DANISH REFUGEE COUNCIL

GUIDANCE NOTE

November 2026 – Copenhagen, Denmark

Forecasting Displacement for Anticipatory Action

A user guide to the AHEAD displacement predictive model



Danish Refugee Council | Lyngbyvej 100 | DK-2100 Copenhagen Ø | Phone.: 3373 5000 | drc@drc.ngo | www.drc.ngo

Executive summary

Anticipatory action aims to act ahead of a crisis, or before its impacts peak, so that people can be supported earlier. For displacement, this benefits from forward looking information of where pressure may be building. The Anticipatory Humanitarian Action for Displacement (AHEAD) model was developed by the Danish Refugee Council (DRC) to contribute to this. It produces short-term, sub-national forecasts of displacement that look three months ahead and are updated every month.

This note is a guide for the teams, partners and organisations that use these forecasts. It explains what the forecasts show, how to interpret them, how thresholds turn them into a signal for action, how well the models perform and how DRC combines them with community-level indicators. It focuses on the AHEAD models whose outputs DRC currently publishes.

The starting point is understanding what a forecast figure represents. Most AHEAD forecasts estimate the total number of internally displaced people expected to be present in an area at the end of each month. A projected rise is therefore a net change in that total, once new arrivals, returns and onward movements have been taken into account. This is a different quantity from the number of people newly displaced, and a modest net change can hide a great deal of movement. Because the model also estimates recent months for which no figures have been reported yet, it can keep a continuous picture even where data are missing or delayed.

Each forecast is also probabilistic. Alongside its central estimate, the model reports an uncertainty range, and both are helpful in order to interpret a forecast well. The central estimate gives the main signal, while the width of the range shows how much caution that signal requires. Even when the model performs well, the actual figure is expected to fall outside the 90% range about one time in ten. Since forecasts are refreshed every month, each round is best understood as the latest assessment given what was known at the time. Following the same area across rounds helps to separate persistent pressure from short-lived changes.

To support decisions, forecasts are compared with thresholds agreed in advance. In DRC's standard approach, a threshold marks a three-month increase that is unusually large compared with recent history in that country. A threshold crossing calls for closer scrutiny and preparation in line with pre-agreed plans. Setting a threshold always involves a balance between catching real increases and avoiding false alarms, and the note sets out the main options and the trade-offs between them.

To help users judge how the forecasts can best be used, each country model is evaluated regularly against what actually happened in the past, using only the information that would have been available at the time. Performance is assessed both for statistical accuracy and for whether the model would have prompted the right decisions at the threshold. Like any forecasting model, AHEAD learns from past patterns. Its forecasts need more careful interpretation when conditions change suddenly or in ways not seen before, and where displacement data are sparse.

This is also why DRC does not base anticipatory action on the AHEAD model alone. The forecasts are combined with community-level indicators, identified with communities and local actors, through pre-agreed anticipatory action plans. This brings together a forward-looking model signal and what communities observe on the ground. The plans also set out in advance how the two are weighed and which actions follow. Anticipatory action always involves acting under uncertainty, and the forecasts are one input that can help such decisions to be taken earlier and on a firmer basis.

This is a live document. It will be updated as the models are refined and extended to new contexts, and as other modelling approaches currently being tested are evaluated.

Content

Executive summary **2**

Introduction **4**

What the AHEAD model forecasts show **4**

How the forecasts are produced and shared 4

Observed, estimated and projected values 7

What the forecast numbers mean 8

Forecasts are probabilistic 9

Interpreting forecast rounds over time 9

Thresholds based on the AHEAD predictive model **11**

Considerations for setting a threshold 11

Forecast thresholds: the percentile approach 12

Interpreting a threshold crossing 13

Evaluating model performance **15**

Statistical accuracy: how close are the forecasts? 15

Anticipatory-action accuracy: performance at the threshold 16

Where forecasts call for more careful interpretation 18

Combining the AHEAD model with community indicators **18**

How the two layers combine 19

Technical annex **21**

The modelling framework 21

Model variants 22

Other modelling approaches being tested 22

Introduction

Anticipatory action is an established and growing part of humanitarian practice. Its foundation is strongest for climate-related hazards, where forecasting methods are relatively well developed. Applying the same logic to displacement, particularly displacement driven by conflict, is inherently more complex. Forecasting is one way to support anticipatory action for displacement, by signalling where risk may be building so that it can be acted on earlier.

Since 2021, the Danish Refugee Council (DRC) has been a global leader in applying forecast-based, proactive interventions in displacement settings, linking displacement early warning models with community-led action. One of these tools is the Anticipatory Humanitarian Action for Displacement (AHEAD) model. It forecasts displacement at sub-national level, usually admin 1 or admin 2, around three months ahead, and was first piloted in Burkina Faso and South Sudan to bring prediction closer to operational decision-making.

With support from the Dutch Ministry of Foreign Affairs, the AHEAD model has since grown well beyond its first pilots, into Niger, Mali, the Democratic Republic of the Congo, Ethiopia, Sudan, Somalia, Uganda, Yemen and Myanmar. Across most of these countries DRC is now connecting the forecasts to anticipatory action frameworks that are locally grounded and shaped by a range of stakeholders.

As this work expands, so does the need to make sure the forecasts are understood and used consistently by the teams and partners who rely on them. A forecast can flag where displacement risk is rising, open earlier conversations about preparedness and anticipatory action and point to where more analysis is needed. But it does this as one part of a larger system, and it is only useful when it's clear what it shows, how to interpret it, where its limits lie and how it works together with what communities observe on the ground.

This note is a user guide for the AHEAD model. It sets out how the forecasts are produced and interpreted, and how thresholds turn them into a signal for action. It is written for end users, including country and regional teams, programme, anticipatory action and emergency colleagues, and partners who draw on the forecasts within wider anticipatory action processes. The note is deliberately general, and how the forecasts feed into decisions varies by country, depending on context, data, response capacity and existing plans.

Importantly, DRC's approach is not to base anticipatory action on the predictive model alone. The forecasts are used alongside community-level indicators, the signs that people on the ground observe before displacement, so that decisions are both data-informed and locally grounded. How those indicators are developed and used is summarised in the final section of this note.

This note is a live document. It focuses on the AHEAD models whose outputs DRC currently publishes, and it will be updated as these models are refined and extended to new contexts and as other modelling approaches, some of which are outlined in the technical annex, are tested.

What the AHEAD model forecasts show

How the forecasts are produced and shared

DRC produces short-term subnational displacement forecasts on a rolling monthly basis. Around the middle of each month, once the latest displacement and contextual data have been processed, the AHEAD model runs and updates the forecasts.

The AHEAD model is a statistical model. In simple terms, it looks at how displacement has changed in the past and how those changes have been linked to conditions such as conflict, food security, climate and protection, and it uses those patterns, together with the most recent available data, to project how the number of displaced people may change over the coming months. For each country, the model learns the relationship between displacement and these conditions from that country's own history and then applies what it has learned to project forward. It is built within a Bayesian statistical framework, so each forecast comes with an estimate of its own uncertainty, and it is designed to handle the irregular and incomplete displacement data that are common in crisis settings.

A few things about how the model works are worth understanding, as they shape how the forecasts should be interpreted. First, the model estimates the relationship between conditions and displacement from the full history of each area, working on the assumption that this relationship is reasonably consistent across the country and over time. This is what allows it to be applied to the most recent data to project forward. Where conditions change in ways not seen before, this relationship can shift

and the forecasts may become less reliable. Monthly updates and regular re-fitting help the models take account of new data, but they cannot fully anticipate changes that have no precedent, which is one reason the forecasts are used alongside community indicators. The section on evaluating model performance describes where forecasts call for more careful interpretation. Where relevant to a given country, the model can also take into account conditions in neighbouring areas as well as in the area being forecast.

What can the forecasts help with?

By providing a forward-looking view, AHEAD displacement forecasts can help to:

- flag areas where displacement pressure may rise before the increase shows up in reported figures;
- prioritise where and when to look more closely, for example where to intensify data collection or field analysis;
- start preparedness discussions earlier and prioritise preparedness resources across areas, within pre-agreed anticipatory action plans;
- give a consistent picture across areas, including where recent displacement data are missing or delayed.

On their own, and in their current form, the forecasts cannot:

- say with certainty whether, when or exactly how much displacement will happen;
- measure humanitarian severity, vulnerability or needs;
- fully anticipate sudden shocks or dynamics that have no precedent in the data.

