KNOWSIS

QUESTIONNAIRE AUDIT

Customer Satisfaction Tracker

Quick-Service Restaurant Chain

Prepared for **[Client Name]**

Knowsis | May 2026

What this audit covers

A desk audit of the questionnaire as a written artefact: question construction, ordering, routing logic, scale design, and analytical readiness against the stated objectives. It does not cover implementation correctness on the survey platform (skip logic firing as designed, quotas behaving, mobile rendering), which needs platform-specific testing once the design is settled.

Summary

The questionnaire is competently structured and covers the territory you would expect of a monthly QSR satisfaction tracker: visit context, overall evaluation, attribute ratings, food quality diagnosis, pricing perception, competitor benchmarking, recovery, loyalty intent, and customer profile. The flow is broadly correct and the length is realistic.

The defects concentrate in three areas. First, three separate items measure value for money on overlapping constructs, which will inflate that attribute's influence in any driver analysis through collinearity. Second, the routing has gaps around comparative judgements (Q13, Q16-Q18) that reference comparator sets not yet established. Third, several open and evaluative items have ambiguous referents or unclear scope.

Once the critical findings below are addressed, the questionnaire will deliver cleanly on its stated objectives and the driver analysis it is structured to support will produce defensible prioritisation rather than a measurement artefact.

Research Design Check

The instrument is a single questionnaire to one audience (recent customers, screened via Q0 to the past 14 days) with a monthly cross-sectional cadence. That design is appropriate for the stated objectives and there is no case for splitting into multiple instruments.

Two design parameters are not stated and should be:

- **Sample size and per-location quotas.** Location-level food quality diagnosis is one of the stated objectives. That requires a meaningful number of responses per location each month, not just a total sample. With 800 monthly responses across 85 locations, the average is around 9 per location, which is too thin for stable monthly diagnosis without quotas. The same applies to driver analysis: with nine attributes in Q8, the analysis cell needs around 200+ respondents.
- **Mode.** The format reads as online self-complete via SMS link, but mobile-specific design implications (grid handling, attention checks, response option count) are not addressed in the implementation notes.

What is Working

Before the issues, the design has real strengths worth keeping:

- The funnel from category context through overall evaluation to specific attributes is broadly correct. Overall satisfaction (Q5) is asked before the attribute grid (Q8), which avoids priming the headline metric with attribute-specific thinking.
- Time anchoring is clean. Q2 uses concrete buckets ("Today", "2 to 3 days ago") rather than vague quantifiers, and Q4 frequency options are specific.
- The recovery section (Q19 to Q21) is properly gated, with a clean skip path for respondents with no problems.
- Q23 / Q24 (visiting less often, gated open-end) is well-constructed and will produce useful diagnostic copy.
- Q27 as a closing open-end is positioned correctly: late, after rapport, framed to capture residual signal.

Critical Findings

These are structural defects. Each one corrupts data, not just style.

CRITICAL Finding 1. Q7 has an ambiguous referent

Q5 is overall satisfaction (1 to 5). Q6 is NPS (0 to 10). Q7 (*"What is the main reason for your score?"*) is sandwiched between them. Some respondents will explain their satisfaction score, some their NPS score, some both. The downstream open-text analysis cannot disentangle them.

Either move Q7 immediately under Q5 and reword as "main reason for your satisfaction score", or move it under Q6 and reword for NPS. Most practitioners attach the diagnostic open-end to NPS because that is where it earns its keep commercially.

CRITICAL Finding 2. Triplicate measurement of value for money

Value for money appears three times on different scales:

- Q8 attribute grid: "Overall value for money" (1 to 5)

- Q12: "How would you rate the value for money of your meal?" (1 to 5)
- Q14: "How acceptable do you find our pricing overall?" (1 to 5)

Q12 is essentially a re-ask of the Q8 attribute. Q14 measures something subtly different (price acceptability rather than perceived value). Aside from respondent burden, this creates a serious analytical problem: when value for money sits in the driver analysis as an independent variable, double measurement on the same construct will inflate its apparent influence and crowd out other attributes through collinearity. The driver analysis will then mislead the prioritisation.

Pick one canonical value-for-money item for the tracker (Q12 sits in a pricing-focused section and is the natural choice), drop the duplicate from Q8, and either retain Q14 as a distinct construct (price acceptability) or remove it.

