

DECEMBER 2020



# 2020 Hindsight?

## The Ecosystem of Humanitarian Diagnostics and Its Application to Anticipatory Action

A FEINSTEIN INTERNATIONAL CENTER PUBLICATION

Erin Lentz, Gregory Gottlieb, Cori Simmons, Daniel Maxwell

Cover illustration: word cloud generated with words from the paper

Citation: Erin Lentz, Gregory Gottlieb, Cori Simmons, Daniel Maxwell. *2020 Hindsight? The Ecosystem of Humanitarian Diagnostics and Its Application to Anticipatory Action*. Boston: Feinstein International Center, Tufts University, 2020.

Corresponding author: Daniel Maxwell

Corresponding author email: [Daniel.Maxwell@Tufts.edu](mailto:Daniel.Maxwell@Tufts.edu)

Copyright 2020 Tufts University, all rights reserved.  
"Tufts University" is a registered trademark and may not be reproduced apart from its inclusion in this work without permission from its owner.

**Feinstein International Center, Friedman School of  
Nutrition Science and Policy**

**Tufts University**

75 Kneeland Street, 8th Floor

Boston, MA 02111 USA

Tel: +1 617.627.3423

Twitter: [@FeinsteinIntCen](https://twitter.com/FeinsteinIntCen)

[fic.tufts.edu](http://fic.tufts.edu)

# Contents

<b>Background and problem statement</b>	<b>4</b>
<b>Purpose of the paper</b>	<b>6</b>
<b>Diagnostic evidence, evaluative evidence, and information systems</b>	<b>7</b>
<b>What has changed about humanitarian diagnostics since 2003?</b>	<b>9</b>
<b>The confusion about diagnostics</b>	<b>12</b>
Components of humanitarian diagnostics	12
Links to action	15
<b>Discussion (and an alternative recommendation)</b>	<b>17</b>
Advances in predictive analytics and machine learning	17
Advances in theory	18
The politics of information and analysis	19
An alternative approach?	19
<b>References</b>	<b>21</b>

# Background and problem statement

In 2019 and 2020 the prevailing situation in eastern Africa resembled many recent situations where early warning systems pointed toward humanitarian crises. On top of long-standing conflict crises in at least five countries and political turmoil in several more, the long rains (*gu* rains in Somalia, etc.) in the first part of 2019 were delayed and, in many cases, well below average despite having been predicted to be about normal. However, by the time the effects of delayed rainfall were beginning to be felt, meteorological early warning systems (EWS) were forecasting a positive Indian Ocean dipole, indicating much heavier than average rainfall in the second half of 2019. Thus, while current information was suggesting that drought could impact livelihoods, early warning was forecasting flooding. At the same time (to anyone who was listening), early warning information identified an upsurge of desert locusts in the Arabian Peninsula that was likely to affect East Africa. All of these came to pass, as did the COVID-19 pandemic, which no humanitarian early warning or information system foresaw. None of this information was *wrong*. Some of it was delayed and, in some cases perhaps, people weren't paying adequate attention. But a significant part of the problem was that there was just *lots* of information, and to many humanitarian decision-makers, what to do with it or how to make sense of it all simply wasn't clear. For much of the ensuing year, a coherent analysis of all factors that could genuinely inform anticipatory—or even early—action was difficult to formulate. By mid-2020, the East African Food Security and Nutrition Working Group noted nearly 54 million people in immediate need of food assistance in the region (ICPAC 2020)—an increase of nearly 40 percent from 2019—and this assessment was missing information from several countries, meaning that the total would be higher. This example reveals the challenges of turning data into an analysis that can meaningfully inform responses to humanitarian crises. It is by no means the only example.

The nature of humanitarian crises has changed over the past two decades. First, the number of people affected has increased fivefold by the most common measure used to count populations in need by the Consolidated Appeals Processes (CAP) or Humanitarian Needs Overviews (OCHA 2020). The budgets to address these assessed needs have also grown rapidly, although not as rapidly as affected populations have grown (Development Initiatives 2020). This has resulted in an ever widening gap between need and response, an ever larger number of people who are more vulnerable, even to shocks of a lesser magnitude. Crises are also increasingly long-lasting—the most recent figures show that 27 countries are in “protracted crises” or crises that have lasted longer than five years (Development Initiatives 2019). In 2011 and again from 2016 to 2018, famine returned in several countries, renewing awareness of and attention to this age-old scourge (de Waal 2018). And both acute emergencies and protracted crises are increasingly driven by multiple causal factors, making any single-factor cause-effect analysis of crises oversimplified at best and perhaps dangerously misguided at worst. Conflict is nearly always one of the main drivers of these crises, with some 70 percent of the case load of humanitarian food security crises driven partly—and sometimes entirely—by violent conflict (OCHA 2019). These multiple factors make the analysis of crisis into what Ramalingam (2013) referred to as a “wicked problem”—complex, uncertain, and, frequently, unique. Of course, the changes brought about by the COVID-19 pandemic only make the situation more complex and uncertain, with the humanitarian sector more reliant on remote data collection and analysis, and social distancing required in the management of responses.

In response to these changes, the demand for anticipatory humanitarian action has increased (Lowcock 2019). Anticipating crises, rather than simply responding to them, could revolutionize humanitarian

action. But anticipating crises also requires much better forecasting *and* a willingness to act without knowing for certain that a crisis will materialize as forecasted. New initiatives such as the World Bank's Famine Action Mechanism (FAM), UNOCHA's Anticipatory Action program utilizing the Central Emergency Response Fund (CERF), and the Red Cross/Red Crescent Climate Centre's forecast-based financing (FBF) initiative, are attempting to do just this. These initiatives are increasingly being back-stopped by advanced predictive analytics (PA), machine learning (ML), and artificial intelligence (AI) (K4D 2020). The new elements in the humanitarian information ecosystem bring with them their own challenges.

Since at least the 1980s, trying to foresee crises before they occur has formally been part of the work of the humanitarian sector (and of course informally, for a long time before that). Broadly referred to as "early warning," the sector has used systems such as the Famine Early Warning Network (FEWS NET), the United Nations World Food Programme's (WFP) VAM (Vulnerability Assessment and Mapping) Unit, and the Integrated Food Security Phase Classification (IPC) system hosted by FAO. More specialized analysis—predominantly seasonal weather forecasting—has long been a mainstay of humanitarian information, though not necessarily decision making.

Making sense of the eco-system of humanitarian diagnostics (of all kinds) and the extent to which they inform humanitarian action is the subject of this brief discussion paper. In 2019, a study on the early warning systems functioning in East Africa found a high degree of confusion about early warning, the kinds of information that were being made available, and the extent to which they facilitated (or in some cases hindered) decision making (Maxwell and Hailey 2020a). Other recent analyses echo similar concerns, which have long dogged the early warning community: late information, leading to late response (Kimetrica 2020); poor data sharing (Hailey et al. 2018); significant political influence over the outcomes of early warning analysis (Maxwell and Hailey 2020b); inherent biases in information systems (Yusuf et al. 2019) or ways that predictive analytics and machine learning may automate sources of bias (K4D 2020); and the decades-old concern about information not connecting to practical action (Buchanan-Smith and Davies 1995). In today's world, this is less about users ignoring information and more about not knowing what to do with the flood of information and/or being hesitant to extrapolate from it.

