Policy Research Working Paper 9413

Stochastic Modeling of Food Insecurity

Dieter Wang
Bo Pieter Johannes Andrée
Andres Fernando Chamorro
Phoebe Girouard Spencer

WORLD BANK GROUP
Fragility, Conflict and Violence Global Theme
September 2020

Policy Research Working Paper 9413

Abstract

Recent advances in food insecurity classification have made analytical approaches to predict and inform response to food crises possible. This paper develops a predictive, statistical framework to identify drivers of food insecurity risk with simulation capabilities for scenario analyses, risk assessment and forecasting purposes. It utilizes a panel vector-autoregression to model food insecurity distributions of 15 Sub-Saharan African countries between October 2009 and February 2019. Statistical variable selection methods are employed to identify the most important agronomic, weather, conflict and economic variables. The paper finds that food insecurity dynamics are asymmetric and past-dependent, with low insecurity states more likely to transition to high insecurity states than vice versa. Conflict variables are more relevant for dynamics in highly critical stages, while agronomic and weather variables are more important for less critical states. Food prices are predictive for all cases. A Bayesian extension is introduced to incorporate expert opinions through the use of priors, which lead to significant improvements in model performance.

This paper is a product of the Fragility, Conflict and Violence Global Theme. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The authors may be contacted at dwang5@worldbank.org.

Stochastic Modeling of Food Insecurity

Dieter Wang*[1], Bo Pieter Johannes Andrée†[1], Andres Fernando Chamorro ‡[2], and Phoebe Girouard Spencer§[1]

[1]World Bank Group; Fragility, Conflict and Violence
[2]World Bank Group; Development Data Analytics and Tools

JEL: C01, C14, C25, C53, O10

Keywords: Food insecurity, famine risk, variable selection, stochastic simulation, panel vector-autoregression, expert opinion.

*Corresponding author. E-mail: dwang5@worldbank.org
†E-mail: bandree@worldbank.org
‡E-mail: achamorroelizond@worldbank.org
§E-mail: pspencer1@worldbank.org

This work was prepared as background for the Famine Action Mechanism (FAM). Support from the State and Peace-Building (SPF) Trust Fund (Grants N. TF0A7049 and TF0A5070) is gratefully acknowledged. The authors would like to express their gratitude to Zacharey Carmichael, Aart C. Kraay and Nadia Piffaretti for their valuable comments, guidance and continued support. The authors also thank Cathrine Ansell, Bledi Celiku, Harun Dogo, Nicholas Haan, Therese Norman, John Plevin, Nicola Ann Ranger, as well as workshop participants of the 14th Annual Workshop of the Household Conflict Network (Medellin, Colombia) for their comments and feedback.

1 Introduction

"Zero Hunger" is the second of the 17 United Nations (UN) Sustainable Development Goals adopted by all Member States in 2015. Achieving this target by 2030 is a central challenge, as the UN Food and Agricultural Organization (FAO) reports more than 2 billion people are currently estimated to suffer from hunger or food insecurity (FAO, 2019). Globally, more than 820 million people are undernourished. About 20% of those undernourished are located in Africa, in particular the Sub-Saharan Africa region, which has also experienced the fastest rise of malnutrition (FAO, 2019). The topic has gained renewed attention due to famines in Somalia in 2011-12 and the "four famines" (Maxwell et al., 2020) that took place in the Republic of Yemen in 2016 and South Sudan, Somalia, and northeast Nigeria since 2017. In their recent report, the World Food Programme (WFP) stresses the additional burden the COVID-19 pandemic imposes on already acutely food-insecure populations (WFP, 2020). As of the time of writing, the FAO raised a food crisis warning as locusts swarms are ravaging crops and pasture in East Africa (FAO, 2020).

In the fight to eradicate hunger, some major recent developments were made in the measurement and classification of food insecurity. For instance, the Food Insecurity Experience Scale (FIES) used in the SDGs distinguishes between three levels of severity: food security or mild food insecurity, moderate food insecurity, and severe food insecurity. In the latter stage, people have typically run out of food and have not eaten for one or more days. Different from this survey-based approach, the Integrated Food Security Phase Classification (IPC) system was developed to support decision-making with a clear analytical focus. The Food Security and Nutrition Analysis Unit (FSNAU) of the FAO pioneered this system in Somalia in 2004. Since then, the IPC framework has been adopted by various international organization and humanitarian response agencies. In our study, we rely on IPC data published by the Famine Early Warning Systems Network (FEWS NET).

