	881	1,485	<i>Mid-tier retail</i>
10,000	264	371	965	1,719	<i>Common ceiling</i>
50,000	270	382	1,046	1,945	<i>Almost identical</i>
100,000	270	384	1,057	1,975	<i>to 10,000</i>
1,000,000+	271	385	1,068	2,000	<i>Asymptote</i>

How to read this table

The highlighted column (95 % confidence, ±5 % margin) is the CX default — pick that unless you have a specific reason not to. Notice the asymptote: once your population is above ~10,000, the required sample barely moves. That is the single most useful fact in this paper.



Chapter 3

Confidence vs precision

How to pick the right confidence and margin

The two knobs interact non-linearly. Going from 90 % to 95 % confidence increases the required sample by roughly 40 %. Going from ± 5 % to ± 3 % increases it by roughly 280 %. Most teams over-engineer confidence and under-engineer margin — when in fact, margin matters more for most CX decisions.

When to use which combination

90 % conf, ± 5 %. Early-stage programmes or directional reads. Quick monthly scanning. Roughly 271 responses for any large population.


95 % conf, ± 5 %. Default for ongoing CX tracking. The number you should be able to defend at any executive review. 385 responses for any large population.

95 % conf, ± 3 %. When the result will trigger investment decisions or performance reviews. 1,068 responses for any large population.

What it costs you

Setting	Sample required	vs. baseline
90 % conf, ± 5 %	271	0.70x baseline
95 % conf, ± 5 %	385	1.00x (baseline)
95 % conf, ± 3 %	1,068	2.77x baseline
99 % conf, ± 5 %	666	1.73x baseline
99 % conf, ± 3 %	1,849	4.80x baseline
99 % conf, ± 1 %	16,641	43x baseline

Pursuing 99 % confidence at ± 1 % requires forty-three times the sample of the default — for a precision gain that no real CX decision needs.



Chapter 4

Response-rate funnel

Working backwards from completed responses

The sample size from Chapter 2 is the number of completed responses — not the number of invitations you send. To know how many invitations are needed, you need an honest estimate of your response rate, which is almost always lower than CX teams expect.

Working formula

Invitations needed = Required responses ÷ Response rate.

Typical response rates

Cold email survey	3 — 8 %
Warm email (existing customer list)	5 — 15 %
Post-transaction email prompt	15 — 30 %
Embedded in-product widget	20 — 40 %
B2B with personal contact	30 — 60 %
Phone (CATI) survey	20 — 35 %

Worked example


Suppose you need 385 completed responses (95 % conf, ±5 %, large population) and you are running a post-transaction email prompt with an expected response rate of 18 %.

Invitations needed = $385 \div 0.18 \approx 2,139$ prompts sent.

Two practical implications

Use the lower end of the response-rate range when sizing the invitation list. Response rates drop over time as fatigue sets in; the number that worked last year will be lower this year. Plan for the worst case.

If your population is small (< 1,000) and your response rate is also low (< 10 %), you will not get to a defensible sample. Either shift channel — phone, in-person, in-product — or accept that this group cannot carry a 95 %/±5 % conclusion on its own.



Chapter 5

Slicing the data

The trap: one number, multiplied

The 385-response rule applies to your headline metric. The moment you slice the data — by region, segment, channel, persona — each slice needs its own 385 responses if you want the same precision for each. This is the single largest under-sampling trap in CX.

The arithmetic

If you want 95 % confidence / ± 5 % margin on each of K segments, the total sample is roughly $K \times 385$. A typical CX programme that wants to slice by five customer segments and three channels — fifteen cells — needs around 5,775 responses, not 385.

Three honest options


- Accept lower precision for segment-level reads. ± 10 % at 95 % takes only 96 responses per segment — fine for directional steering.
- Pre-define a small number of important segments (3-5), not all the ones marketing might want.
- Use rolling 12-month data for segments, not just last month — accumulates more responses per cell.

Per-segment requirement at common precisions

Margin	95 % conf	90 % conf
± 10 percentage points	96	68
± 7 percentage points	196	138
± 5 percentage points	385	271
± 3 percentage points	1,068	752
± 2 percentage points	2,401	1,691

For segment-level reads, ± 10 percentage points at 95 % is the honest default. The result will move 4x faster than your headline metric — that is feedback, not noise.

Reading movements



The smallest movement worth reporting

Your monthly NPS dropped from 41 to 37. Is that real, or noise? The honest answer depends on three things: how many responses are behind each number, the variability inside each period, and whether you are testing against a single threshold or fishing for any change.

A practical rule for NPS

For NPS, a movement under approximately 5 points between two monthly readings of ~400 responses each is within statistical noise. Anything below that is not worth reporting as "down" or "up" without further investigation.

The rule of thumb table below

Minimum real NPS movement at 95 % confidence, given equal sample sizes in both periods.