CRITICAL Finding 3. Q8 contains a double-barrelled item

"Food quality (freshness, temperature, consistency)" is three constructs in one rating. A respondent who finds the food fresh but cold, or hot but inconsistent across visits, has no clean way to answer. The downstream diagnostic value is also compromised: a poor rating gives no signal about which of the three drove it.

Either split into three separate items (freshness, temperature, consistency) or pick the construct that matters most to the brand and ask only that. The other two can sit in the food quality diagnostic section.

CRITICAL Finding 4. Comparative questions reference unestablished comparison sets

Q13, Q16, Q17 and Q18 all rely on the respondent having a meaningful frame of comparison.

- **Q13** ("Compared to similar restaurants, do you feel our prices are...") is asked before any competitor visits have been established. The respondent has to assemble their own comparison set on the fly, which produces noise.
- **Q16 / Q17 / Q18** explicitly reference "these restaurants" or "other similar restaurants", but Q15 (the competitor question that establishes the set) allows "None of these". A respondent who selects "None of these" will hit Q16 with no referent and either drop out or default-answer.

Two fixes. First, move Q15 (competitor visits) before Q13. Second, gate Q16, Q17 and Q18 to respondents who selected at least one brand in Q15.

CRITICAL Finding 5. Q11 is not gated to respondents who reported issues

"If you experienced any issues, please describe them briefly." The "if" is doing all the work, which is exactly the kind of routing logic that should be enforced by the platform rather than left to the respondent.

Respondents who selected "No issues" in Q10 will see this question and either leave it blank, invent something, or interpret it as a general complaint prompt. Gate Q11 to respondents who selected anything in Q10 other than "No issues".

CRITICAL Finding 6. Q19 overlaps with Q10 without a clear scope distinction

Q10 captures food-specific issues. Q19 ("Did you experience any problem during your visit?") captures problems more broadly. The relationship is not made explicit, so a respondent who flagged a food issue in Q10 may say "No" in Q19 (already covered) or "Yes" (problem was real). Both readings are reasonable.

Reword Q19 to make the scope clear, for example: "Beyond any food issues already mentioned, did you experience any other problems during your visit?" That way Q10 owns food, Q19 owns everything else, and the recovery section is properly anchored.

Important Findings

These will not corrupt the dataset, but they will weaken the analysis.

IMPORTANT Finding 7. Grid hygiene in Q8

Q8 is a 9-item rating grid on a 5-point scale. Without disruption, this is exactly the structure that produces straight-lining on mobile, which is where most responses will come from. Three small additions will materially improve the data:

- Randomise the row order (the implementation note already mentions this; make sure it actually fires).
- Insert one attention check inside or immediately after the grid ("To show you are paying attention, please rate this row as 4").
- Consider including one reverse-coded item ("It was hard to find what I wanted on the menu") so straight-liners can be identified for cleaning.

IMPORTANT Finding 8. Q9 is largely redundant with Q8

Once Q8's "Food quality" item is split or scoped properly, Q9 covers the same ground (fresh, hot, taste, consistency, looks appealing) in multi-select format. The multi-select gives less analytical power per attribute than the rated grid, and the duplication adds length without adding signal. Remove Q9 entirely and let Q8 carry the food attribute load, with Q10 handling specific issue diagnosis.

IMPORTANT Finding 9. Region list is too coarse for South African geographic analysis

Q25 lists Gauteng, Western Cape, KwaZulu-Natal and "Other". That collapses Eastern Cape, Free State, Mpumalanga, North West, Northern Cape and Limpopo into a single bucket. If location-level diagnosis or regional cuts are part of the use case, this will not support them. The full nine-province list adds five seconds to the questionnaire and substantially more analytical headroom.

IMPORTANT **Finding 10. Consent text is too light for POPIA**

The introduction states the survey is confidential and may be used for follow-up. POPIA expects more: explicit informed consent, stated purpose, retention period, and a contact for queries. Tightening this up is half a paragraph of work and removes a compliance risk that grows as the dataset accumulates over monthly waves.

Routing Audit Summary

Pulling the routing issues into one place for the implementer:

Question	Issue
Q7	Ambiguous referent (which score?). Reword and reposition under the score it is meant to explain.
Q11	Not gated. Should be conditional on Q10 selecting any issue.
Q13	Makes a comparative judgement before the comparator set has been established. Move Q15 earlier.
Q16-Q18	Gate to respondents who selected at least one brand in Q15.
Q19	Reword to clarify scope (non-food problems only, or everything-including-food).
Q24	Correctly gated to Q23 = "Less often". No change needed.