Second, because the forecast looks several months ahead, the model needs values for the drivers over those future months, and these are not yet known. Assumptions therefore have to be made about how conditions such as conflict or food insecurity are likely to develop over the forecast window. These assumptions can draw on existing forecasts, recent trends and, where relevant, qualitative and contextual judgement about the situation. This does not mean the model assumes it knows exactly what will happen. It means each forecast is conditional on a set of assumptions about the drivers.

How driver levels are set for the forecast window varies by context and can change from one round to the next. One common situation is where recent averages or trends understate what may come. If the forecast assumes that conflict over the coming three months will stay close to its recent average, while country office analysis and external reporting point to a likely escalation, for example armed groups mobilising ahead of an expected offensive, the forecast may understate the risk. In such cases the assumed conflict levels can be raised to reflect the expected escalation, or the forecast can be interpreted knowing that it is likely to be on the low side. If conditions develop differently from what was assumed, the next round takes the new data into account and the forecast changes.

The model is also built around the way displacement data actually arrive. Displacement figures are frequently reported at irregular intervals, with gaps of several months and no fixed schedule. Rather than setting aside the months without a reported figure, the AHEAD model estimates them as part of the same process that produces the forecast, so a continuous monthly picture can be maintained even when reporting is incomplete. How this works, and how to tell estimated months apart from reported ones, is explained in the next section and a more technical description of the overall framework is given in the technical annex.

Most of the contextual information that feeds into the model comes from open-source humanitarian datasets, although the exact data sources differ by country. The data generally span a few broad themes, combining displacement figures themselves with indicators of the conditions associated with displacement. The table below sets out the main themes and the kinds of indicators and sources typically used. It is illustrative rather than exhaustive, and the specific indicators and sources are chosen for each country according to what is relevant locally and what data are available.

Each model update looks three months ahead. A forecast produced in mid-June covers June, July and August; the

next one, in mid-July, covers July, August and September. The window moves forward every month.

A forecast is never produced once and left to stand. It is refreshed as new information arrives, including new displacement figures, shifts in conflict, food insecurity, seasonal factors and other signals. Each round should therefore be understood as the latest assessment of likely trends, given what was known at the time.

	Example indicators	Typical sources (e.g.)
Displacement	Number of internally displaced people; number of refugees and asylum seekers; new displacement movements; rapid alerts of population movements	Displacement tracking systems (e.g. IOM DTM), CCCM, national and partner reporting; IDMC data and rapid humanitarian alert systems in some contexts
Conflict and violence	Violent events; fatalities; attacks on civilians; abductions; remote violence; violence in neighbouring areas; changes in territorial control	ACLED, UCDP, INSO and other conflict-event datasets
Food security	Share of population in higher food-insecurity phases; food prices	Cadre Harmonisé / IPC, WFP
Climate and environment	Vegetation and drought indices, seasonal and climatic stress; extreme rainfall; fires in settlements	FAO, ICPAC, remote-sensing sources (e.g. CHIRPS rainfall, NASA FIRMS fire detections)
Protection	Feelings of safety, protection incidents	Protection monitoring, including DRC's own data in some contexts

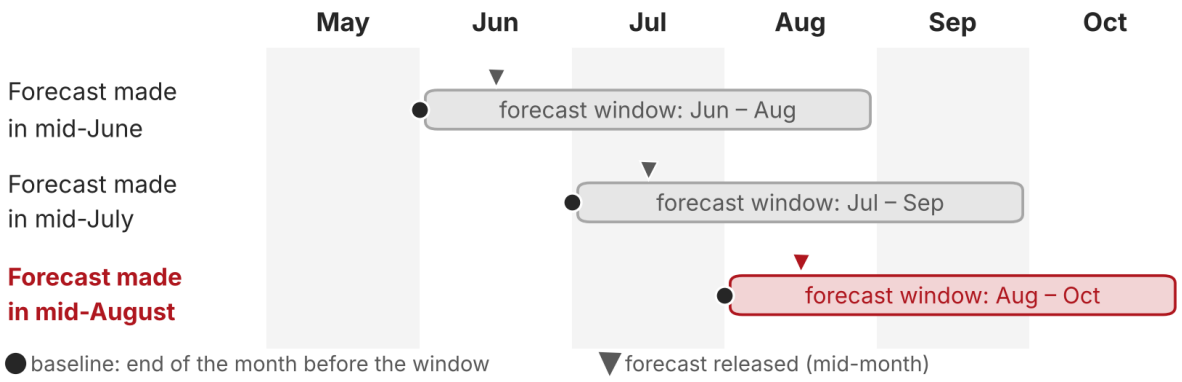


Figure 1. *Rolling forecast windows. Each month the model looks three months ahead from a new baseline, so successive rounds overlap and the window advances. The baseline is the end of the month before each window, and each forecast is released around the middle of its first forecast month.*

There is also no single "AHEAD model". There is one shared way of building it, fitted separately for each country, using that country's own displacement figures, its own administrative geography and its own mix of drivers. The model for the territories of eastern DR Congo and the model for Sudan are two distinct models that that share the same logic. Throughout this note, "the model" means the model fitted for the specific country the user is looking at. The technical annex describes the shared framework for readers interested in the technical detail.

This note focuses on the models whose outputs DRC currently publishes.

The forecasts are produced at sub-national level, with the exact administrative tier depending on the country and its data. In some places, they are produced at admin 1; in most others at admin 2; and in a few contexts at a finer level where the data allow.

Forecasts are shared in a few formats, each with a different purpose. The monthly snapshot gives a short narrative of the main projected trends, key hotspots and drivers and is meant for quick interpretation and discussion. The underlying datasets, available through HDX where relevant and through dashboards or APIs, let users dig into the numbers in detail.

Other modelling approaches

Alongside the models described in this note, DRC is testing other approaches in some contexts. These include:

- classifiers that estimate the probability that displacement in an area will exceed a given level over the coming months;
- ensembles that combine statistical and machine-learning models;
- additional data sources, including earth observation data and information extracted from text such as reports and alerts.

These approaches are evaluated against the existing models before they inform any operational decisions. The technical annex gives more detail.

Where to find the existing forecasts

Monthly snapshots. A short narrative summary of each forecast round, covering the main projected trends, hotspots and drivers: AHEAD model background info, forecast reports, etc. | Danish Refugee Council

Datasets on HDX. The forecast data by area and month, available for download for further analysis: Danish Refugee Council Humanitarian Data | 12 Datasets | HDX

Interactive dashboard. The forecast data visualized and presented in an interactive and easy to read format: AHEAD – Global Dashboard

API. The forecast data can also be accessed programmatically through an API, allowing users to integrate the forecasts directly into their own systems and analyses: Home - DRC API Management - developer portal

Observed, estimated and projected values

The model does more than project the future. It also estimates recent and current displacement levels when data are missing, delayed or irregular. To produce monthly forecasts, the model first has to build a complete monthly picture of how displacement has evolved up to now. Where an official or partner-reported figure exists, it is used directly after quality assessments. Where there is a gap, the model estimates the most likely value for that month from the available historical and contextual data and then projects forward from that continuous series.

This is why a single output can contain three kinds of values, and they should be interpreted differently. An observed value is a reported figure from a data source. An estimated value is the model's estimate for a month where data were missing or not yet available. A projected value is the model's forecast for a future month. The more of a district's recent picture rests on estimates rather than reported data, the more cautiously its forecast should be interpreted.