The IPC framework is constructed to track and alert international and civil organizations to food insecurity situations, ultimately aiming to prevent the most extreme case (famine) at all costs. The IPC scale has five phases of acute food insecurity, where Phase 1 represents none/minimal and 5 represents catastrophe/famine (Table 1). A population enters this phase if, among other criteria, the crude death rate is at least 2 per 10,000 per day and the under-five death rate is at least 4 per 10,000 per day (IPC Global Partners, 2019). Urgent action is required to avoid populations reaching Phase 3 or higher. The latest version of the IPC framework 3.0 introduces a forward-looking "Famine Likely" phase to call for urgent action. Nonetheless, this may already be too late as mortality has already set in and claimed many lives. The 2011 famine that took place in parts of Somalia remains a lamentable reminder for the consequences of inaction (Salama et al., 2012).

The rationale behind our work is fully aligned with that of famine prediction and prevention. What distinguishes it from existing approaches is its statistical framework, forward-looking perspective and stochastic treatment of food insecurity. The quantitative, empirical framework also allows us to identify the most relevant agronomic, environmental, conflict or economic drivers behind famine dynamics among a vast set of possible candidates, in a systematic way (Pape et al., 2018). As Lentz et al. (2019) demonstrate, systematically incorporating such information leads to significant improvements over existing methods.

A quantitative model also allows analyzing alternative scenarios. Once a model has been constructed and its parameters estimated, it is possible to explore counterfactual situations. For instance, one can construct an entire stochastic catalog that assigns probabilities to events of interest. Quantifying such events has clear benefits for decision makers and contingency planners. A particularly attractive feature of a model is the ability to trace its predictions back to the individual drivers of famine or the assumptions made at the outset. This keeps the model transparent and its results interpretable. Furthermore, if the model incorporates the time-dependent characteristics of famine, we can use it to obtain probabilistic forecasts several steps ahead. Once again, the benefits for decision makers lie in the ability to explore plausible scenarios and adjust accordingly.

Food insecurity is a complex problem with multiple causes, several concurrent dynamics and catastrophic consequences. A wide range of experts such as humanitarian organizations, agricultural specialists, geographers, conflict experts, poverty researchers and economists are involved to tackle this multi-faceted challenge. Region- and domain-specific knowledge are central to understand the drivers of famine and predict its likely developments. One of the key features of our proposed model is that it leaves prominent room for expert opinion. This is not only essential for policy makers to adapt the model to situation-specific conditions. It is particularly important when facing unprecedented scenarios that are not contained in the data, yet obvious to human decision-makers. Moreover, expert judgement is necessary to augment the short recorded famine history. Our model therefore seeks to complement existing efforts in food insecurity prevention by providing a statistically rigorous foundation.

The relevance of approaches to predict food insecurity and evaluate risks over multiple time horizons is likely going to increase in coming years as key drivers of food insecurity are expected to worsen into the 21st century. Historical successes in eradicating undernourishment have largely occurred alongside substantial pressure on environments (Andrée et al., 2019; Stern et al., 1996). Further degrading environments, environmental change (Ingram et al., 2010; Myers et al., 2017), more frequent weather extremes (Diogo et al., 2017), desertification (Grainger, 1990), at the backdrop of growing populations, will continue to put further pressure on the ability of future agricultural systems to produce and distribute food. These developments will be particularly felt by the rural poor that rely on local natural assets for food consumption and income (Barbier & Hochard, 2018; Barrett & Bevis, 2015; Duraiappah, 1998).

The remainder of this paper is structured as follows. Section 2 describes the data used and introduces a convenient transformation to model famine risk distributions. Section 3 introduces the single- and multi-country stochastic models that will be used throughout the analysis. Section 4 then describes how we narrow down 1,670 possible predictors and select a relevant set of around 30 predictor variables. Section 5 employs the estimated model and selected variables to make expanding one-step ahead predictions for each point in time. Section 6 integrates expert opinion into the model and demonstrates significant improvements for the case of Afghanistan.