# Purpose of the paper

The objective of this paper is to reduce some of the confusion arising from cases like the eastern Africa situation above by clarifying what has changed, noting the role of different components of humanitarian diagnostics, and demonstrating the links to action. Many kinds of “data” are available—resulting in confusion about what is what: confusion about what is actual empirical data and what is predictive or presumptive information, hesitation regarding available methods and the reliability of information, and uncertainty about the extent to which information and analysis is influenced by political actors and politicized motives. A 2003 paper by one of the authors attempted to address some of the confusion around humanitarian information systems as they existed then (Maxwell and Watkins 2003). Much has changed since—including our own views of humanitarian information systems and the role these systems can and should play in predicting, anticipating, mitigating, and responding to humanitarian crises.

This paper attempts to do five things. First, we revisit a 2003 paper on humanitarian information systems, to underline some of the things that have changed. Second, we clarify the difference between diagnostic evidence and evaluative evidence—and focus this

paper on the former. Third, we expand the typology of the kinds of diagnostic evidence and explain how they logically relate to each other. Fourth, based on findings from a number of studies, we attempt to summarize the problems of linking evidence to action, specifically linking early warning to anticipatory action—both humanitarian response and mitigation/prevention. And finally, we attempt to assess the contributions—potential and actual—of predictive analytics and machine learning to this rapidly expanding field.

A 2003 paper by Maxwell and Watkins illustrated an “ideal type” of humanitarian information system. This paper updates that paper by delineating the ways in which diagnostic evidence informs action (or not), outlining some of the constraints, and making (modest) suggestions for improvement. This paper is not intended to be an empirical report of any specific study, but draws on numerous studies (Lentz et al. 2019; Maxwell and Hailey 2020a; Maxwell et al. 2020).<sup>1</sup>

---

<sup>1</sup> It is also based on numerous interactions with humanitarian decisions-makers in the context of the COVID-19 pandemic who are overwhelmed with both information and demands for action but in which the former is simply not informing the latter.

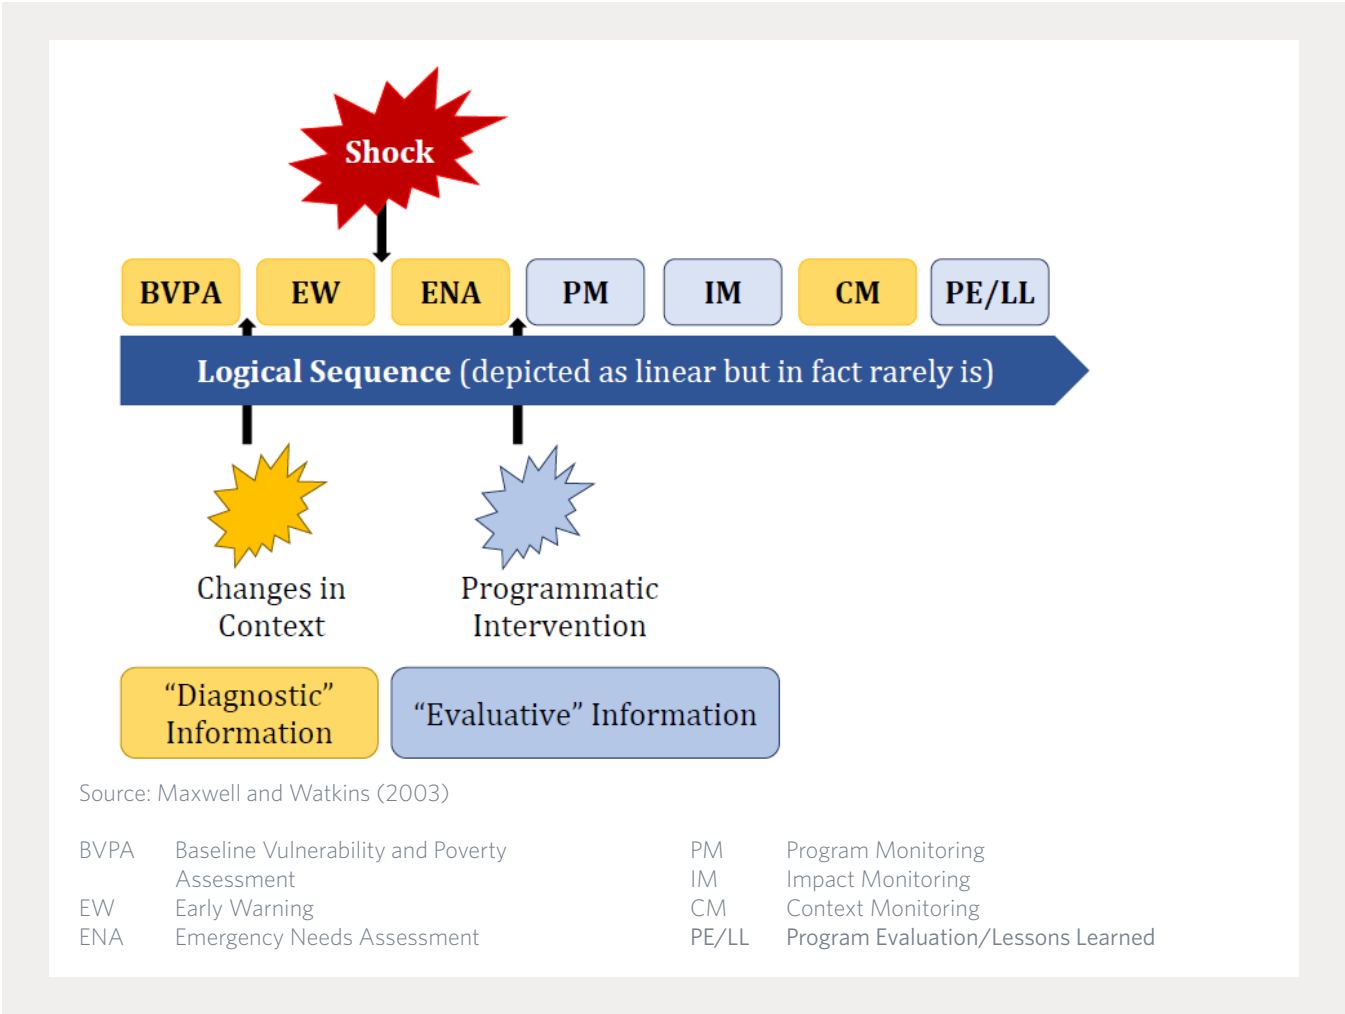
# Diagnostic evidence, evaluative evidence, and information systems

The 2003 paper attempted to include many kinds of humanitarian information collection and analysis activities and labeled them (misleadingly, at least by current standards—dare we say with “2020 hindsight”?) as “pre-programmatic” and “programmatic” information. Rather, it is more appropriate to label these as “diagnostic” and “evaluative” information

and evidence. Figure 1 is derived from the 2003 paper and updated with the more accurate language.