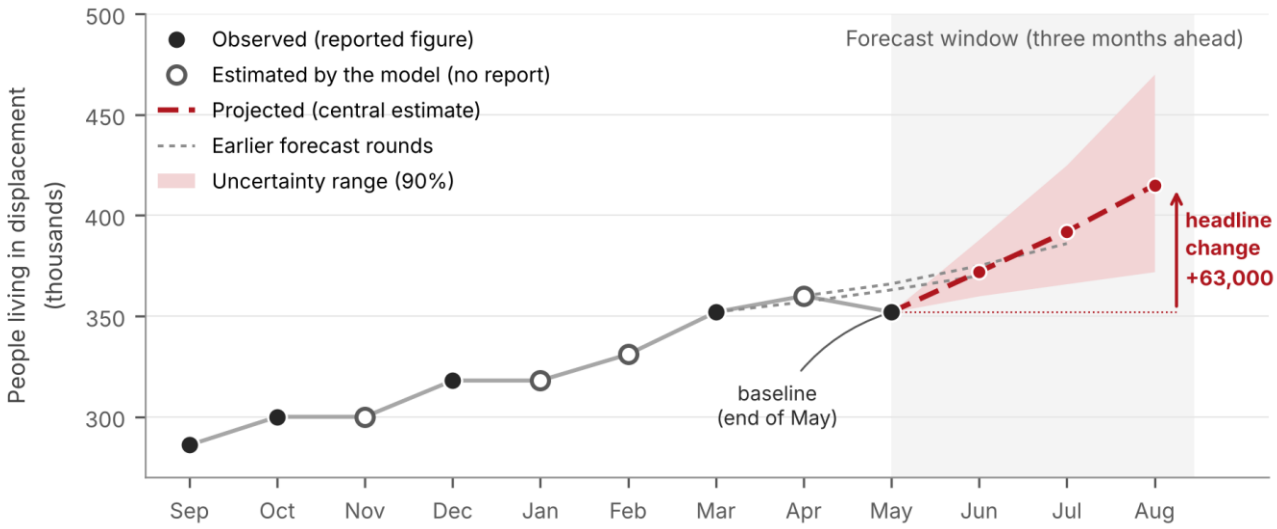


Figure 2. *A single district over time. Illustrative data. Solid points are observed (reported) figures; hollow points are estimated by the model where data were missing; the dashed line and shaded band are the projection and its uncertainty range. The headline change runs from the May baseline to the August projection. Dotted grey lines show the projections from the two previous forecast rounds, illustrating how each round updates the last.*

What the forecast numbers mean

Most AHEAD model forecasts estimate the total number of internally displaced people expected to be present in an area at the end of each forecast month. In some countries the model instead, or additionally, forecasts cross-border movements. Depending on the context, this may be the number of refugees present in a location at a given time, or the cumulative number of arrivals since a given date. What the forecast represents therefore depends on the country and the output in question, and it is worth being clear which of these a given forecast refers to before interpreting it.

The rest of this section uses IDPs as the running example, since most models forecast internal displacement. In this note we call that projected figure the total number of IDPs, or the number of people living in displacement. It is not the number of people newly displaced during the month, which is one of the most important things to get right.

A forecast value for August does not mean "the number of people who will be displaced during August." It means the estimated total number of IDPs expected to be present in that area by the end of August. That total reflects many different movements at once (e.g., new displacement into the area, secondary displacement from elsewhere, returns, onward movement, people returning home). If an area is projected to have 100,000 IDPs by the end of August, this does not mean 100,000 people will arrive in August. After arrivals, returns and other movements are accounted for, around 100,000 people are expected to be living in displacement there by then.

The same logic applies to the three-month change figures in the snapshots. When a snapshot says the number of IDPs is projected to rise by 10,000 between May and August, it means the expected total at the end of August is 10,000 higher than the total at the end of May. It does not mean exactly 10,000 movements will happen over June, July and August. The real number of movements can be far higher, especially where new displacement, secondary displacement and returns are all happening together.

This is why a modest rise in the total can sit on top of a great deal of movement. An area may take in many newly displaced people while others return or move on, and in the final figures these partly cancel out. The result can be a relatively small net increase even when the underlying situation is highly unstable. The forecasts are therefore best understood as estimates of how the overall IDP caseload is expected to change in each destination area.

Net change remains a useful measure. A large projected net increase indicates that displacement pressure in an area may be building, whatever mix of movements lies behind it. In many contexts, the total number of people living in displacement is also the most consistently available displacement measure, which allows forecasts to be produced and evaluated in a similar way across countries.

Two further cautions follow from this. First, the forecast describes the destination, not the origin. It does not, by itself, say where displaced people came from. If a territory is projected to see an increase in the number of IDPs, we interpret it as a signal that more people may be present there and may need support. The displacement itself may have been triggered elsewhere. The drivers may lie in that area, in neighbouring ones or along the routes between them.

Second, the forecast does not always identify the type of movement behind a change. As mentioned, a projected increase is a forecast of the net change in the total number of IDPs in a destination area, not of new displacement events, new arrivals or secondary movements specifically. That net change is often largely by new displacement. It may also be shaped by onward movement, repeated displacement or shifts in reporting, and the forecast does not always separate these or show how much each contributes. So a projected rise should not be taken as a direct count of "newly displaced households."

Forecasts are probabilistic

The forecasts do not give one certain future. Instead, the model tests many possible futures based on the data available at the time. Some of these possible futures show a smaller increase in displacement, while others show a larger increase. The main number shown in the forecast is the central estimate. It is the model's best summary of these possible futures and is best understood as the most likely direction and scale of change. But it is not an exact prediction, and the actual number may be higher or lower.

This is why AHEAD also reports an uncertainty range. The range shows where the model expects the final number to fall in most cases. A narrow range means the possible futures are relatively close to each other, so the model is giving a more precise estimate. A wide range means the possible futures differ more, so there is greater uncertainty. In current AHEAD outputs, DRC usually reports a 90 percent uncertainty range. In simple terms, this means the model expects the final value to fall within this range in most cases, based on the data available at the time. Even when the model is working well, the actual number is expected to fall outside a 90 percent range in roughly one case in ten, and DRC checks this as part of evaluation. Outcomes outside the range become more likely if a major new shock occurs or if important data are missing.

For users, the important point is to look at both the central estimate and the uncertainty range. The central estimate gives the main forecast signal. The uncertainty range shows how cautious users should be when interpreting that signal. How well the forecasts have performed in the past is covered in the section on evaluating model performance.

There is no fixed cut-off below which forecasts are not produced. In some contexts, however, gaps in publicly available displacement data are so large that the uncertainty ranges become very wide and the central estimate alone is not a reliable guide to magnitude. In these cases the forecast should be interpreted with particular caution, putting less weight on the size of the projected figures and more on the expected direction of change, considered together with conflict trends and other contextual information.

Interpreting forecast rounds over time

Because forecasts update every month, users might want to compare rounds, which is useful but also needs care. Each round has its own baseline and its own three-month window, and the rounds overlap rather than adding up. If the June forecast round projects an increase of 2,000 IDPs between May and August, and the July forecast round projects 1,500 between June and September, these do not combine to 3,500. The July forecast is a fresh estimate from a new starting point with new information; it updates the previous view rather than extending it.

The safest comparison is to track the same area and the same target month across rounds. A district projected to rise again and again might point to persistent pressure. A district that appears as a hotspot one month and not the next is worth checking to understand whether it is a genuine change in the context, new data, a revised baseline or simply model uncertainty.

At the same time, end users should not wait for perfect consistency before making specific decisions. Conflict-induced displacement can change quickly and forecasts exist to support earlier response, not to deliver complete certainty. A projected increase is a signal that calls for further analysis, always considered alongside other evidence. The AHEAD model can show where displacement pressure may rise, though it cannot explain everything on its own.

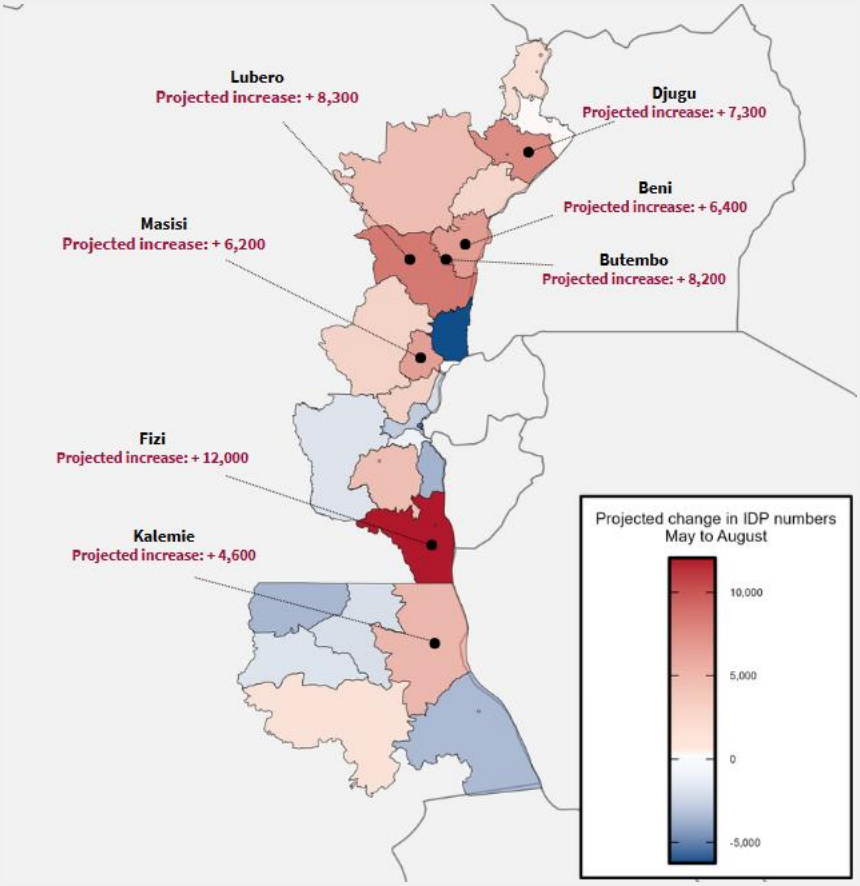
In practice: eastern DR Congo snapshot, June 2026

The June 2026 snapshot projected displacement for June, July and August across 28 territories in North Kivu, South Kivu, Ituri and Tanganyika. Among other things, the snapshot notes that *"Displacement is expected to rise across several territories in South Kivu province. Fizi is projected to record the largest increase, with around 12,000 additional IDPs by August 2026".* This means that the model projects an increase of around 12,000 in the total number of IDPs recorded in Fizi. It does not mean that 12,000 people are expected to arrive in Fizi during August, or that 12,000 new displacement movements will take place.

