

The cash caseload, and the rule that produced it

Food security classification and assistance targeting

Organisation	Food security and livelihoods programme across four districts
Prepared for	Food security cluster and the cash working group
Updated	July 26, 2026
Sectors	Food Security
Standards and methodologies	Integrated Food Security Phase Classification (IPC) · Sphere Standards · Core Humanitarian Standard (CHS) · Results-Based Management (RBM)
Data	food-security-survey-2024

DECISION THIS INFORMS

Which households enter the cash assistance caseload, and on what stated criteria.

Every dataset behind this report is synthetic. No record traces to a real person, and no figure here describes a real programme.

Read the project online at data-analysis.cassion.dev/projects/food-security-classification

This report is licensed CC BY 4.0. You may adapt and redistribute it, including commercially, with attribution to Cassion.

Contents

Context	2
1 Recommendation	2
2 Why no single indicator gives the caseload	2
3 The rule	3
4 The boundary is not a finding	3
5 Who the rule includes	3
6 What this does not establish	4
7 Recommended next steps	4
About this edition	4

Context

The problem

The caseload for the next cash distribution round is fixed at roughly a fifth of surveyed households, but the share classified as food insecure moves from one percent to seven percent depending only on which Food Consumption Score thresholds are applied. The targeting list, and therefore who eats, changes with a decision nobody had written down.

1 Recommendation

Assist the **422 highest-ranked households** on the composite need score, and publish the rule alongside the list.

The caseload size is a budget decision, not a finding: 422 is a fifth of the 2,112 households surveyed, and the budget decided that. What this analysis provides is the ranking and — more importantly for a cluster review — the argument for why the line is where it is.

Three things must travel with the list, and each is a sentence the cash working group will be asked about:

- Displaced households are 19.2% of the surveyed population and 33.4% of the caseload. No displacement term appears anywhere in the score.
- Nord-Ouest takes 219 of the 422; Sud takes 11. That is a twenty-to-one ratio between districts and it will be challenged.
- Nothing in the data distinguishes the last household included from the first excluded.

2 Why no single indicator gives the caseload

The obvious approach is one threshold on one indicator. Applied to this survey it does not work, and seeing why is the argument for the composite.

Rule	Households	vs caseload of 422
FCS poor (≤ 21)	17	−405
FCS poor or borderline (≤ 35)	465	+43
HHS moderate or worse (≥ 2)	889	+467
HHS severe (≥ 4)	325	−97
rCSI at or above median	1,074	+652

$\text{FCS} \leq 21$ selects 17 households and cannot spend the budget. $\text{HHS} \geq 2$ selects more than twice the caseload. $\text{rCSI} \geq \text{median}$ selects half the survey by construction and is not a threshold at all.

$\text{FCS} \leq 35$ lands within 43 of the caseload — and **choosing it because it lands close would be fitting the criterion to the budget rather than to need**. That is exactly what a cluster review asks about, and it is why the criterion was rejected. This paragraph is the record that it was considered.

The deeper reason is that the three indicators disagree on purpose. FCS measures dietary diversity, HHS measures whether people went without, rCSI measures what they did to cope. Any single-indicator threshold discards two of the three measurements.

3 The rule

```
need score = ( -z(FCS) + z(rCSI) + z(HHS) ) / 3
```

Each indicator is standardised across the 2,112 surveyed households so none dominates by scale, and FCS enters negatively because a higher score is better consumption. Full definitions, weights and exclusion rules are in `docs/threshold-justification.md`.

Three assumptions are visible in that formula and all three belong in the report. The weights are equal, which says the three indicators matter equally — an assumption, not a finding. The standardisation is sample-relative, so the score is a ranking device inside this survey and carries no meaning across surveys. And the missing-data fills are conservative in opposite directions, so a household with an incomplete form is neither pushed into nor pulled out of the caseload by the gap itself.

The caseload it produces contains 16 of the 17 FCS-poor households and 284 of the 325 with severe hunger. The composite does not discard the single-indicator signal; it stops that signal being the only one.

4 The boundary is not a finding

The score at the cut is **0.663**. The forty households either side span **0.630 to 0.697** — under a tenth of a point — and adjacent pairs differ by thousandths.

Nothing in the data distinguishes the last household included from the first excluded. Two consequences the report must carry:

1. Do not describe the caseload as "the food-insecure households". It is the 422 highest-ranked households under a stated rule.
2. If the programme can extend at the margin, the boundary list is the list to extend with, in score order, without re-running the analysis.

5 Who the rule includes

Displacement status

Status	Population	Caseload	Difference
IDP	19.2%	33.4%	+14.2 pts
Returnee	7.2%	12.3%	+5.1 pts
Host	3.7%	3.1%	−0.6 pts
Resident	69.9%	51.2%	−18.7 pts

District

District	Surveyed	Selected	Share
Nord-Ouest	528	219	41.5%
Centre	530	131	24.7%
Artibonite	528	61	11.6%
Sud	526	11	2.1%

Female-headed households are 35.6% of the population and 38.9% of the caseload. Casual-labour households are over-represented by 8.6 points; subsistence farmers under-represented by 7.2.

These are outputs of the indicators, not targeting choices. The districts were surveyed with near-equal sample sizes, so the twenty-to-one ratio reflects measured need rather than sampling. It still has to be stated, because a district authority will read the caseload as a statement about who deserves assistance.

6 What this does not establish

The survey frame is not the population. The caseload is drawn from households the enumeration reached. Displaced people outside settlements and households absent on survey day are not in this ranking and cannot be targeted from it.

The caseload size is a budget decision. This analysis ranks need; it does not say that a fifth of households is the right number to assist.

No seasonality. A single round cannot separate a lean-season household from a chronically food-insecure one, and cash targeted on the first should not be presented as addressing the second.

This is not an IPC classification. IPC phases rest on convergence of evidence across outcome and contributing indicators with an analysis workshop. No household here should be labelled with a phase.

7 Recommended next steps

1. **Publish the rule with the list.** A targeting list without its criterion is a decision nobody can contest, and it will be contested anyway.
2. **Take the district ratio to the cluster before distribution,** with the sample sizes, so it is read as a measurement rather than an allocation.
3. **Hold the boundary list.** If the budget moves at the margin, it extends in score order and needs no new analysis.
4. **Resolve the 12 duplicate enumerations** before the next round. They cannot move a cut at 422 of 2,112, but the count should not be allowed to grow.

About this edition

This report is generated from the same source as the project page, so the two cannot drift. The online edition carries the analysis notebook, the dataset and any corrections issued since this file was produced.

Read the project online at data-analysis.cassion.dev/projects/food-security-classification

This report is licensed CC BY 4.0. You may adapt and redistribute it, including commercially, with attribution to Cassion.

Generated July 26, 2026.