This figure provides a useful jumping off point to identify how humanitarian information systems have changed, where gaps remain, and what opportunities to improve humanitarian diagnostics exist. In Figure 1, we label the information related to moni-

Figure 1. A Hypothetical Humanitarian Information System



toring and evaluation, mostly on the right side of the diagram, as “intervention” evidence (blue in color). Much of the debate about “evidence” in contemporary humanitarian action has focused on this side of the diagram. It is heavily methodologically driven, mostly using, on the one hand, randomized controlled trials (RCTs) or quasi-experimental designs to test and compare the impact of interventions or, on the other hand, “systematic reviews” to compile and compare the results of multiple studies. This is extremely important information, because it builds up the evidence base about “what works.” Numerous actors in the humanitarian sector focus largely or even exclusively on the evaluative side of the picture: the Humanitarian Evidence Program (Krystalli and Emerson 2015), Evidence Aid (Allen 2020), Cochrane Reviews, for example.

The left side of Figure 1 and the box labeled “CM” (for context monitoring) are about diagnostic information (yellow in color). This information doesn’t address the question “What works?” but rather “What’s wrong?” or “What is likely to go wrong?” and “Who is (or is likely to be) affected?”, “How badly?”, “For how long?” and, crucially, “How certain are you about all of this?” Understanding the difference between these two different kinds of evidence—as well as the analytical practices that link them—is one important step towards sorting out the confusion about the situation in eastern Africa in 2019–20. Understanding the difference will also

help improve analysis for future crises. Like evaluative evidence systems, numerous organizations deal almost exclusively with the diagnostics side of the humanitarian evidence picture—FEWS NET, WFP/VAM, IPC, ACAPS, REACH, to name a few. This kind of evidence relies on different methods and has very different methodological limits.

The 2003 paper attempted to identify and define all these different kinds of information gathering activities and to fit them together into some kind of logical construct—demonstrating how easy it was to misinterpret information if one or more of these components was missing.

On the “diagnostic” side of Figure 1, several things in humanitarian information systems have changed. But one constant is that several information collection and analysis activities occur concurrently. While all ostensibly contribute to the same general goal, this patchwork of data collection and analysis does not seem to add up to a coherent “system.” Much of this is labeled “early warning,” but there is considerable confusion about what “early warning” is and isn’t. Early warning (EW) is about tracking causal factors and trying to determine with some degree of accuracy how likely those factors are to lead to shocks that negatively affect people. Early warning is only one component of humanitarian diagnostics (even though the whole of diagnostics is sometimes referred to simplistically as “early warning”).



# What has changed about humanitarian diagnostics since 2003?

We have observed twelve major changes in approaches to humanitarian diagnostics and analysis—and their links to early action—since 2003. They are outlined here, beginning with those more related to early warning and diagnostics and followed by those more related to the link to early action.

- **Substantial investment in improving current-status assessment.** This began with the SENAC (Strengthening Emergency Needs Assessment Capacities) project led by WFP in the mid-2000s, and includes the SMART initiative (Standardized Methods and Assessment of Relief and Transition), the development and institutionalization of the Integrated Food Security Phase Classification system (IPC Partners 2019), and the rise of agencies like ACAPS and REACH. The main tools resulting from the SENAC project include the Food Security and Nutrition Monitoring Survey (FSNMS) and the Emergency Food Security and Nutrition Assessment (EFSNA). SMART surveys predominantly measure the prevalence of malnutrition, but can also include measures of mortality, health, caring practices, and sometimes food security as well.
- **The invention and rise of Integrated Food Security Phase Classification (IPC).** IPC was invented as a current-status assessment and as a classification system with a mechanism for counting the population in need (or PIN number), and it was more recently formalized as the means of projecting future classifications as well. The concept of “projections” wasn’t common in humanitarian diagnostics in 2003.
- **Incorporating real-time monitoring.** Real-time monitoring (RTM) combines “context monitoring” (which we described as EW activities in the 2003 paper, but are no longer “early” in the context of a given shock or crisis) and “impact monitoring” or what we would probably today call “outcome monitoring.” **Outcome monitoring** tracks short-term outcomes like food access (through indicators like the Food Consumption Score or the Reduced Coping Strategies Index) and nutrition (GAM prevalence). RTM can also track on-going early warning indicators—climate, conflict, or markets (as well as more novel hazards such as COVID-19, locusts, etc.). While this certainly has implications for the immediate-term future, RTM is fundamentally about monitoring what is happening in the “here and now.” Some kinds of RTM include monitoring the validity of the **assumptions** on which projections are made. And finally, a related function of RTM can be “**hotspot identification**,” a hotspot being a location where a crisis is growing rapidly worse, even if it was not identified by early warning.
- **More publicly available data.** In 2003, satellite data was relatively expensive and not publicly available. Since then, data of all sorts has become much more widely available—not only remote sensing data, but price data, agricultural production data, and to some extent even conflict data (such as ACLED or Uppsala).
- **Improved predictive analytics (PA).** The introduction of much higher-speed computers and increasing availability of data has allowed predictive analytics (including machine learning) to take off (Varian 2014). Predictive analytics (PA) can be used at multiple stages in the humanitarian diagnostic and analysis system. They can be used to predict hazards (e.g., future drought) or predict outcomes (e.g., future food security). These forecasts could feed into scenario building

or directly trigger responses. PA can also be used to estimate current status for locations where data are not available (e.g., WFP's Global Hunger Map). PA is one category of **artificial intelligence (AI)** used in humanitarian contexts. Other AI tools can also facilitate reaching conclusions about an assessment, for example by weighing analysts' responses in an online forum (e.g., artificial swarm intelligence).

- **From CAP appeals to Humanitarian Needs Overview (HNO)/Humanitarian Response Plans (HRP).** These are annual processes in most countries in protracted crises. They estimate needs for the following year and identify programming to meet those needs while appealing for funding. In 2003, we noted that an emergency needs assessments—which we now label “current-status assessments” (CSA)—were triggered by a “shock.” At that time the idea of early warning was to predict, however crudely, emergent shocks. Now, most current-status assessments are regularly scheduled and do not wait for shocks to occur. The schedule is based partly on the presumptions of seasonality, but also partly in anticipation of the need for information for the HNO/HRP. But specific—and sometimes acute—shocks still happen in protracted crises so the need for hotspot identification remains, which can still be identified by traditional EW, or by RTM (or both, with the latter confirming the former). Some kinds of rapid emergency needs assessment are frequently still called for, completely outside the annual cycle defined by HNO/HRPs.
- **Building an inter-sectoral analysis.** While to date the IPC is the best organized diagnostic tool for estimating the population in need, it is focused on food security and nutrition only, so it does not really provide an overall estimate of needs. Various intersectoral analysis tools have been introduced, including the multi-cluster rapid assessment (MCRA), the multiple indicator cluster survey (MICS) and others. The most recent attempt at this is the Joint Intersectoral Analysis Framework (JIAF), which is not yet rolled out but proposes a common framework and methods for intersectoral, multi-agency needs analysis to inform decision-making and

response planning. The point of all these initiatives is a joined-up analysis that enables decision-making across the sector for a coherent response to crises—and to inform HNO/HRP processes.