The latest available DTM assessment at the time, cycle 16, reported 307,368 IDPs in Fizi in February 2026. Because there were no updated assessments for March, April and May, the AHEAD model first estimates the missing monthly values up to the present. It does this by using Fizi's own displacement history, displacement patterns in neighbouring areas, conflict dynamics and other relevant drivers.

Based on this, the model estimates that there were 319,193 IDPs in Fizi in May 2026. This May estimate becomes the baseline for calculating the projected three-month change. The model then projects 323,897 IDPs in Fizi for June, 327,494 for July and 331,224 for August. The projected increase reported in the snapshot is calculated by comparing the August projection with the May baseline.

This is why the snapshot reports an increase of around 12,000 IDPs by August 2026. The figure is a projected net increase in the total number of IDPs hosted in Fizi between May and August. It may reflect different underlying movements, including arrivals, onward movement, returns or changes in displacement patterns.



Thresholds based on the AHEAD predictive model

Forecasts become operationally useful when they are tied to clear decisions. On its own, a forecast shows where displacement is expected to rise. Combined with a pre-agreed threshold, it also helps judge when an expected rise is large enough to warrant a response. That is the purpose of thresholds and triggers.

In anticipatory action, thresholds are normally agreed in advance. A threshold is the level of risk or expected impact that must be reached before action is considered. In AHEAD, this is the projected three-month increase in the number of IDPs that is considered unusually large compared with recent historical patterns.

The threshold has to be interpreted together with everything in the previous section. The model is not forecasting individual movements or new arrivals; it is forecasting the expected total number of IDPs in a destination area at the end of each month. The threshold is therefore applied to the projected net change in that total over three months. If the number of IDPs in a district is expected to rise by more than the threshold between the baseline and the end of the forecast window, this is treated as a signal of an unusually large increase in displacement pressure.

Considerations for setting a threshold

There is no single correct threshold. Where it is set is a practical choice, and it directly shapes how often the system will be triggered. Setting a threshold well means being clear about what makes a threshold 'good', how to choose the level and what geographic scope to set it across. Each is discussed below.

What makes a good threshold. Setting a threshold always involves a trade-off. Set it low and the system flags more potential crises, but it also raises more false alarms and may exhaust the capacity and credibility needed to act. Set it high and it produces fewer false alarms, but it risks missing real events or leaving too little time to prepare. It's not easy to remove this trade-off; the aim is to strike a balance that fits the context. A good threshold is therefore one that is credible to the people who must act on it, transparent enough to be explained and defended, operationally useful given the response that can realistically follow and suited to the quality and quantity of the available data. Where the cost of missing an event is very high, a more sensitive threshold may be justified; where response resources are scarce or a false activation is costly, a more conservative one may be preferable. Put simply, the cost of acting when displacement does not follow has to be weighed against the cost of not acting when it does. Because response resources are generally scarce, limiting false alarms often carries considerable weight. At the same time, where a forecast signal is followed by further checks before full activation, such as community indicators, a somewhat more sensitive forecast threshold can be acceptable, because those later steps help filter out false alarms.

How to choose the level. Several approaches can be used to set the level of the threshold. They are not mutually exclusive, and in practice more than one can be used together:

- A percentile of recent history. The threshold is set at a high percentile of the displacement changes observed in the recent past, so that it flags changes that are unusually large relative to what a country has actually experienced. This is grounded in real data, adapts automatically to each context and is straightforward to explain and apply consistently across areas. Its main dependency is data, as it needs enough good-quality history to be stable, and by definition it assumes the recent past is a reasonable guide to what counts as unusual, which may not hold when dynamics change.
- An absolute figure. The threshold is a fixed number, for example a set increase in displaced people, chosen because it is operationally or programmatically meaningful, such as the caseload at which a response plan is designed to activate. This is simple, stable and transparent, and easy to communicate to partners and donors. But the same absolute number can mean very different things in a small area and a large one, and it may not reflect what is actually unusual or severe for a given context.
- Operational capacity. The threshold is set with reference to the level of change the response can realistically prepare for or act on in a given area, so that triggers stay aligned with what teams can actually deliver. This keeps triggers actionable and avoids raising alarms that cannot be resourced. Where information is available, it can

take into account humanitarian response capacity as well as absorptive capacity and the services already in place, for example those provided by government or local actors. Its weakness is that it risks setting the bar according to what is convenient rather than according to risk.

Whichever of the above provides the starting point, the candidate threshold should be tested against historical data to see how it would have performed, how many real events it would have caught and how many false alarms it would have raised, and then adjusted. This is the link to the evaluation section below.

What geographic scope to set it across. A second choice, linked to the choice of threshold level, is the geographic level at which the threshold is defined, whether one shared threshold for the whole country, one per region or one per district. Each has advantages and drawbacks, and the right choice depends heavily on how much data is available. The two choices interact, since the same percentile can produce a very different threshold when it is calculated across a whole country than when it is calculated for a single district.

A pooled threshold for the whole country is rather stable as it draws on all areas' data at once. It's easy to explain and to apply consistently and works even where individual districts have sparse or irregular data. Furthermore, it would direct teams to focus response in the areas that see the highest levels of displacement. However, it treats a large increase as the same everywhere, so it may be too high for areas that are usually calm and too low for areas that are usually volatile.

Regional thresholds can also be useful as they reflect genuine differences between regions while still pooling enough data within each region to stay reasonably stable. However, this needs enough data within each region to be reliable, and there is more to maintain and explain than with a single country threshold.

District-specific thresholds reflect each area's own displacement pattern most closely, which can be appealing in principle. But in practice they are often unstable. A single district's displacement history might have few data points, so one past spike can push the threshold so high that future risks are missed, while a district with little recorded displacement can have a threshold so low that minor changes trigger alerts constantly.

More elaborate approaches can blend these, letting a threshold reflect local patterns while borrowing strength from the wider country trend, which reduces the influence of any single spike or data gap. These can perform well, but they are harder to explain and less transparent for operational use, which is itself a cost. As with the choice of level, a key guiding principle is that the best scope is the one that is stable given the available data, credible to those who act on it and no more complex than it needs to be.

Forecast thresholds: the percentile approach

DRC's standard approach uses the percentile method mentioned above, defined at the pooled country level. It defines an unusually large increase from the displacement changes actually observed in each country over the recent past.

One way to approach the calculation of the threshold in this case is to look at the previous two years of data. For each district, we calculate how much the total number of IDPs changed over every possible three-month period. For example, we calculate the change from January to April, then from February to May, then from March to June, and so on. These are rolling three-month changes as the calculation moves forward one month at a time.

This gives a long list of historical three-month changes across all districts. The list includes periods where IDP numbers decreased, periods with little change, and periods where a district saw a sharp increase. Together, these values show what three-month displacement changes have looked like in the country during the recent past.

All of these changes are then put together and ordered from the lowest to the highest. The threshold is set at the 90th percentile of this ordered list. This means the threshold is the point where around 90 percent of recent three-month changes were lower and around 10 percent were higher.

In practical terms, the 90th percentile is used to identify changes that are large compared with recent history. It means that, in the recent historical data, only about one in ten three-month changes was larger than this value. For example, imagine that after calculating all rolling three-month changes across all districts, most changes are between -1,000 and +2,000 IDPs. Some are larger, such as +5,000 or +8,000. When the list is ordered from lowest to highest, the value at the 90th percentile is +6,000. In that

case, +6,000 becomes the threshold. A district forecasted to increase by +7,000 IDPs over the next three months would thus cross the threshold as the projected increase is larger than what has usually been observed in recent three-month periods. But the threshold is a statistical signal and does not measure humanitarian severity, vulnerability or response priority. A district can stay highly vulnerable even with a projected increase below the threshold, and a district can cross the threshold yet still face real questions about whether anticipatory action is feasible.