Analytical Readiness

Walking the stated objectives one by one:

Track NPS, satisfaction, value perception monthly. Q5, Q6 and one of (Q12, Q14) cover this. Once the value-for-money triplication is resolved, the headline metrics are clean.

Diagnose food quality issues at location level. Q1 plus Q8 / Q10 / Q11 (once Q11 is gated) supports this in principle. The constraint is sample size per location per month. Without per-location quotas, the diagnosis will only work for high-volume locations, and the long tail will need to roll up over multiple months to be analysable.

Identify drivers of satisfaction. This is where the questionnaire is structurally well-set-up but where the redundancy issues bite. Q5 supplies the dependent variable and Q8 supplies the independent variable battery. Once the double-barrel in Q8 is split and the duplicate value-for-money item is removed, the structure is exactly what driver analysis requires: ranking attributes by importance and quadranting them for prioritisation. Knowsis offers driver analysis as a service to clients and through the Data Alchemy Platform for self-service users.

Monitor pricing perception drift. Pick one canonical pricing item per the value-for-money note above. Tracking three pricing items will produce three drift lines that move differently for measurement reasons rather than market reasons.

Benchmark against competitors. Q15 to Q18 do this descriptively. Q17 / Q18 ("which areas do we perform better / worse") produces multi-select counts, fine for a top-line read but cannot quantify the gap. If competitive benchmarking will influence brand positioning, the section needs more analytical depth than it currently has.

Enable targeted customer recovery. Q19 to Q21 plus the implementation note flagging Detractors and low-satisfaction respondents covers this well.

Suggested Order of Fixes

If you want a prioritisation rather than a flat list:

1. Fix the Q7 referent. One-minute fix. Removes a measurement ambiguity.
2. Fix the value-for-money triplication. Pick one item, drop the others. This is the single most important fix for the driver analysis to be trustworthy.
3. Split the Q8 double-barrel.
4. Re-order Q15 before Q13 and gate Q16 to Q18 properly.
5. Gate Q11 to respondents with issues.
6. Reword Q19 for scope clarity.
7. Add an attention check, and ideally one reverse-coded item, into Q8.
8. Expand Q25 to all nine provinces, and decide what additional demographics belong in the tracker.
9. Tighten the consent text for POPIA.
10. Remove Q9 once Q8 is cleaned up.

Once these are in, the questionnaire will deliver on its stated objectives cleanly, and the driver analysis it is designed to support will produce defensible prioritisation rather than an artefact of measurement redundancy.

Where to take this further

This audit identified driver analysis as the highest-value next step on this dataset. A few ways we can help you act on the findings:

- **Run the analysis for you.** Once the questionnaire fixes are in and the data is collected, Knowsis can take it from there: driver analysis, segmentation, composite indices, and the analytical narrative your stakeholders will ask for. Book a 30-minute conversation to discuss what your data could deliver.

- **Do the analysis yourself.** If you run surveys regularly, the **Knowsis Data Alchemy Platform** packages a full set of analytical tools into a self-service flow. It also includes the Data Alchemist, a built-in AI assistant you can ask analytical questions of and get guidance on the best approach

for your specific data. **Sign up for a free trial**.

- **Talk through the findings.** If you want to walk through this audit and discuss what to prioritise, **book a free 30-minute call**.

Boundary

This audit covers the questionnaire as a written artefact: wording, structure, routing intent, analytical readiness. It does not cover:

- Implementation correctness in the survey platform. Whether the platform enforces skips, whether randomisation is implemented as specified, whether quotas conflict with routing. This needs platform-specific QA: a soft launch and a manual respondent walk.
- Sample design. Whether the panel composition matches the analytical sub-groups intended. Sample sizing for sub-group analysis was flagged but not fully specified.
- Cleaning and weighting rules. These should be set before fielding, not retrofitted.

Run a pilot of 10 to 30 responses before fielding the next wave at scale. The pilot will catch comprehension issues and surface the Q8 problems described above in the data, which is more persuasive than a desk review when there is internal disagreement about whether the changes are worth making.

Greg Streatfield

Founder & Chief Data Alchemist, Knowsis

greg@knowsis.net | knowsis.net | apps.knowsis.work