- **The introduction of response analysis.** Response analysis aims to identify which response is appropriate and assumes that some degree of choice is possible. Response analysis was not a major part of information systems in 2003. Its importance grew after cash- and market-based programming became alternatives to in-kind food aid (Barrett et al. 2009, Maxwell et al. 2013). It is now integrated into many sectors of response.
- **The use of “triggers.”** These are “automatic” links between information and interventions. Triggers are intended to increase the speed of releasing traditional donor funds—and therefore mitigation or response. They are often linked to novel financing mechanisms such as insurance schemes or disaster bonds that require a clear “signal” for being released, rather than to a general “scenario.” Triggers were not in use in 2003; interest in triggers grew following the Somalia famine in 2011 (Maunder 2013).
- **Improved contingency planning.** Contingency planning—the primary link to early action—was around long before 2003, but the assumption in 2003 was that information was enough to set contingency plans in motion—even though a landmark study eight years earlier had said this was frequently not the case (Buchanan-Smith and Davies 1995). Now the planning is better linked to diagnostic information (Choularton 2007).
- **Anticipating rather than reacting.** The link to early action was implicit in the 2003 paper, aside from the fact that you couldn't “act” early if you didn't have information early and that lacking certain kinds of information made action even more difficult (for example, current-status information in the absence of baseline information). Anticipatory action is now the humanitarian *idée du jour*: with good predictive information linked to financing, intervention can precede a crisis, prevent human suffering, and save money on

humanitarian operations. This has various corollaries, such as “**forecast-based financing**,” etc. As noted above, it has been pilot tested by the World Bank, OCHA, the IFRC, and a variety of other agencies. Much of the anticipatory action has focused mainly on climatic factors and thus is mainly applicable to drought or flooding emergencies, although some funders look to changes in IPC classifications as possible “triggers” (see below). It is also heavily focused on “training” data from a limited number of contexts, particularly Somalia. Given a general fear of misallocating resources in an uncertain environment, a variation of anticipatory action is “**no regrets**” programming—early interventions that will add value, even if an anticipated shock does not develop or if its impact is not as serious as forecast.

- **Increased emphasis on “value for money” (VfM).** VfM is simply “cost effectiveness” redefined. But in this context, the notion that early intervention saves money has become an important rationale for improving information systems and preventing losses to livelihood assets and humanitarian conditions. Investments in improved information and analysis systems, but also in resilience programming and early interventions in mitigating humanitarian emergencies, are now justified as much on the basis of saving money as saving lives (e.g., FAO 2018). Put another way, in a context where there is never enough money to meet all needs, meeting humanitarian needs in the most cost effective way is critical.

# The confusion about diagnostics

The changes outlined above show that the nature of humanitarian diagnostics has improved significantly. Inadvertently, some of these changes, combined with lack of clarity about different kinds of information/analysis, contribute to the “data confusion” described in the first section. This confusion indicates a need for a new approach to thinking about how all of these components interact to ultimately inform interventions—both anticipatory and responsive. We consider two questions here in an attempt to address some of this confusion:

1. What are the existing components of a humanitarian diagnostics “system?” How do these various things fit together? Part of the current confusion is that it is not clear how these components are meant to work together.
2. Where, exactly, would predictive analytics or machine learning fit into or assist such a system?

## Components of humanitarian diagnostics

In addition to the changes outlined above, humanitarian diagnostics can now be broken out much further than the 2003 paper implied. Although still a component of the systems, **baseline analysis** (referred to in the 2003 paper as *baseline vulnerability and poverty assessment*) has probably declined in importance since then, particularly in countries where data on current status over a period of time is abundant. Baseline data was (and remains) important as a point of comparison for both early warning and current-status assessment, but in contexts having many years of trend data, the more accurate comparison is with those trends, rather than with a particular point in time presumed to represent a “baseline.”

**Early warning** (EW) has always tracked hazards and assessed the risk of those hazards causing

damage to people and their livelihoods—i.e., *causal factors*. The assumption is still that we can track long- and short-term trends, seasonality, and relatively fixed drivers like geography and infrastructure, as well as changing factors such as climatic and environmental drivers, macro-economic and political factors, production estimates, markets and prices, population movement and conflict . . . and predict when and where hazards will manifest, which populations will be affected, and how likely crises are to occur. These predictions are typically “scenarios” that focus on most-likely outcomes, but with an emphasis on “*likely*,” underscoring the probabilistic nature of EW.

Note that triggers are not the same thing as early warning as such, although they serve a similar purpose. Triggers are thresholds which, when breached, set in motion pre-arranged actions such as a payout for an insurance policy, or a scaling up of cash transfers in a social protection program. Triggers have also been suggested as one way to reduce the time for decision making or to take some of the politics out of decision making based on scenario analysis. For the most part (at least so far) triggers have worked best when linking a single threshold for a single hazard to a single response. However more recent efforts have considered using FEWS NET-projected transitions from IPC Phase 2 to 3 as a trigger, which would turn a scenario based on an in-depth analysis of multiple factors into a “trigger” for a multi-pronged response. But this is something of a departure from what “trigger” usually means.

The system has invested substantial money, time, and human resources into improving **current-status assessment** since 2003, particularly with the institutionalization of the Integrated Food Security Phase Classification (IPC) in the later 2000s and the supporting surveys that go with IPC analysis. IPC regularly reports figures on the current status of populations, classifying them into *phases* or severity categories,

either by livelihood zone or administrative zone, and providing a *population in need* (PIN) figure in each phase for each geographic unit. Anyone in IPC Phase 3 or higher is counted in the PIN for humanitarian food assistance. IPC has been instituted in some 35 countries—and a technically identical analytical protocol, Cadre Harmonisé, is used in 17 additional West African countries (IPC Info N.D.). IPC analyses take place usually once or twice a year, covering entire countries at an ADMIN2 level. They are usually based on WFP FSNMS or EFSNA surveys, supplemented by SMART surveys etc. IPC therefore compiles an impressive amount of data, but as a result, it is time-consuming and is always a bit out of date by the time the data is collected, cleaned, and analyzed; the situation is classified and vetted; reports are written and cleared; etc. Even in the best case scenario, the information is likely to be at least two months old by the time a report is finally issued—in extreme cases, the information may be up to a year out of date (and frequently information is simply not available at all).

At least partially as a result, the **projections** have become a much more important part of IPC analyses than they originally were. Projections take the current-status information as a kind of short-term baseline (not to be confused with BVPA information in Figure 1), and draw on early warning information to craft the most likely scenario for the short- and medium-term future (2–3 months and 4–6) months, and then “project” the number of people likely to be in each IPC phase by geographic zone (so the projections appear in the same form as the current-status assessment), numbers of people in different IPC phases, and an overall phase classification for geographic units of analysis.