What a 90th percentile threshold implies in practice also depends on the context. Because it is pooled across districts and applied to every district each month, how often a country can expect a threshold crossing depends on how many districts are covered and how displacement is distributed across them. The percentile is therefore a starting point that can be revised. Setting thresholds is a live process. Thresholds are reviewed as new data accumulate and experience with activations grows, and can be adjusted where they flag too many or too few increases.

When a district's forecast crosses the threshold, this is flagged to a pre-defined group of focal points, and the steps that follow are those set out in the pre-agreed anticipatory action plan for that country. The final section of this note gives examples of how the forecast signal can be combined with community indicators.

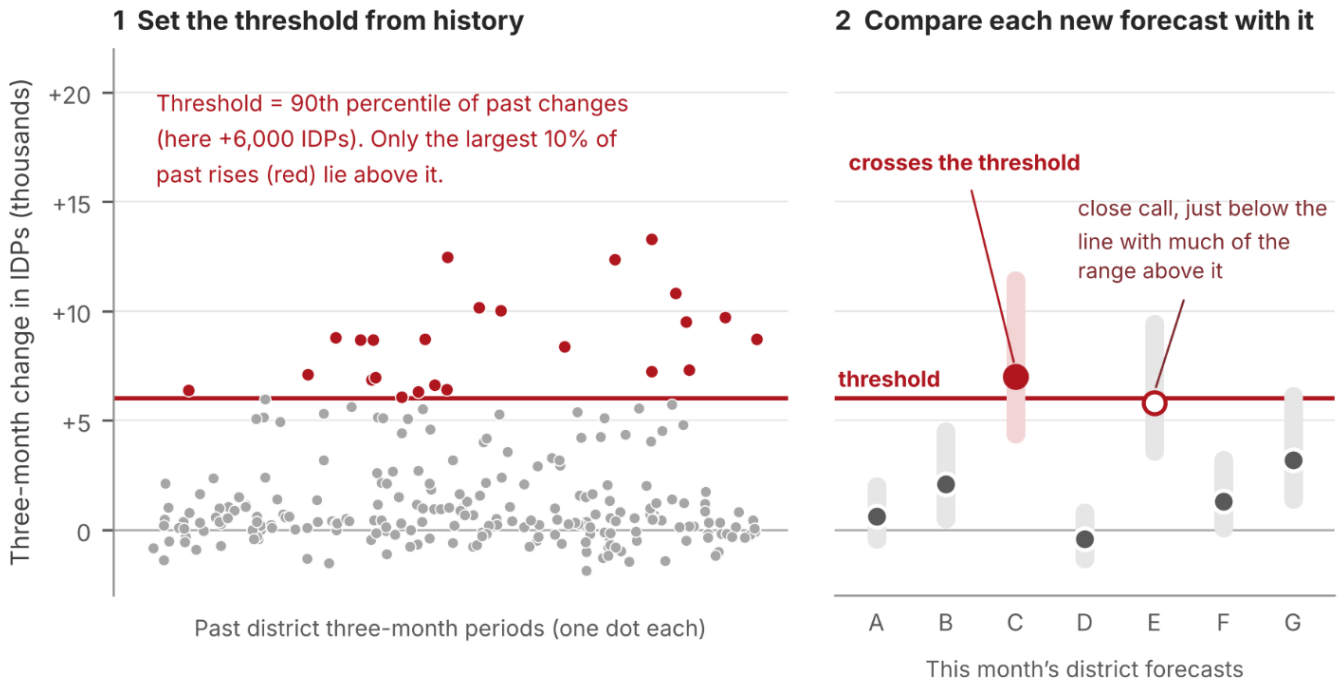


Figure 3. *Setting and applying a threshold. Illustrative data. In the left panel, each dot is one district in one three-month period over the past two years, and its height is the change in IDPs over that period. The forecast threshold is set at the 90th percentile, the level exceeded by only the largest 10% of past three-month rises (red). In the right panel, the same threshold is applied to each district's new forecast. District C crosses it, while district E falls just below it with much of its uncertainty range above it.*

Interpreting a threshold crossing

A threshold crossing means the forecast three-month increase in the IDP caseload is unusually large compared with recent patterns. It does not mean the model is certain, that the increase will happen exactly as forecast, or, importantly, that the forecast number is the largest possible increase.

That last point deserves emphasis. If a projected increase is just above the threshold, the actual increase could still be much higher if the crisis escalates, if displacement accelerates or if the uncertainty range is wide. As mentioned before, forecasts are probabilistic and the central estimate is a signal and should not be treated as a ceiling. When an area crosses the threshold, we need to look beyond the central forecast to the uncertainty range, the recent trend, the narrative analysis and the wider context. If the threshold is 6,000 and a district is projected to rise by 6,500, one should not prepare for exactly 6,500 more people. The range may point to a far larger increase,

and field teams may already be reporting worsening security or new movement restrictions. Planning should hold a span of outcomes in view, including a higher-impact one. This is also why the target population for a response should not be defined by the forecast figure alone.

The same caution applies just below the line. A projected increase of 5,800 against a threshold of 6,000 should not be dismissed automatically. If the range is wide, if the area has appeared repeatedly in recent rounds or if community indicators are worsening, it can still warrant close monitoring or preparatory steps. DRC does not currently apply a fixed rule for these close calls, for example one based on how much of the uncertainty range lies above the threshold. For now, they are best handled through discussion among the focal points and teams responsible under the anticipatory action plan, taking into account the width of the range, recent rounds and field information. More specific guidance may be added as experience grows.

Lead time deserves a careful look too. A mid-June forecast covering June, July and August has its threshold applied to the change between the end of May and the end of August. That does not mean the displacement is expected only at the end of August; it could come at any point across the three months, and some may already be under way when the forecast is released. So the practical lead time is often shorter than three months — perhaps ten or eleven weeks to the end of the window, but less before displacement begins, accelerates or peaks. How much usable lead time there is also varies by context and by the type of conflict and displacement. Sudden escalations, such as a rapid offensive, may leave very little time between the first signs and large-scale movement, while displacement driven by slower-building pressures may become visible in the forecasts earlier. Expected lead time should therefore be discussed for each context when anticipatory action plans are designed. This is also why the forecast threshold is useful, as it can give teams a reason to prepare before the full increase shows up in reported data. And pairing it with community and field indicators is what helps narrow the timing.

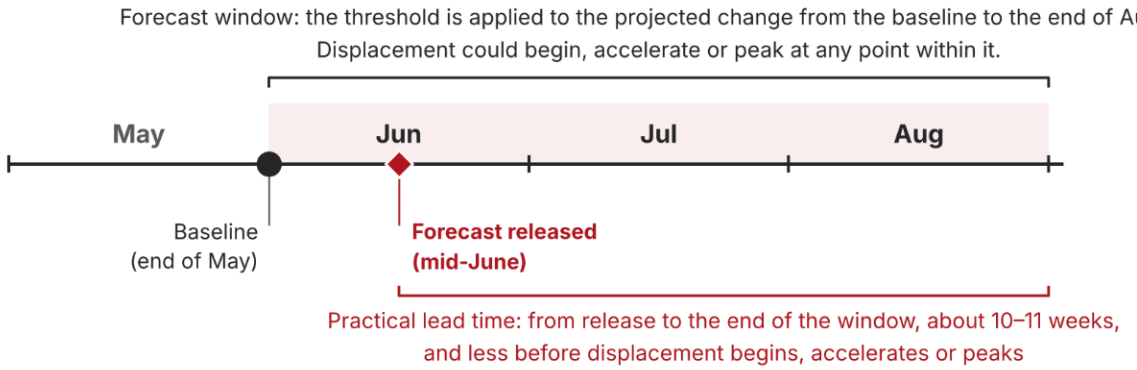


Figure 4. *Lead time from a three-month forecast. A forecast made in mid-June uses the end of May as its baseline and looks across June, July and August, with the threshold applied to the projected change over the whole window. If the projected change over the window is above the threshold, the forecast is flagged. Because displacement could begin, accelerate or peak at any point in the window (and some may already be under way when the forecast is released) the practical lead time is usually shorter than three months.*

In practice: a threshold crossing in Téra, Niger

In the May 2026 AHEAD forecast round for Niger, the model projected displacement for May, June and July, using the end of April as the baseline. For Téra, the model projected that the number of IDPs at the end of July would be more than 2,600 higher than at the end of April.