Figure 1: Overview of IPC distributions for 15 all countries

The individual graphs show the entire distribution of IPC1 (green), IPC2, (yellow) and IPC3+ (red) for all 15 countries based on FEWS NET data. This distribution is based on the population-weighted IPC levels in subnational administrative regions. Prior to 2016, the reports were reported four times a year, afterwards three times a year (the vertical lines demarcate the reporting dates). The vertical axis runs between 0% and 100%.

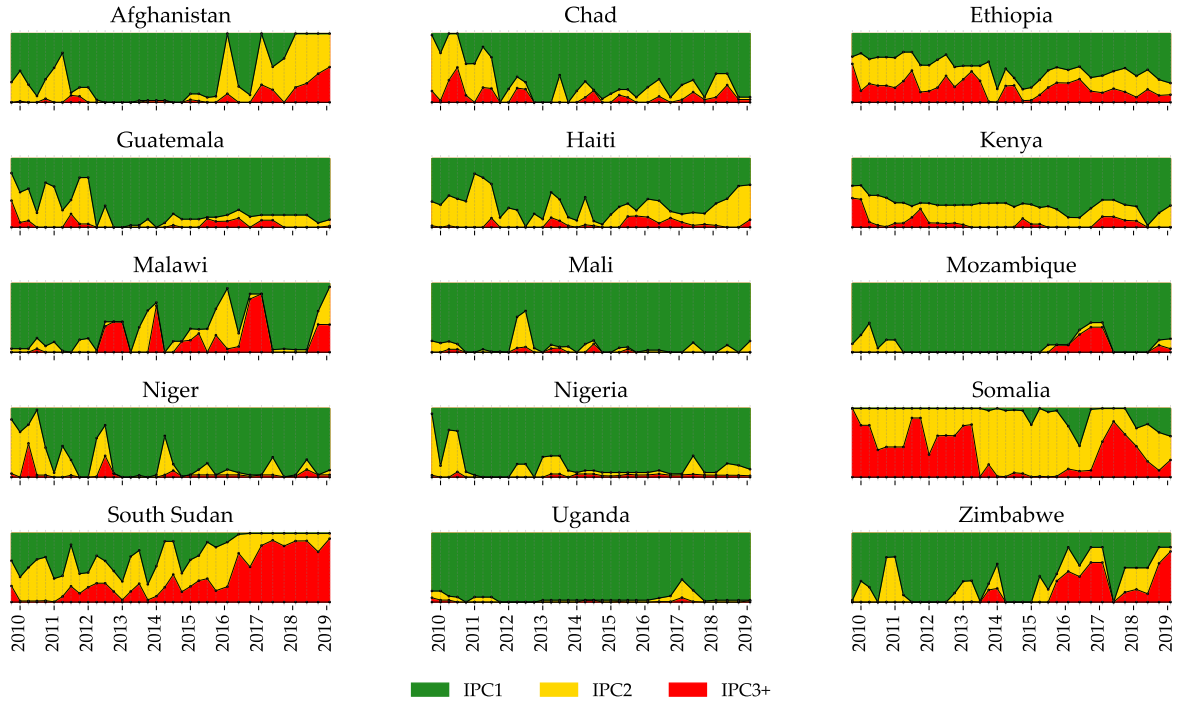

2 Data

In our study we rely on the Integrated Phase Classification Acute Food Insecurity (IPC) framework for the definition of famine risk and food insecurity (see Table 1). We obtain the data from FEWS NET as shown in Figure 1. These data are reported on the first-level administrative country subdivisions (districts).