Minimum real movement (NPS)

Responses per period	Min movement
100 each	~14 NPS points
200 each	~10 NPS points
400 each	~7 NPS points
1,000 each	~4 NPS points
5,000 each	~2 NPS points
10,000 each	~1.5 NPS points

How to report movements

Three forms work. "Within statistical noise" — when the change is below the threshold. "Material movement" — when it exceeds the threshold but not by a wide margin. "Significant shift" — when it exceeds the threshold by a factor of two or more. Use these terms consistently in dashboards.



Chapter 7

Six statistical mistakes

Six mistakes CX teams make with their numbers

These are the patterns we see most often. None is a maths problem — every one of them is an operating-discipline problem. They get fixed by a small number of rules, applied consistently across every dashboard the team ships.

01 "We have 38 responses — that's plenty."

It is not. Below ~270 responses you cannot defend any conclusion at 95 % confidence with reasonable precision. "Directional read" is the honest label; treat it as such.

02 Reporting a movement without a confidence interval.

A point estimate without uncertainty around it is misleading by default. Every reported metric should travel with its margin of error — even if only in the footnote.

03 Comparing periods without a significance check.

A drop from 7.2 to 6.9 on small samples is almost certainly noise. Compare against the minimum-movement table in Chapter 6 before reporting anything as a trend.

04 Using the whole-population sample for segment reads.


If the 385-response sample is your only data, you cannot read segments from it. Either accept lower per-segment precision or collect more responses per cell.

05 Demanding 99 % confidence on everything.

Over-engineered confidence triples your sample-size cost for precision that no real CX decision requires. Default to 95 %; go to 99 % only for results that will drive material spend.

06 Stopping the survey once you hit the magic number.

Once a programme is running, keep going. Continuous data lets you spot real movement earlier; capping at the minimum sample creates artificial blind spots between collection windows.



Chapter 8

Self-assessment

Are your CX numbers actually defensible?

Twenty statements across the five areas of statistical discipline. Score one point for each statement you can honestly answer "true today" — not "in progress."

HEADLINE METRIC

- We have at least 385 responses per period for our headline metric.
- The 95 % confidence interval is reported alongside the headline.
- We use the same survey instrument across periods (no version drift).
- We can defend our sample-size choice to a CFO if asked.

RESPONSE-RATE PLANNING

- We model invitations sent from required responses, not the other way around.
- Our response rate assumption is taken from the lower end of historical range.
- Response rates are tracked monthly; we know when fatigue is setting in.
- When response rates fall below a threshold, the survey is paused or redesigned.

SEGMENT READS

- Each segment we report has its own minimum sample, not borrowed from the total.
- We pre-defined a small number of segments — not 20 — before launching.
- Segment-level precision is reported as $\pm X$ percentage points, not just a number.
- Rolling windows are used for segments to accumulate sufficient responses.

MOVEMENT INTERPRETATION

- Every reported change is checked against a minimum-movement threshold.
- We use consistent language: 'within noise', 'material', 'significant shift'.
- Period-to-period comparisons include the test-statistic or p-value where applicable.
- Single-point movements are never treated as a trend.

REPORTING DISCIPLINE

- Dashboards show confidence intervals or margins, not just point estimates.
- Sample sizes are visible on every chart that drives a decision.
- We have a written threshold for when a movement triggers action.
- Stakeholders know what "directional" means in our dashboards.

Interpreting your score

16-20 · Defensible programme; CFO and CEO can rely on the numbers. 10-15 · The headline metric is sound but segment and movement reads need discipline. Below 10 · You are reporting noise as signal; start with the headline sample and the movement threshold.

Closing thoughts

Statistical discipline in CX is not about being clever with formulas. It is about being honest with the data you have. The teams that defend their numbers to a sceptical CFO are not the ones running the most sophisticated maths — they are the ones applying three rules consistently.

Know your minimum sample before you launch. Look up the row in the Chapter 2 table, model the invitations backwards from your response rate, and only then design the survey. Not after.

Travel the metric with its uncertainty. Every number you publish carries a margin of error. Showing it is not an admission of weakness — it is the difference between a defensible report and a guess.

Treat segment and movement reads with extra care. These are where most CX teams quietly over-claim. Use the per-segment table for slicing and the minimum-movement table for trends. If the data does not support the claim, do not make the claim.

How InsightSofa can help.

InsightSofa applies these rules automatically. Confidence intervals appear on every chart by default. Sample sizes are visible everywhere. Movement reads are flagged against the minimum-movement threshold for the underlying sample. Segment dashboards refuse to display reads below the defensibility threshold — so your team cannot accidentally over-claim. Across roughly forty client engagements, the most common early win is the same: an executive who finally trusts the CX numbers because the dashboard refuses to lie.

The fastest start is a 30-minute working session with our CX strategy team.

Contact +420 777 661 368 · insightsofa.com

Sample-size figures use the standard proportion formula with $p = 0.5$ (worst case) and finite population correction. NPS movement thresholds are approximate, derived from the standard error of the difference between two NPS scores with equal sample sizes.