The forecast threshold for Niger was set at about 2,100 IDPs. This threshold was calculated from rolling three-month changes in department-level IDP caseloads over the previous two years. Because Téra's projected increase was above 2,100, Téra crossed the forecast threshold. This means that the model projected an unusually large increase in the number of IDPs present in Téra by the end of the forecast period, compared with recent displacement patterns in Niger.

The threshold crossing should therefore be understood as an early warning signal. It shows that the projected increase is large enough, relative to recent history, to warrant the steps set out in pre-agreed anticipatory action plans. It does not show exactly when within the three-month window movement will happen, where people will come from or whether the change is driven by first-time displacement, onward movement, secondary movement, returns or changes in reporting. Those questions require community indicators, field information and other analysis.

Evaluating model performance

Before a country model is used to support decisions, it is evaluated to understand how well it performs and where its limits lie. Evaluation is frequent. A comprehensive evaluation round is carried out at least twice a year, and in practice the models are also evaluated every time they are changed or re-fitted. Performance figures can therefore change relatively often, as new models are developed and tested to improve their value for anticipatory action; the figures below reflect the latest evaluation of each country model (a few country models are currently undergoing development and testing and their metrics will be incorporated into the guidance once available). Evaluation matters for two reasons, and they map directly onto the two things the model is asked to do. First, are the forecasts accurate enough to be relied on at all? Second, and just as important, how does the model behave at the threshold, that is, would it actually have prompted the right decisions in the past? A model can be reasonably accurate on average yet still perform poorly at the threshold, and a model can look strong at the threshold while being noisy elsewhere, so both need to be checked.

All evaluations are done out of sample, that is, on held-out history. The model is asked to forecast past periods as if they were the future, using only the information that would have been available at the time, and its forecasts are then compared with what actually happened. This is sometimes described as hindcasting.

Performance is assessed in two complementary ways, namely general statistical accuracy and accuracy from an anticipatory-action perspective. The two are described below.

Statistical accuracy: how close are the forecasts?

The first question is simply how far the displacement forecasts fall from the figures that eventually materialised. Several metrics are used, each capturing a different aspect of error.

Mean absolute percentage error (MAPE). The typical error expressed as a share of the actual level. Because it is a percentage, it can be compared across areas of very different sizes. A MAPE of 15% means the forecast is, on average, about 15% away from the actual displacement figure in a given area/country. It is the most intuitive headline measure of accuracy, but it can be volatile where actual displacement figures are very small.

Mean absolute error (MAE). The typical error expressed in the same units as the displacement figures themselves, that is, numbers of displaced people. This is useful when absolute magnitudes matter operationally, since a 10% error means something very different in an area of 5,000 than in an area of 500,000.

Continuous ranked probability score (CRPS). The two metrics above judge only the single central forecast figure. But, as mentioned earlier, the model is probabilistic, i.e., it produces a range, not just a point. The CRPS evaluates the whole forecast range against what happened, rewarding forecasts that are both accurate and appropriately confident. A precise-looking forecast whose uncertainty range was badly judged will score poorly here even if its central figure happened to be close.

These metrics help determine whether the central forecast is close and whether the uncertainty around it can be trusted, which is what allows the forecast range to be used for planning. To make sense of the numbers in absolute terms, model performance is also compared against simple benchmarks, for example a no-change forecast, which assumes displacement stays at its current level, and a recent-trend forecast, which assumes it keeps changing at its recent rate. These are the kinds of implicit guesses one might make without a model, so a useful model should outperform them. The table below records these metrics for each country model.

	Resolution	MAPE	MAE	CRPS
Somalia	Admin 2	12	1,284	-
Uganda	Settlement	0.65	717	1127
Sudan	Admin 1	10.77	35,592	62,433
Niger	Admin 2	9.95	959	1,080
Yemen	Admin 2	N/A	18.84	46.44
Myanmar	Admin 3	17.98	2,274	2,325

Anticipatory-action accuracy: performance at the threshold

Statistical accuracy is necessary but not sufficient. For anticipatory action, what matters is also whether the model, used with the threshold set for it, would have led to the right decision, that is, raising the alarm when a large increase was coming and staying quiet when it was not. This ties evaluation directly back to the threshold discussed earlier. The same threshold that will be used operationally is applied to the model's past forecasts, so that its performance is judged on the exact decision it will be used to support.

The method is to replay history at the threshold. For every past area and period, we ask two yes/no questions. Did the model's forecast cross the threshold, and did the change that actually happened cross it? Comparing the two places every case into one of four outcomes, as shown in the example below where we assume a threshold of 5,000 individuals.

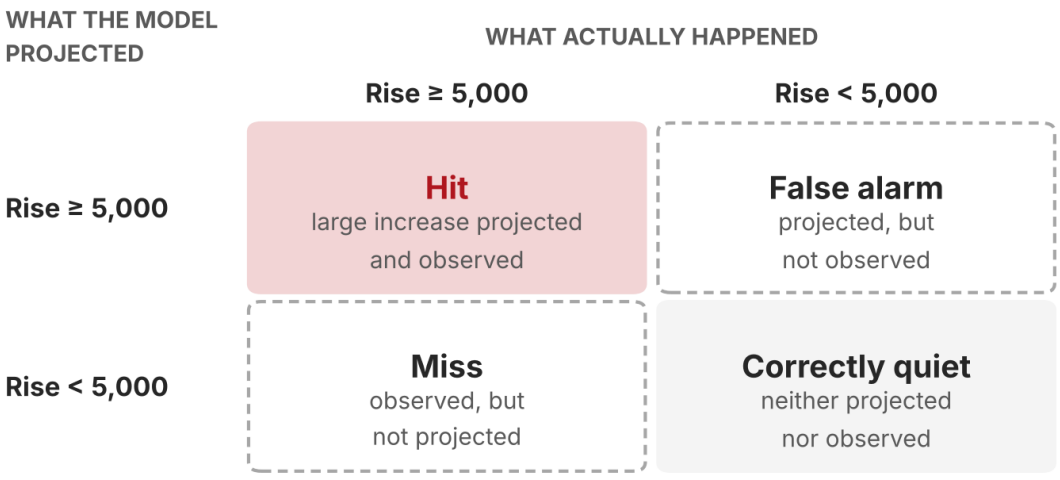


Figure 5. *The four possible outcomes when past forecasts are compared with what actually happened, using an example threshold of 5,000 IDPs.*

From these four counts we derive three metrics that together describe how the model behaves at that threshold. They are the operational counterpart of the statistical metrics above. Where MAPE and the rest ask how close the forecast was, these ask whether it would have led to the right call.

- **Hit rate.** Of all the large increases that actually happened, how many did the model flag in advance? A hit rate of 70% means the model caught roughly seven of every ten large increases. The miss rate, the share of large increases the model did not flag, is simply the complement, so a hit rate of 70% means a miss rate of 30%.
- **False-alarm rate.** Of all the times the model crossed the threshold, how many were not followed by a large increase?
- **Overall accuracy.** How often the model made the right call either way, correctly flagging large increases and correctly staying quiet otherwise. It should always be interpreted alongside the other two, since a model can score high on overall accuracy simply by staying quiet in a context where large increases are rare, while missing most of the few that occur. For this reason, it should not be used as a headline measure on its own.

Crucially, these three metrics move together as the threshold changes, which is what connects evaluation back to the choice of threshold. Lowering the threshold usually raises the hit rate but also the false-alarm rate; raising it does the reverse. This is precisely why the threshold is set against evidence. The evaluation shows, for a candidate threshold, how many real events would have been caught and how many false alarms raised, and the threshold can then be tuned to strike the balance the context calls for. It is also why the same country may revisit its threshold as more data accumulate and the picture sharpens. The table below records these metrics for each country model at its chosen threshold.

	Threshold used	Hit rate	False alarm rate	Overall accuracy
Somalia	2,400	27%	68%	89%
Uganda	3,500	48%	9%	94%
Sudan	73,000	33%	51%	90%
South Sudan	2,400	19%	90%	92%
Niger	2,400	36%	19%	93%
Yemen	1,000	8%	80%	99%
Myanmar	3,500	16%	44%	90%

Taken together, the two sets of metrics give a rounded picture. The statistical metrics show whether the forecasts are close and whether their uncertainty can be trusted; the anticipatory-action metrics show whether the model would have supported the right decisions at the threshold. A model that performs well on both can be used with more confidence as one input into decisions. Where a model is strong statistically but weak at the threshold, or vice versa, that tension is itself useful information, and usually points either to a threshold that needs adjusting or to a context where the model should be weighted more cautiously against other evidence.