One of our main innovations is to model food insecurity distributions for a given country, rather than a single index metric.[1] An example for this would be the established IPC framework, which classifies districts in a country based on the IPC stage the worst 20th percentile of the population finds itself in. For illustration, consider the two populations in Figure 2. The left population is distributed as 60%, 10%, 10%, 10% and 10% in increasing severity phases from IPC1 to IPC5. The overall IPC phase for the entire region would, applying the standard classification threshold of 20%, be IPC4, as the most severely affected 20% is in IPC4 or worse. However, the second population in the right panel receives the same classification, following the same rationale as before, even though its situation is undoubtedly more acute. Our distributional approach models dynamics of the distribution of district populations across all phases, rather than a single index. To calculate this food insecurity distribution for a country, we take the population share

[1]Such single metrics could be calculated as (weighted) averages, specific quantiles or other index construction methods that yield in a single metric.

Figure 2: Hypothetical scenarios leading to IPC4

These two scenarios illustrate how two entirely different situations of food insecurity lead to the same classification under the 20% rule, where a region or country is classified based on which IPC classification the 80th-percentile of its population finds itself in. Our framework addresses this shortcoming by using a compositional transformation and leverages the full distributional information in the data, thereby accurately reflecting the more detrimental situation in Scenario 2.

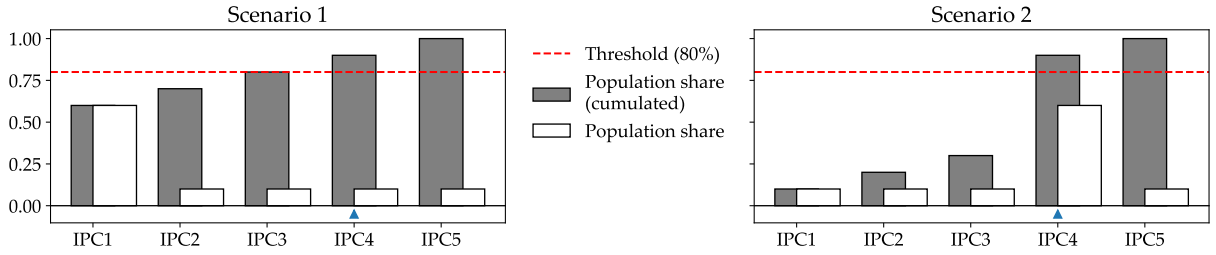

in the respective IPC phases over all districts.[2] The reported IPC values are net of humanitarian assistance impacts as the model is intended to estimate outcomes if urgent action is not taken.

As previously outlined, food insecurity is driven by complex interactions between conflict, poverty, extreme weather, climate change and food price shocks (D'Souza & Jolliffe, 2013; Headey, 2011; Misselhorn, 2005; Singh, 2012). This means that a large number of possible covariates may explain dynamics in the food insecurity distribution.

Agronomic and weather-related variables reflect the agricultural conditions. Historically, crop failures and fluctuations in crop yields have been a key reason for worsening food conditions. We use satellite data to calculate stress indicators and anomalies in vegetation, humidity and soil moisture. To highlight anomalies from historical norms, we also calculate differences with respect to long-term averages for a particular month.

Closely related are prices of staple food items, whose composition may differ between countries. We constructed a staple food price index to capture the market conditions for food commodities. Volatility in food prices are seen as major drivers for food insecurity (IPC Global Partners, 2019). The treatment of the raw data is discussed by (Andrée et al., 2020). We obtain food price data from the FAO and the WFP. We additionally include price data that are seasonally adjusted.

Political instability and conflicts have been other key drivers of famines in recent history, even with ample food supply. For instance, Devereux (2000) argues that famines in Sub-Saharan Africa were the result of the combination of natural disasters, such as droughts, and political triggers, such as civil wars. We gather data on conflict events and the number of associated fatalities from Armed Conflict Location & Event Data Project (ACLED). Due to the sparsity of conflict events, and their possible wider regional impact, we use inverse-distance weighting to proxy the effect of violent outbreaks close to a district of interest.

[2]To illustrate, the formula below shows how the share of people in IPC2 is calculated for a given period.

$$\frac{1}{pop_c} \sum_d pop_{c,d} \cdot \mathbb{I}(IPC_{c,d} = 2)$$

$IPC_{c,d}$ is the IPC phase of district d in country c. $pop_{c,d}$ is the population in district d of country c and $pop_c = \sum_d pop_{c,d}$. $\mathbb{I}(cond)$ is an indicator function, which takes on the value 1 if the condition $cond$ is satisfied and 0 otherwise.