The difference is that the current-status assessment is based on real numbers (i.e., empirical data) and the projections are based on assumptions about what is likely to happen to the current numbers. Those assumptions are ideally based on a thorough analysis of early warning factors, the development of scenarios, and a judgment about which is the “most likely” scenario. Having determined a “most likely”

Figure 2. Diagnostics: Relationships between EW, Projections, CSA, and Real-Time Monitoring




Table 1. Activity, Information Type, and Frequency

Information activity	Type of information					Frequency
	Causal factors	Predicted outcomes	Actual outcomes	Hotspot identification	Assumption verification	
Early warning	X			X		Ongoing
Current-status assessments (IPC)			X	X (often late)		Every 6 or 12 months
Projections		X				Every 4 months
Real time monitoring	X		X	X	X	Monthly
Baseline vulnerability assessment	X		X			Every 5 years/after a major change

Source: Authors

scenario, the process by which projected PIN numbers are assigned to each phase by livelihood zone or administrative unit is less transparent than the current PIN number, because it is entirely a question of human judgment—there is no algorithm. According to a few published articles (e.g., Choularton and Krishnamurthy 2019), FEWS NET and IPC, who both do this, get the projections right between two-thirds and three-quarters of the time—but that is only referring to the phase projected, not the PIN.

Nevertheless, the **projected PIN** is probably the single most important piece of actionable information that comes out of the entire system—because it refers to the future and at a range of time when governments, donors, agencies, and even local communities can still act. At a minimum, it says “humanitarian agencies are going to need to provide food assistance to this number of people in this place to deal with food insecurity 4–6 months from now.” (Note that there actually isn’t a good way to check the accuracy of the projected PIN.) Sometimes, projections can be very effective at provoking early action. In early 2017, when FSNAU projected famine in Somalia (for the second time in six years), it helped to trigger additional resources and led to a

more rapid response, and famine did not recur. (This is not to say that FSNAU information was solely responsible for preventing renewed famine: Somalia had a functioning government in 2017; a number of people in the humanitarian community determined not to let famine happen again; donors weren’t as concerned about counter-terrorism measures, etc.)

But several points can be made about IPC projections: (1) the methodology for coming up with the PIN number is opaque—even the people who do it all the time admit that it is more art than science (although some of them are pretty good artists!); (2) no clear means have been developed to monitor the assumptions that go into the projections; and (3) they are based on very expensive and time-consuming data-collection processes that can only be mounted once or twice a year—and a lot can change in between those times.

This is where the need for **real time monitoring** (RTM) arises. If done well, RTM can actually track changes in the context and note whether current humanitarian conditions are improving or deteriorating and thus serve as a form of “hotspot identification”—which is not really early warning because it is

real-time but may be the only means of identifying rapidly worsening situations. RTM can also supplement information useful to early warning analysis (in fact in some situations it is indistinguishable from EW). And RTM can be utilized to track the extent to which the assumptions in scenario development—and therefore projected PIN numbers—are borne out in reality, and therefore whether the projection are about right, too high, too low, etc. Unfortunately, RTM often doesn't do all of those things equally well.

Sometimes RTM systems operate where no EW system exists, but that may mean that information is not generated in time for early action, and at best can influence rapid response. But frequently, and importantly, many kinds of real time monitoring information is simply collected, processed and put out for general consumption, with no real analysis provided for what it means, and often with confusing links back to the other parts of the formal system. Real time monitoring is relatively new in humanitarian information systems and is still being developed in many contexts. (See annexes for several examples of RTM systems.) To reiterate, we propose that clarifying the role of diagnostic information could help to clarify the data confusion within the humanitarian sector, ultimately leading to more actionable outputs of the information system as a whole.

Figure 2 presents diagnostics information, which expands the left side of the Humanitarian Information System in Figure 1 to recognize that different types of information are collected at different times and inform different information activities.

Table 1 describes the types of information and how frequently it is collected for each information activity on the right hand side of Figure 2. Real time monitoring, for example, often occurs monthly and can include any combination of information types.

The reality is more complex than shown in Figure 2 and Table 1. Some humanitarian information systems do not treat these types of information as feeding into separate information activities. For example, “nowcasting,” a type of predictive analytics (PA), can support out-of-sample predictions for locations lacking current outcome data. Nowcasting models can be built using causal factors, correlated factors, and actual outcomes—either from other locations or from historical data from that location—to predict

current status. The ability to do out-of-sample predictions could cut information gathering costs.

However, unless and until there is a clear way of seeing how all the different bits work—or could work or *should* work—together, humanitarian diagnostics run the risk of producing a lot of noise and disjointed information and ideas without a clear idea of what it means or how it informs action.

## Links to action

---

The links to early or anticipatory action are not always clear. This includes the consideration of using triggers or scenarios as the link, but also the notion of “no regrets” programming (whether linked to *either* triggers or scenarios), the role of response analysis, and communications with affected communities (a seemingly obvious but frequently overlooked link to action).

Some early action decisions are based on a **trigger** system. These signal-driven systems may use CSA information or EW information on hazards. A trigger does precisely what is implied—it triggers a pre-set action. The pre-set action may be a more in-depth assessment of the situation, or it may be an actual response.

**Scenarios** are a more in-depth assessment of the situation, noting multiple causal factors and potentially multiple outcomes. They are more useful for an overall response, rather than for a single action. But they also require further judgment with regard to the appropriate response. Knowing food security will get “worse” or become “bad” (or that a drought will worsen) does not answer the questions decision-makers really want to understand (how many people? where? for how long? how severe? how much money is required?) and perhaps even what the appropriate response would be. This is what scenario building is about.

The “**no regrets**” approach to anticipatory action has grown in popularity—at least in principle—in recent years. This involves engaging in early action that will mitigate a crisis and will have beneficial impacts even if the crisis does not develop as anticipated. Constraints to early action (including “no-regrets”

action) have long been assumed to relate simply to a lack of finances to enable a response. While finance is one critical component, having a strong contingency plan in place that lays out exactly what has to happen is critical in cases requiring more than one single response (in other words, if the *only* response is to set cash transfers in place, a detailed contingency plan may be less critical than if there are cash transfers, asset protection programs, commercial livestock offtake initiatives, or other activities to be sequenced or undertaken at the same time). And, it is critical to have the implementation capacity to rapidly put the plan into action.

Linking early warning to early action also requires **response analysis**, or determining the most appropriate response or set of responses to a rapidly changing situation (Barrett et al. 2009). Ideally response analysis would be included in contingency planning, but in rapidly changing situations, the best response may change as well. Strong awareness of response options and when each is the most appropriate is a key consideration. For instance, providing cash transfers has proven to be an important intervention to mitigate food insecurity among urban populations during the COVID-19 pandemic, but in several cases—for reasons perhaps unrelated to the pandemic as such—food price inflation made cash transfers a less useful intervention than in-kind food or vouchers. It wasn't always clear that all of these considerations were factored into contingency planning.