Where forecasts call for more careful interpretation

Two limitations should be kept in view when interpreting any of these metrics. First, evaluation rests on historical data, so it measures how the model would have performed under conditions resembling the past. A model can be well-evaluated and still be caught out by an unprecedented shock or a sudden change in dynamics, a challenge shared by all displacement forecasting, where the largest and most sudden crises are often the hardest to anticipate.

Second, the figures are only as good as the displacement data behind them. Where reporting is sparse or irregular, performance is estimated on fewer real observations and should be interpreted with more caution.

Like any forecasting model, AHEAD performs best when the future resembles the patterns it has learned from. In some situations its forecasts call for more careful interpretation, including the following.

- Sudden and large increases. The larger the relative increase in displacement, the larger the forecast error tends to be. The model may pick up the direction of a sharp escalation while underestimating its scale, and the uncertainty range and community indicators help to put the forecast in perspective.
- Areas that have been quiet for a long time. Where displacement has been very low for an extended period, there is little recent movement to learn from. Because the models pool information across areas, they can still draw on patterns observed elsewhere in the country.
- Rapid changes in conflict conditions. Sudden political or military developments can change conditions faster than monthly data can capture. Monthly updates allow the forecasts to take account of new data, although they may still lag behind fast-moving events.

Performance also differs between countries, largely for the same reasons. Forecasts tend to be strongest where displacement data are reported regularly and consistently, where displacement responds in a fairly stable way to conditions such as conflict and food insecurity, and where the history covers a range of past escalations. They call for more caution where reporting is sparse, where a large share of recent values are estimates, or where displacement is dominated by sudden, large-scale events.

Combining the AHEAD model with community indicators

The chapters above have focused on the AHEAD model, what it forecasts, how to interpret it, how thresholds turn a forecast into a signal and how its performance is evaluated. This final section places the model within DRC's wider approach to anticipatory action. The central point is that DRC does not base anticipatory action on the model alone. The forecast is one layer of early warning, used together with a second layer built on community-level indicators.

For anticipatory action to be locally owned and data-informed, the mechanisms need to be built on community-level indicators together with the forecasting model. Community-level indicators can be signals in the community that conflict and displacement are likely to happen in the near future. They are identified with communities and local actors, based on what people have observed before previous episodes of conflict and displacement. They help ground the early warning mechanism in the realities and lived experience of the communities most exposed to risk.

This section explains, at a high level, what community indicators are, how they are developed and used, and how they combine with the model. The detailed methodology, how indicators are identified, how data are collected and monitored, and how community thresholds are set and reviewed, is set out in DRC's internal guidance on community-based anticipatory action indicators. This section gives an overview of how this second layer fits

alongside the model in DRC's overall anticipatory action approach.

Community indicators are identified through direct engagement with communities and local actors, for example through key informant interviews and workshops. The starting point is to discuss past episodes of conflict and displacement, and to draw out the signs that people saw beforehand. These discussions typically focus on which signs appear in the lead-up to conflict and displacement, how much time usually passes between those signs appearing and displacement reaching its peak, and whether there are tipping points, specific levels or changes in a sign after which significant displacement tends to follow.

The signs themselves vary by context. Some are drivers of displacement in their own right, such as killings, abductions or cattle raiding. Others are indirect signals that are linked to rising risk without directly causing displacement, such as movement restrictions, rising market prices or feelings of insecurity. Not every sign becomes an indicator. Selection weighs how meaningful a sign is for anticipating displacement, whether data on it can be collected credibly and safely, and whether relevant information is already being collected through existing systems such as protection monitoring.

Once indicators are selected, they are monitored on a regular basis, most commonly through structured key informant interviews, and sometimes direct observation, so that changes can be compared consistently over time. Monitoring is preceded and accompanied by a risk assessment, since collecting information on conflict-related signs can be sensitive for communities, key informants and staff. The frequency of data collection is adjusted to the situation - it intensifies when risk appears to be rising, including when the model has flagged an area, and can be scaled back in periods known to be calm.

Each indicator is given a community threshold, the level at which it is considered to be signalling imminent risk. Initial thresholds are often based on what communities identify as tipping points, and are then checked and refined against data as it accumulates, by looking at how indicator levels related to conflict and displacement that actually followed. As with the predictive model threshold, this involves a trade-off between catching real events and raising false alarms, and the aim is a threshold that is credible, timely and useful for decision-making, accepting that some uncertainty will remain. DRC also regularly assesses which community indicators give reliable early signals of increases in conflict and displacement, for example within two weeks of data collection.

How the two layers combine

The model and the community indicators can be combined in different ways, and there is no single formula. How they are brought together depends on the context, the indicators available and the response the plan is designed to support.

In some configurations, the two are used at different stages of decision-making. A model signal might set readiness in motion, that is, the preparatory steps that come before full anticipatory action, such as alerting focal points, reviewing the plan, preparing resources and intensifying community-indicator monitoring in a flagged area, while fuller anticipatory action activation depends on community indicators also pointing to rising risk. Sequencing the two this way can help balance early awareness against the risk of acting on a single signal. In other configurations, a strong model signal alone may justify a degree of early action; or community indicators may carry more weight in the decision to act, particularly where they have historically given reliable warning in that area. The balance between the two, the order in which they are considered and the specific actions attached to each all can vary from place to place.

In every case, how the two layers combine, what each contributes to a decision and which actions follow are defined in the pre-agreed anticipatory action plans for that country or area. Agreeing this beforehand means that when a signal does arrive, the response follows a framework that has already been thought through and validated.

The mechanism developed in Akobo, South Sudan, illustrates how the two layers can work together. There, community protection indicators identified with local peace-committee members, such as killings and abductions outside the town, movement restrictions and pre-emptive movement into town, were monitored through key informant interviews and combined with the displacement model. The diagram below shows how the two triggers interact.

The two triggers are kept distinct in this example. The model has its own trigger and the community indicators have theirs. It is the combination that determines how much is activated. Either trigger on its own leads to a

partial activation, a proportionate set of early actions, while both triggering together leads to complete activation and the full package. This allows the mechanism to act on partial evidence without committing everything and to escalate as the picture becomes clearer.

The specific indicators, thresholds and actions in Akobo are particular to that context and were agreed with local structures as part of its anticipatory action plan. Other countries and areas will have different indicators, different threshold values and different configurations of readiness and activation. The constant is the underlying approach, a forward-looking model signal and a locally grounded community signal, combined through a pre-agreed plan, so that anticipatory action is neither driven by the model in isolation nor left waiting until a crisis is already fully visible.

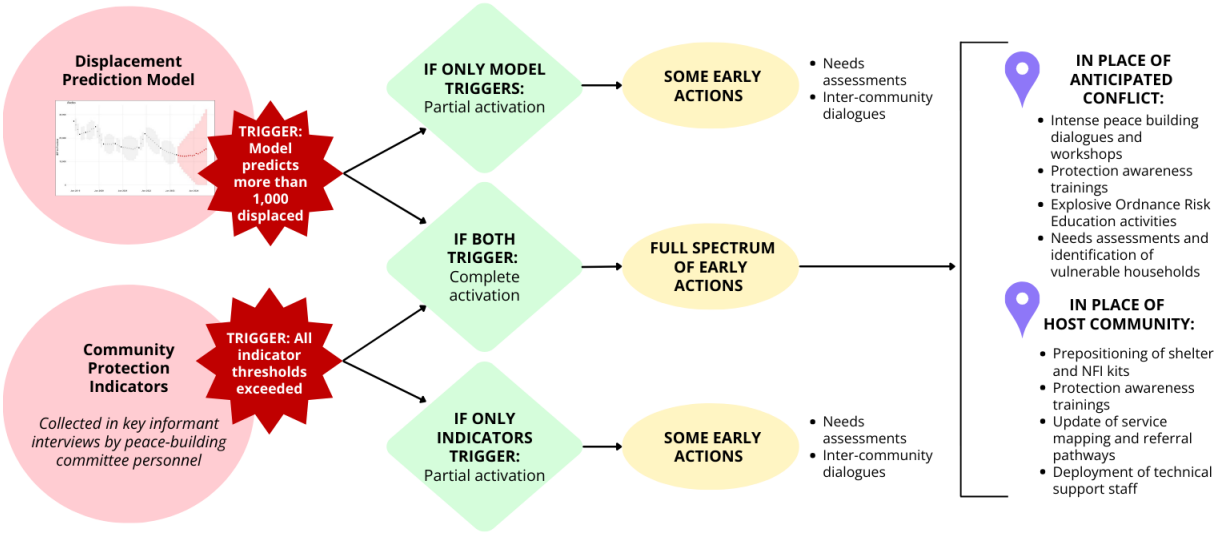


Figure 6. *The two layers of DRC's early warning system (Akobo, South Sudan example)*

Technical annex

This annex sets out how the AHEAD displacement-forecasting models are built. It describes the modelling framework, gives the core equations of a representative model and outlines other modelling approaches that DRC is testing.