Table 1: IPC Acute Food Insecurity Phase Descriptions

For more details we refer to IPC 3.0 Area Phase Classification (IPC Global Partners, 2019). The left-most column describes the classification used in this article. The rationale being that IPC3+ phases are identified as requiring "urgent action" (IPC Global Partners, 2019) and individual high IPC stages lacking sufficient historical observations.

IPC1	Phase 1	Minimal	Households are able to meet essential food and non-food needs without engaging in atypical and unsustainable strategies to access food and income.
IPC2	Phase 2	Stressed	Households have minimally adequate food consumption but are unable to afford some essential non-food expenditures without engaging in stress-coping strategies.
IPC3+	Phase 3	Crisis	Households either have food consumption gaps that are reflected by high or above-usual acute malnutrition, OR, are marginally able to meet minimum food needs but only by depleting essential livelihood assets or through crisis-coping strategies.
	Phase 4	Emergency	Households either have large food consumption gaps which are reflected in very high acute malnutrition and excess mortality, OR, are able to mitigate large food consumption gaps but only by employing emergency livelihood strategies and asset liquidation.
	Phase 5	Famine	Households have an extreme lack of food and/or other basic needs even after full employment of coping strategies. Starvation, death, destitution, and extremely critical acute malnutrition levels are evident.

Table 2: Exogenous variables and transformations

This table shows the collected exogenous variables in different categories and the transformations we apply (feature engineering). The order of transformations and aggregations follow the vertical structure of the table. Spatial aggregations are applied to the district-level data to compute country-level values. Temporal aggregations are computed on a monthly frequency and we then select the values that are contemporaneous with the reported IPC values.

Category	Argonomic stress and weather	Conflict	Economic
Variable	Rainfall Normalized difference vegetation index (NDVI) Evapotranspiration (ET) Evaporative stress index (ESI)	Number of violent events Number of fatalities	Consumer price index (CPI) Gross domestic product (GDP) Staple food price index
Transformations	Average Anomalies (anom)	Count Inverse-distance weighted (IDW)	
Spatial aggregation	Average over regions Standard deviation over regions		
Temporal aggregation	Percentage changes (%Δ) Rolling average over 3- or 6-months Rolling standard deviation over 3- or 6-months		

These data inputs are compiled for 15 countries on a district level. We then apply several transformations and aggregations to highlight different features of the data set that may be predictive for famine risk dynamics. In particular, we aggregate district data to the country level and monthly data to the reporting frequency of FEWS NET. Table 2 summarizes the raw data inputs and transformations applied. In total we obtain 1,670 candidate variables.

3 Empirical framework

Which set of variables is relevant for a particular country or a specific level of food insecurity is not obvious a priori. One strategy is to rely on experts to select the appropriate set of variables for each case. Not only is this route tedious, but it is also prone to human error and biases. We instead rely on variable selection methods that have been widely employed in the statistical, economic, medical, and machine learning literature. Using these methods, we select the most important drivers of famine risk for each country and each level of famine risk. The selection process is entirely data driven and has the explicit goal of avoiding over-fitting. That is, we aim at selecting variables that are descriptive for historical data but also general enough to extrapolate their effects.

It is important to emphasize that our framework was not to constructed to supersede expert opinion. On the contrary, one of the main strengths of our model is the ability to explicitly and transparently incorporate expert judgements. This is not only necessary due to the shortness of recorded famine history. Augmenting the data with expert knowledge is a solution that major policy institutions have chosen to overcome their data challenges, for instance, when calculating macroeconomic forecasts. Subjective judgements also make it possible to deviate from historical patterns when unprecedented events are looming on the horizon. Events such as pandemic outbreaks, conflict situations, or natural disasters have direct consequences on famine risk. These shocks, immediately obvious to human experts as impactful, would escape a purely data-driven model. Reserving a channel for such opinions is therefore essential for an operational model.