Finally, only a few mechanisms actually **disseminate early warning information to at-risk communities**. Failure to take any one of these into account can lead to the failure of early or preventative action. All of this has long been known (e.g., Choularton 2007) but the reason for the failure of early action has long been presumed to be poor information or delays in finance, rather than a lack of strong, actionable contingency plans.

Often, some early actions are ruled out by the misalignment between the timeframe of the early warning projection and the timeframe required for action.

For example, if a fodder project requires six months to roll out, but the existing EW system provides information only three months ahead of time, fodder projects will not be feasible. A longer time-window of early warning will be less precise, but will expand the set of feasible responses. Some of the feasible responses in the expanded set will require longer lead time, but may also be lower cost. Just as there are tradeoffs between EW accuracy and timeframe, there can be tradeoffs between response costs and the timeframe to set them up. Understanding the feasible set of options in response to common or severe hazards and the timeframe required to implement them can help inform the time and information requirements for the EW system, and clarify what sort of tradeoffs decisionmakers can live with (e.g., certainty versus speed and cost). In other words, the usefulness of EW system information is linked to the set of potential responses available.

Two additional points are worth noting here. First, sometimes the range of potential mitigation or response measures are not known. And second, care has to be taken that known interventions are not compounding biases or marginalizing specific groups. In both cases, human judgment and analysis will remain essential to understanding the impacts of hazards that are data-sparse and/or rare, or when a unique combination of hazards—which could either amplify or dampen one another—has not been modeled before. Projections for the hard-to-predict hazards might be missing until they are absolutely critical (e.g., few systems were systematically monitoring the impact of a large-scale public health emergency until the COVID-19 pandemic actually occurred). Systems need the flexibility to ramp up quickly, and skilled people need to figure out how to incorporate them into the more standard EW system. An organizational implication is that, at a minimum, someone or some institution has to assess the net impact of all these factors holistically. Currently many of these analyses are spread across multiple institutions and systems.



# Discussion (and an alternative recommendation)

Several issues arise that merit further discussion. The first is how recent advances in predictive analytics and machine learning can enhance the early warning-early action (EW-EA) equation. A second is how recent advances in theory can be incorporated into EW-EA. A third concern, not discussed in depth here, is about the political influences on both information and analysis (Maxwell and Hailey 2020b). And finally, we suggest an alternative to specifying the necessary components of an information system.

## Advances in predictive analytics and machine learning

Predictive analytics (PA) can be applied throughout the diagnostic system. First, PA as well as other analytical techniques (e.g., artificial swarm intelligence, in-person Delphic approaches, and food security analysts compiling scenarios from a wide range of materials) can support converting disparate streams of data into an information activity, such as producing RTM findings or generating EW scenarios. Second, PA has recently received a great deal of interest in the practitioner, researcher, and donor communities because some PA models offer a way to synthesize large amounts of data to generate diagnostic evidence. The diagnostic evidence can include (1) estimated outcomes related directly to food security, such as PIN and IPC phases, or (2) estimated outcomes associated with food insecurity, such as the likelihood of hazards (e.g., meteorological forecasts) or other factors (e.g., conflict). When used as part of a “signal” approach, PA estimates can directly trigger a response. Currently, trigger-based models are used mostly for single-hazard/single-response actions, such as triggering drought insurance or flood mitigation or triggering funding due to a change in IPC status. Whether PA can generate reliable and valid triggers for multiple hazards remains to be seen.

The potential of PA has not yet been reached in either triggers or scenario-based diagnostic information. Questions remain about the role of PA in assessments that incorporate multiple streams of information, some of which may not be incorporated in PA models. Incorporating predictive analytics into a humanitarian diagnostics system will require modelers and decisionmakers to make decisions about what to prioritize in their models. There is not yet adequate consensus on which tradeoffs matter, but the choices depend, in part, on the goals of the humanitarian diagnostic system. Examples of such tradeoffs include the following:

- Should PA models be highly specific to certain livelihoods or hazards or be consistent across a country or region? The latter offers comparability. The former may ensure accuracy for factors that are otherwise difficult to monitor.
- Some modeling techniques (e.g., linear and logistic regressions) allow for easy interpretation of the coefficients on variables on the right hand side of the equations, supporting analysis of drivers. Other modeling techniques (e.g., neural networks) maximize accuracy (i.e., fit) but may be less helpful for identifying the contribution and importance of individual variables (Paul et al. 2018). The latter models are considered to be more “black-box.” Thus, when should modelers and decisionmakers prioritize interpretability over fit? Ensemble models, which average across predictions, may provide balance (Varian 2014).
- Do decisionmakers prioritize minimizing errors of exclusion or minimizing errors of inclusion? PA models can be weighted to favor one or the other.

Beyond model development, there are also tradeoffs and difficult choices around best supporting users and decision makers:

- From a capacity and cost perspective, should models be developed to rely on secondary data or require primary data collection? Primary data can be tailored to meet the needs of humanitarian diagnosticians but can be costly and require ongoing or multi-year investments to be most useful. Secondary data can be slow to arrive without well-defined sharing protocols, which can limit the usefulness of PA for early warning.
- PA models require data platforms and capacity building. Do interest and support exist for building the capacity to process data, update models, interpret outputs, and integrate PA findings into existing systems? Or is it better to strengthen existing practice? The question of who will populate, manage, and update these platforms is, perhaps, a much more expensive and onerous question than building the models themselves.
- Should the outputs generated be triggers or should they be a component in a broader scenario? Paul et al. (2018) warn of “excessive trust” in machine learning, which they define as “unquestioning acceptance of model results, which can result in misinformed choices when models do get it wrong” (p. 40). At the same time, triggers, when properly identified, can offer clear-cut decision-points for intervention and early action.

Finally, it is rare to involve affected populations in model development, in reporting, in feedback on model outcomes, or in ensuring a machine learning (ML) product that is understandable and interpretable (K4D 2020). Van den Homberg et al. (2020) report “when accountable artificial intelligence is lacking even in non-emergency contexts in the Global North, the likelihood of artificial intelligence in emergency contexts in the Global South harming vulnerable populations is dramatically increased” (p. 2). Yet, including communities’ feedback on model design (e.g., on possible triggers or overlooked causal factors) and building in reporting mechanisms (e.g., through cell-phone or radio alerts) could significantly enhance the impact of early warning, support fairness and inclusivity, and reduce biases.

Regardless of the decisions made, these tradeoffs and the modeling assumptions ought to be made explicit. Coyle and Weller (2020) argue “demanding that ML systems be explainable is likely to make the

tradeoffs between different objectives far more explicit than has been the norm previously” (p. 1434).

Further, the value of PA will be greatest when they are set up to work with the existing information systems—IPC, WFP/VAM, FEWS NET, SMART etc. Employing PA to address “gaps” in current systems—rather than inventing new systems—might be its most effective use.