Many models

The models differ from country to country. Each is fitted to that country's own displacement series, administrative geography, data availability and conflict dynamics, and the structural choices (which covariates enter, whether flow data are used, how the count of displaced people is modelled) vary accordingly. What follows is therefore an example of the shared framework rather than a single specification.

The modelling framework

State-space formulation. Most country models are Bayesian state-space models for the number of internally displaced people present in an area. Let $D_{s,t}$ denote the displacement level in district *s* at month *t*. The latent level evolves multiplicatively from one month to the next, driven by a set of covariates and a stochastic term:

$$D_{s,t} = D_{s,t-1} \cdot \exp(\beta^T x_{s,t} + \alpha_s + \varepsilon_{s,t})$$

Here $x_{s,t}$ is a vector of covariates (conflict intensity, food insecurity, climatic stress and others, measured in the district and in its neighbours), β the associated coefficients, α_s a district-specific intercept (where relevant). Equivalently, the conditional mean of the level is log-linear in the previous level and the covariates:

$$\log \mu_{s,t} = \alpha_s + \beta^T x_{s,t} + \log D_{s,t-1}$$

Because a lot of internal displacement is local or between adjacent areas, covariates are constructed from both the district itself and its neighbours, allowing conditions next door to inform a district's near-term trajectory. The coefficients are shared across districts and time, so the historical relationship between drivers and displacement can be learned from the full panel and projected forward.

Gamma observation model. Displacement levels are strictly positive and strongly right-skewed, with modest month-to-month changes in most cases and occasional very large surges. The observation model is therefore a gamma distribution, parameterised so that its mean follows the latent state while its dispersion varies by district (and over time/predictor values in some cases):

$$D_{s,t} \sim \text{Gamma}\left(\frac{\mu_{s,t}}{\phi_s}, \frac{1}{\phi_s}\right), \qquad \text{E}[D_{s,t}] = \mu_{s,t}, \quad \text{Var}[D_{s,t}] = \mu_{s,t}\phi_s$$

The mean is $\mu_{s,t}$ and the variance scales with it through a district-specific dispersion ϕ_s. This delivers the asymmetric predictive intervals seen throughout the forecasts, which reflects how displacement behaves. The dispersions are themselves drawn from a shared hierarchical prior, so information is pooled across districts.

Covariate dynamics and lags. Drivers do not act instantaneously. A spike in violence may translate into displacement over the following weeks. Several country models therefore include distributed lags, in which each covariate enters at the current month and, down-weighted, at the previous month. The weights control how much of last month's covariate still matters and are themselves estimated. This lets the model capture momentum without a fully specified autoregressive lag structure. Lags were tested in all country models and kept where they improved forecast performance, with the testing guided by what is known about displacement dynamics and registration practices in each context.

Gap-filling: missing observations as parameters. Sub-national displacement series are irregular, with frequent gaps and figures that are sometimes carried forward administratively. Rather than dropping or naively interpolating missing months, the models treat each missing value as a latent parameter to be inferred jointly with everything else:

$$D_{s,t} = y_{s,t} \text{ if observed}, \qquad D_{s,t} = y^{\text{mis}}_{s,t} \text{ (a parameter) if missing}$$

Because the latent level exists in every month whether or not it was measured, the sampler reconstructs the most plausible path of displacement consistent with the observed points and the learned dynamics. Estimated months naturally carry more uncertainty than reported ones, and that uncertainty propagates into the forecasts.

Integrating flow data where available. In some countries partner reporting provides direct estimates of movements (or new displacement events) alongside the caseload series. Where such flow observations exist for a given month, they are incorporated as a second likelihood term

that links the modelled change in level to the reported flow:

$$e_{s,t} \sim N(D_{s,t} - D_{s,t-1}, \sigma^2) \quad \text{when a flow observation is available}$$

This term is only active when a flow observation is present, so the model uses richer information where it exists and falls back on the stock dynamics where it does not. The combination of a stock likelihood and an optional flow likelihood is one of the main ways the framework adapts to differing data landscapes across countries.

Model variants

The framework described above supports markedly different specifications. Three illustrative variants:

Variant A — flow-augmented stock model. A gamma stock model with distributed lags on the covariates and an additional Gaussian likelihood linking modelled changes to observed displacement flows. Missing stock values are inferred as parameters. Suited to countries with partial flow reporting.

Variant B — hierarchical stock model with staggered starts. A gamma stock model with per-series intercepts drawn from a hierarchical prior and an explicit start index per district, so series that begin at different times are handled cleanly. Suited to countries where coverage was rolled out gradually.

Variant C — hurdle count model. For settings where the outcome is a count of events or arrivals with many zeros, a hurdle model separates whether anything happened from how much. Occurrence is modelled with a logistic component and intensity with a zero-truncated negative-binomial component. The occurrence component carries its own autoregressive terms (the log of last month's count and an indicator for a zero last month) so persistence and intermittency are both captured.

The three variants share the same Bayesian, hierarchical, gap-aware philosophy while differing in the observation model and dynamics, which is exactly how the framework adapts to each country.

Other modelling approaches being tested

Alongside the models described in this note, DRC is testing other modelling approaches in some contexts. Some of these may later be brought into operational use once they have been evaluated against the existing models. They are described briefly here, and this section and the guidance in general will be updated as the work develops.

Event classifiers. Instead of forecasting the number of people living in displacement, classifiers model a binary outcome, such as whether new displacement in an area will exceed 1,000, 5,000 or 10,000 people within the next one to three months, and estimate its probability. In some contexts, event data such as rapid displacement alerts and flow figures are more plentiful than periodic stock assessments. Approaches under test include hierarchical Bayesian logistic regression and gradient-boosted decision trees, which can capture non-linear relationships and interactions between drivers. Some specifications use a two-part structure that models whether an event occurs separately from its size. Classifiers are being tested for both conflict- and climate-induced displacement.

Combining statistical and machine-learning models. Different models capture different aspects of displacement dynamics. Ensemble methods such as stacking combine the predictive distributions of several models, with weights chosen according to their past out-of-sample performance, for example using the continuous ranked probability score. Hybrid designs are also being tested, in which the event intensity predicted by a machine-learning model enters the statistical model as an additional driver, so that the machine-learning component informs timing while the statistical model anchors the level.

Origins, destinations and returns. In some contexts, displacement data record where displaced groups came from as well as where they arrived, and some also record returns. DRC is exploring ways to use this information, including testing models of return flows and origin–destination count models of the gravity type, which relate movements to travel time and to the characteristics of origin and destination areas. Representing displacement, onward movement and returns jointly within one state-space model is technically demanding, because each process is observed incompletely and at different times, and this work remains exploratory.

Additional data sources. DRC is testing whether additional data sources improve forecast performance in some contexts, in particular earth observation data and information extracted from text. Each source is assessed by how much it improves out-of-sample performance and is retained only where it adds value.



Founded in 1956, the Danish Refugee Council (DRC) has worked for 70 years to support displaced people in their pursuit of a dignified life. As Denmark's largest humanitarian NGO and a leading international organisation, DRC works in nearly 30 countries with a dedicated team of around 6,000 employees and thousands of volunteers.

DRC advocates for the rights of and solutions for displacement-affected communities, and provides assistance during all stages of displacement: In acute crisis, in exile, when settling and integrating in a new place, or upon return. DRC supports displaced persons in becoming self-reliant and included into hosting societies. DRC works with civil society and responsible authorities to promote protection of rights and inclusion.

Our more than 5000 volunteers in Denmark make an invaluable difference in integration activities throughout the country.

DRC's code of conduct sits at the core of our organizational mission, and DRC aims at the highest ethical and professional standards. DRC has been certified as meeting the highest quality standards according to the Core Humanitarian Standard on Quality and Accountability.

HRH Queen Mary is DRC's patron.

To read more about what we do, see: www.drc.ngo

DRC DANISH REFUGEE COUNCIL

· We are there