3.1 Notation

Let $X_{ct} = [X_{ct}^1, ..., X_{ct}^k, ..., X_{ct}^K]$ denote a vector containing the population share of country c in K different IPC stages. For example, the Scenario 1 of Figure 2 would be represented as $X_{ct} = [0.6, 0.1, 0.1, 0.1]^\top$. In this study, we consider $K = 3$, which corresponds to IPC phases 1,2 and 3+. IPC3+ is the sum of populations in IPC3,4,5 and represents the population share where urgent action is needed.

3.2 A convenient transformation

While modeling the entire distribution X_{ct} provides a holistic picture, we cannot treat it as a conventional dependent variable in a time-series model. For instance, let $X_t = AX_{t-1} + e_t$ be a "naive" model with some appropriate autocorrelation matrix A. The shortcomings of this model become evident once we consider its predictions $\hat{X}_t$. First, in order to qualify as a distribution, the vector elements have to sum up to one, $\sum_k \hat{X}_t^k = 1$. Second, all elements of $\hat{X}_t$ have to be non-negative. Or in mathematical terms, $\hat{X}_t$ must lie in a $K = 3$ dimensional unit simplex Δ^3. See Figure 3 for a graphical representation of this concept for South Sudan. Neither of these

Figure 3: Food insecurity situation X_t of South Sudan, 2009-2019

This figure traces the food insecurity situation, i.e. the vector X_t, for the example of South Sudan over time. The three corners represent the situations where 100% of the population would be in either IPC1, IPC2, or IPC3+. The cross in the center represents the case where the population is equally divided into IPC1,2 and 3+. The compositional transformation (see Appendix) ensures that the predicted values $\hat{X}_t$ always remain within this triangle.

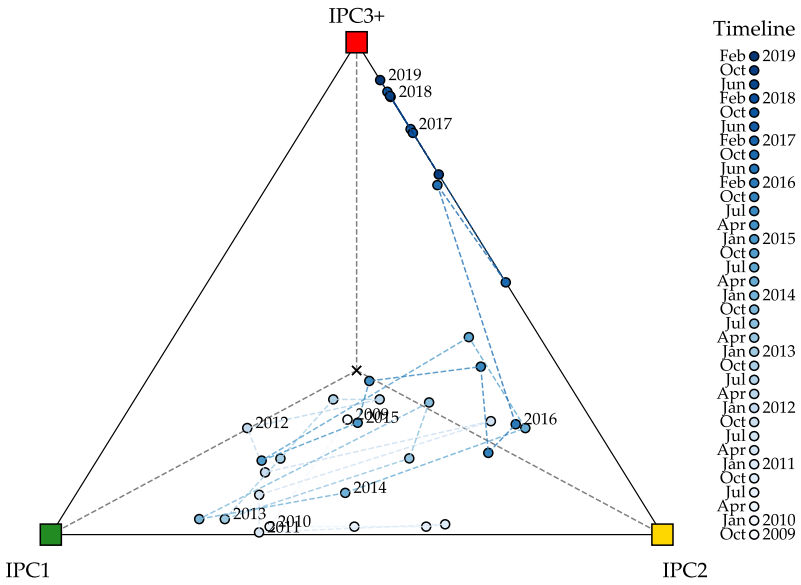

requirements is guaranteed in the naive model.

To ensure that the predicted values $\hat{X}_t$ lie within the triangle of Figure 3, we transform the constrained food insecurity distribution vector X_t into an unconstrained *famine risk vector* $F_t = G(X_t)$. We achieve this by removing one of the redundant dimensions of X_t through a vector-valued compositional transformation $G : \Delta^3 \to \mathbb{R}^2$ (see Appendix).[3] The elements of the 2-dimensional F_t can intuitively be understood as F_t^L for "low" and F_t^H for "high" famine risk states.

Working with an unconstrained F_t greatly facilitates the modeling task. We can now employ conventional multivariate time-series methods on F_t and make predictions $\hat{F}_t$ which are not subject to any constraints as X_t. After obtaining the predictions of interest, e.g. $\hat{F}_{t+1}$, we then retrieve the full food insecurity distribution $\hat{X}_{t+1} = G^{-1}(\hat{F}_{t+1})$, which satisfies all distributional requirements, through the inverse compositional transformation G^{-1}.