Humanitarian diagnostics may need **multiple approaches for multiple hazards**: perhaps the attention to automation ought to be towards those factors that are the most predictable and for which there is the greatest amount of (publicly available) data. It may be most effective to focus human capacity on the less easily predictable hazards or on interpreting the impact of multiple hazards. This raises a challenge of how to support analysts needing to integrate findings from different approaches in a valid and replicable way. Without a process on how to integrate techniques like PA into existing systems, there is a risk of duplication of effort and of inconsistent messaging, not unlike the 2019–20 multiple-hazard case in East Africa.

## Advances in theory

An additional challenge is to take recent theoretical insights and make them useful for prediction. Much of the recent advances in information collection and data management are essentially atheoretical in nature. For instance, although information on different sectors might exist, little is done to interpret even the relationship between food security and nutrition, let alone the complex intersection of food security, water, health, and nutrition. Likewise, even though the livelihoods analytical framework has been in use for decades, it is only systematically used by a handful of information systems to interpret the likely impact of a shock. Many information systems simply equate a shock with food insecurity or other outcomes. Likewise, famine early warning doesn’t particularly take into consideration recent insights on famine causation, including attention to idiosyncratic factors or what Howe (2018) labeled “the hold” in famine systems. This is the difference between being

“data driven” and “understanding how things work” and driving the response accordingly.

The role of theory in PA varies by machine learning technique; there may be tradeoffs between predictive accuracy and inclusion of key drivers. Some models incorporate variables based on theories of what contributes to food insecurity and famines. Other approaches may prioritize accuracy, either by looking for associations among possible proxies (e.g., cell phone top-up rates) or by selecting models that do not lend themselves to interpretation of coefficients (e.g., more “black box” modeling techniques such as neural networks). Such models trade off insights into the causal drivers of food insecurity for higher rates of predictive accuracy. As a result, these models are less able to assist in identifying factors that could be changed through longer-term development programming aimed at addressing the drivers of food insecurity. As long as mechanisms exist for identifying these interventions, selecting improved accuracy over interpretability could be a worthwhile tradeoff.

A significant risk of PA is that the easy availability of certain data can over-focus predictive attention on these data, at the risk of missing other important factors. For example, meteorological forecasts are substantially richer than they were twenty years ago (Funk et al. 2019). These improvements have helped spur development of forecast-based financing (Coughlan de Perez et al. 2019). Yet, other factors may be equally or more important for predicting food insecurity, such as conflict. Our understanding of the relationship between types of conflict and food security is under-developed (Maxwell and Hailey 2020b). This lack of understanding is a likely contributor to the limited number of food security models that either incorporate conflict or find a meaningful relationship between the two.

## The politics of information and analysis

---

This paper doesn't touch on concerns related to the politicization of information. For a detailed discussion, see Maxwell and Hailey (2020b). However, those concerns are germane to this discussion. For

example, one justification for predictive analytics and machine learning is that it “takes politics out of the equation.” It is, however, equally likely that PA/ML moves the politics to other individuals (modelers and their managers), who may not be well placed to even identify or understand the political nudges they might be getting. Paul et al. (2018) warn of “excessive trust” in ML, which they describe as “unquestioning acceptance of model results, which can result in misinformed choices when models do get it wrong” (p. 40). Further, implicit biases can be built in to PA. For example, PA that relies on cellular phone data needs to recognize that who owns a cell phone varies both within households by gender and across households by income and access to charging points (Paul et al. 2018; K4D 2020).

## An alternative approach?

---

Numerous changes in the humanitarian sector have enhanced data availability, analytical procedures, and mitigation/response options in the past fifteen years. Given expanding needs and shrinking budgets, identifying the emerging and ongoing gaps in humanitarian diagnostics and resolving the resulting “information confusion” is of critical importance. One approach to resolving these gaps is to reframe the goal of humanitarian diagnostics. Much of the approach to early warning and information systems generally is about collecting many kinds of data and assuming that they add up to some kind of picture of what is happening or is likely to happen, in effect asking, “What data are needed for early warning or RTM?” Perhaps the more important question for humanitarian diagnostics should be, “What is the minimum information needed to enable a preventative or early response?”

We propose reframing humanitarian diagnostics to be more response-focused. In other words, if the set of feasible early and anticipatory action responses and their required timelines can be clearly identified and specified, the questions for diagnostics become, “What do we need to know and how far in advance do we need to know it in order to set in motion either specific actions (triggers) or a more complex contingency plan (scenarios)?” Not all information necessarily needs to be tracked. Such an approach

would start with known/expected hazards (but would have to retain sufficient flexibility to be able to engage novel hazards—if we've learned nothing else from 2020, let us at least learn this!). For known and anticipated hazards, the questions for mitigative or preventive actions would be as follows:

- What actions can/should be taken to mitigate them?
  - What contingency plans would be needed to mitigate them?
  - What information would be needed to put mitigation/contingency plans into action?
  - How much time would be needed in advance for that information to be programmatically useful?
  - What degree of certainty must that information have?
  - What is the relationship between the expected magnitude and severity of the shock on the one hand and the cost of preventive or early action on the other hand?
  - What else (besides contingency plans) would be needed to mitigate the identified hazards?
  - What learning and feedback (links to evaluative information) would need to be built in?
- What actions can/should be taken to respond to a humanitarian emergency?
  - What contingency plans would be needed to respond to the emergency?
  - What information would be needed to put response plans into action?
  - How far ahead of time does that information need to be gathered?
  - What degree of certainty must that information have?
  - What is the relationship between the expected magnitude and severity of the shock on the one hand and the cost of response on the other hand?
  - What else (besides response plans) would be needed to mitigate the emergency?
  - What learning and feedback (links to evaluative information) would need to be built in?

For actual humanitarian response, a similar set of questions drive information needs:

Addressing those questions in advance, and on a repeating basis, would enable a much leaner, more flexible information and analysis system and would identify information needs more accurately and precisely. This in turn may help to reduce the confusion related to humanitarian diagnostics and enable more effective links between information systems and anticipatory action.