3.3 Single-country model

We now describe the modeling approach for a single country, before extending it for multiple countries. In the single country model we suppress the c subscript for clarity. As can be seen in Figure 1, food crises are a rather persistent phenomenon although we also observe strong heterogeneity between countries. The statistical figures in Table 6b corroborate this finding by

[3]Note that we only require $K-1$ elements to determine the full distribution, since $X_t^1 = 1 - \sum_{k=2}^{K} X_t^k$.

highlighting the predictive power of preceding data. This motivates us to specify a dynamic framework, where previous values are indicative for subsequent developments.

$$F_t^L = \mu^1 + a_{LL}F_{t-1}^L + a_{LH}F_{t-1}^H + Z_t^1\beta^1 + u_t^1 \tag{1}$$
$$F_t^H = \mu^2 + a_{HL}F_{t-1}^L + a_{HH}F_{t-1}^H + Z_t^2\beta^2 + u_t^2 \tag{2}$$

or in vector notation

$$F_t = \mu + AF_{t-1} + Z_t\beta + u_t \tag{3}$$

This model is a two-dimensional vector autoregression with exogenous variables (VARX). The autocorrelation matrix A is 2×2. It is not diagonal, i.e. $A_{LH} \neq 0, A_{HL} \neq 0$, which reflects the possibility for low famine risk states to transition to high famine risk states, and vice versa. The errors are bivariate normal $u_t \sim N(0, \Sigma)$. All exogenous variables are contained in Z_t and as equations (1) and (2) show, the set of relevant variables may differ between high and low food insecurity states, $Z_t^1 \neq Z_t^2$. Consequently, their coefficients can differ as well, $\beta^1 \neq \beta^2$.

Ideally, we would be able to estimate the single-country model to tailor its prediction for each country. As mentioned earlier, this is not feasible due to limitations imposed by data availability. In Section 6 we will demonstrate how single-country models become feasible once we incorporate expert opinions in a Bayesian framework.

3.4 Multi-country model

The single-country model conveys the dynamic framework we stipulate for each country. However, given the short time series of 34 periods and our interest of famine risk patterns across countries, we specify a panel extension, i.e. a panel vector autoregression with exogenous variables (PVARX).

$$F_{ct}^L = \mu^1 + a_{LL}F_{ct-1}^L + a_{LH}F_{ct-1}^H + Z_t^1\beta_c^1 + u_{ct}^1 \tag{4}$$
$$F_{ct}^H = \mu^2 + a_{HL}F_{ct-1}^L + a_{HH}F_{ct-1}^H + Z_t^2\beta_c^2 + u_{ct}^2 \tag{5}$$

or in vector notation,

$$\boldsymbol{F}_t = \mu_N + \mathbf{A}_N\,\boldsymbol{F}_{t-1} + Z_tB + U_t \tag{6}$$

where $\mu_N = \iota_N \otimes \mu$ with $\otimes$ denoting the the Kronecker product and ι_N is a N-dimensional vector of ones. Moreover, $\mathbf{A}_N = \iota_N \otimes A$, which implies the same autocorrelation matrix A across all countries. The bold vector $\mathbf{F}_t$ denotes a stacked vector containing each countries famine risk state vector F_{ct} for $c = 1, ..., N$. Correspondingly, the disturbances U_t follow a multivariate normal with block-diagonal covariance matrix $I_N \otimes \Sigma$.

4 Variable selection

Based on the transformations and aggregations described in Table 2 we constructed a total of 1,670 candidate variables. Including all of them would certainly lead to over-fitting and low external

validity of the results. While the literature and the IPC framework itself suggest certain variable types to be included, their exact transformations remain less obvious. We therefore opt for a statistical variable selection method, also called feature selection procedure. This is usually done using either L1-regularization or LASSO regression (Tibshirani, 1996), L2-regularization or Ridge regression (Hoerl & Kennard, 1970), or a combination thereof, such as the Elastic-Net regression (Zou & Hastie, 2005), or the adaptive LASSO (Zou, 2006).