# References

- Allen, Clare. 2020. "Evidence Aid and COVID-19 – Where Should We Go Now?" Global Alliance of Disaster Research Institutes.
- Barrett, Christopher, Robert Bell, Erin Lentz, and Daniel Maxwell. 2009. "Market Information and Food Insecurity Response Analysis." *Food Security* 1(2): 151–68.
- Buchanan-Smith, Margie, and Susanna Davies. 1995. *Famine Early Warning and Response: The Missing Link*. London: Intermediate Technology Publications.
- Choularton, Richard, 2007. "Contingency Planning and Humanitarian Action: A Review of Practice." HPN Network Paper 59. London: Humanitarian Practice Network.
- Choularton, Richard, and Krishna Krishnamurthy. 2019. "How Accurate Is Food Security Early Warning? Evaluation of FEWS NET Accuracy in Ethiopia." *Food Security* 11: 333–44.
- Coughlan de Perez, Erin, Maarten van Aalst, Richard Choularton, Bart van den Hurk, Simon Mason, Hannah Nissan, and Saroja Schwager. 2019. "From Rain to Famine: Assessing the Utility of Rainfall Observations and Seasonal Forecasts to Anticipate Food Insecurity in East Africa." *Food Security* 11:57–68. <https://doi.org/10.1007/s12571-018-00885-9>
- Coyle, Diane, and Adrian Weller. 2020. "Explaining" Machine Learning Reveals Policy Challenges." *Science*. 368(6498): 1433–34. <https://science.sciencemag.org/content/368/6498/1433>
- Development Initiatives. 2019. "Global Humanitarian Assistance Report." Bristol, UK: Development Initiatives.
- Development Initiatives. 2020. "Global Humanitarian Assistance Report." Bristol, UK: Development Initiatives.
- de Waal, Alex. 2018. *Mass Starvation. The History and Future of Famine*. Cambridge: Polity Press.
- FAO. 2018. "Acting Early to Prevent Humanitarian Emergencies." Rome: FAO. <http://www.fao.org/news/story/en/item/1145645/icode/>
- Funk, Chris, Shraddhanand Shukla, Wassila Mamadou Thiaw, James Rowland, Andrew Hoell, Amy McNally, Gregory Husak, Nicholas Novella, Michael Budde, Christa Peters-Lidard, Akhlalil Adoum, Gideon Galu, Diriba Korecha, Tamuka Magadzire, Mario Rodriguez, Miliaritiana Robjhon, Endalkachew Bekele, Kristi Arsenault, Pete Peterson, Laura Harrison, Steven Fuhrman, Frank Davenport, Marty Landsfeld, Diego Pedreros, Jossy P. Jacob, Curt Reynolds, Inbal Becker-Reshef, and James Verdin. 2019. "Recognizing the Famine Early Warning System Network: Over 30 years of drought early warning science advances and partnerships promoting global food security." *American Meteorology Society*. June: 1011-1027.
- Hailey, Peter, Jeeyon Janet Kim, Erin McCloskey, Maria Wrabel, and Daniel Maxwell. 2018. "The Constraints and Complexities of Information and Analysis in Humanitarian Emergencies. Evidence from Somalia." Medford MA: Feinstein International Center, and Nairobi: Centre for Humanitarian Change.
- ICPAC (IGAD Climate Prediction and Application Centre). 2020. "Food Security and Nutrition Working Group Report for September 24, 2020." Nairobi: ICPAC.
- IPC Info. N.D. <http://www.ipcinfo.org/>
- IPC Partners. 2019. *The Integrated Food Security Phase Classification Technical Manual. Version 3*. Rome: FAO.
- K4D (Kevin Hernandez and Tony Roberts) 2020. "Predictive Analytics in Humanitarian Action: A Preliminary Mapping and Analysis." K4D Emerging Issues Report 33. Brighton, UK: Institute of Development Studies.
- Kimetrica. 2020. Phase II Report. "Ethiopia Early Warning System Enhancement: Executive Summary." Nairobi: Kimetrica.

- Krystalli, Roxani, and Lauren Emerson. 2015. "Synthesizing Practices of Evidence Appraisal in the Humanitarian Field." *Humanitarian Evidence Program*. Medford MA: Feinstein International Center.
- Lentz, Erin, Hope Michelson, Katherine Baylis, and Yang Zhou. 2019. "A Data-Driven Approach Improves Food Insecurity Crisis Prediction." *World Development* 122: 399–409.
- Lowcock, Mark. 2019. "Anticipation Saves Lives: How Data and Innovative Financing Can Help Improve the World's Response to Humanitarian Crises." Lecture delivered at London School of Economics, 2 December. New York: OCHA.
- Maunder, Nick. 2013. "Triggers for Early action: A Scoping Report." CHASE Report. London: DFID.
- Maxwell, Daniel. 2019. "Famine Early Warning and Information Systems in Conflict Settings: Challenges for Humanitarian Metrics and Response." Conflict Research Programme Occasional Paper. <https://sites.tufts.edu/reinventingpeace/files/2019/11/CRP-Famine-Early-Warning-and-Information-Systems-in-Conflict-Settings-D-Maxwell-final-20191119.pdf>
- Maxwell, Daniel, Peter Hailey, Abdullahi Khalif, and Francesco Checchi. 2020. "Determining Famine: Multi-Dimensional Analysis for the 21st Century." *Food Policy* 92 Article 101832.
- Maxwell, Daniel, and Peter Hailey. 2020a. "Towards Anticipatory Information Systems and Action: Notes on Early Warning and Early Action in East Africa." Medford MA: Feinstein International Center, and Nairobi: Centre for Humanitarian Change.
- Maxwell, Daniel, and Peter Hailey. 2020b. "The Politics of Information and Analysis. Humanitarian Information Systems and Analysis in Famines and Extreme Emergencies." Medford MA: Feinstein International Center, and Nairobi: Centre for Humanitarian Change.
- Maxwell, Daniel, Heather Stobaugh, John Parker, and Megan McGlinchy. 2013. "Response Analysis and Response Choice in Food Security Crises: A Roadmap." HPN Network Paper 73. London: Humanitarian Practice Network.
- Maxwell, Daniel, and Ben Watkins. 2003. "Humanitarian Information Systems and Emergencies in the Greater Horn of Africa: Logical Components and Logical Linkages." *Disasters* 27 (1): 72–90.
- OCHA (UN Office for the Coordination of Humanitarian Affairs). 2019. 2020 Humanitarian Needs Overview. New York: OCHA.
- OCHA UN Office for the Coordination of Humanitarian Affairs). 2020. 2020 Humanitarian Response Plan. New York: OCHA
- Oxfam. 2017. "From Early Warning to Early Action in Somalia. What Can We Learn to Support Early Action to Mitigate Humanitarian Crises?" Nairobi: Oxfam.
- Paul, Amy, Craig Jolley, and Aubra Anthony. 2018. "Reflecting the Past, Shaping the Future: Making AI Work for International Development." USAID. <https://www.usaid.gov/sites/default/files/documents/15396/AI-ML-in-Development.pdf>
- Ramalingam, Ben. 2013. *Aid on the Edge of Chaos*. Oxford: Oxford University Press.
- Van den Homberg, Marc, Caroline Geveart, and Yola Georgiadou. 2020. "The Changing Face of Accountability in Humanitarianism: Using Artificial Intelligence for Humanitarian Action." *Politics and Governance* 8 (4): 456–67.
- Varian, Hal. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives*. 28(2): 3–28.
- Yusuf, Badra, Guhad Aden, and Stephen McDowell. 2019. "Qualitative Study of Vulnerability, Inclusion and Exclusion: Final Report." Mogadishu, Somalia: Inter-riverine Research and Consulting.

The Feinstein International Center is a research and teaching center based at the Friedman School of Nutrition Science and Policy at Tufts University. Our mission is to promote the use of evidence and learning in operational and policy responses to protect and strengthen the lives, livelihoods, and dignity of people affected by or at risk of humanitarian crises.

Twitter: @FeinsteinIntCen

[fic.tufts.edu](http://fic.tufts.edu)