These methods all have the common goal of reducing model complexity, but differ in what is considered complex. The LASSO regression aims at selecting as few parameters is possible while the Ridge regression shrinks the parameters towards zero as much as possible. The benefit of the LASSO lies in the parsimonious nature of the resulting model and greatly facilitates interpretation of the few resulting parameters. However, this comes at the cost of possibly discarding variables that still may be relevant (Zou, 2006). The Ridge regression, in contrast, avoids this type of over-discarding variables as it rarely excludes variables completely. This means the set of active variables remains stable, but also very large which makes interpretation more ambiguous. We choose the LASSO regression since our goal is to better understand the drivers behind famine risk dynamics.

A key technical consideration is the choice of the regularization parameter. This is usually done either with cross-validation or information criteria. However, the panel setup of our model is not commonly dealt with in the machine learning or statistical literature. The panel structure implies additional dependency structures compared to simple cross-sectional models that are often employed to select genomic markers or imagery features. The additional time dimension induces temporal dependency which demands special attention (Bergmeir & Benítez, 2012). Arguably, our validation also has to account for a spatial dependency component between countries (Roberts et al., 2017). For instance, areas close to each other are likely to share similar agronomic conditions and weather variables. This consideration may be relevant for district level analyses. On the country level this is less relevant as many spatial dependencies are aggregated out. Moreover, the 15 countries considered are rather dispersed and do not share many common borders.

It is therefore problematic to rely on conventional best practices for cross-validation selection (Zhang & Yang, 2015). The results may be strikingly different, depending on the method chosen. Bergmeir et al. (2018) advocate the use K-fold cross validation for time series problems. However, their argument for the validity of these approaches relies on a correct-specification argument which can be met more easily in their non-parametric context, but becomes unrealistic in our application that uses only a modest number of parameters to model the data. Moreover, our application is particularly interested in introducing expert information through the use of priors. Such methods introduce biases that help improve model performance for certain events, but persist in the limit and can thus not be reconciled with a correct-specification argument in the parametric context (Andrée, 2020; Blasques & Duplinskiy, 2018). We therefore opt for a group K-fold cross-validation which splits the sample into train and testing folds while accounting for the panel structure.

4.1 Selection results

The Tables 3 and 4 show the selected variables after employing a LASSO regression for (4) and (5), separately. In the selection procedure we excluded the autoregressive components and the constant, meaning we always select them. The results are grouped following the same logic as they were constructed (Table 2) and are ordered across the columns according to the number of monthly lags. In each cell, the vertical bar separates the (refitted) coefficients from the relative importance of the variables.[4] We do not report standard errors due to the inferential issues following a regularized regression (Berk et al., 2013; Lee et al., 2016). For low famine risk states, we select 29 variables and 31 variables for high famine risk states. This number is stable across randomized subsets of candidate variables.

We make several interesting observations. In both famine risk states we find that the autoregressive variables are among the top five most important features. The weakly positive serial correlations support our time-dependent model architecture. The "Model" components account for about a third of overall variable importance in both cases. We also find evidence that the cross-dependency is asymmetric: Low famine risk states are more likely to transition to high famine risk states ($\hat{a}_{HL} = 0.201$) than vice versa ($\hat{a}_{LH} = 0.097$).

For low famine risk states, in Table 3, conflict events tend to play a minor role with 10.74% overall importance, while (lagged) agronomic stress indicators account for 38.28%. Interestingly, the two- and six-month lags tend to be most predictive. Economic variables contain two of the top five most relevant variables, namely average returns on staple food prices and the GDP dispersion across districts.

A different pattern emerges for high famine risk states, in Table 4. Conflict variables gain in importance (13.81%) and contain two of the five most relevant variables. For both conflict and economic variables, the most recent values are also most informative. Weather and agronomic variables lose in relevance but still account for a third of overall importance and have a dispersed lag structure.

The plots of Figure 8 show the residuals of $\hat{F}_t^L, \hat{F}_t^H$ after estimating the PVARX with the selected variables. We can roughly discern a bi-variate normal distribution from the contours of the scatter plot. The marginal plots show that while the distributions are mostly symmetric, they exhibit excess kurtosis compared to the standard normal distribution. The higher probability mass around the mean and fatter tails indicate that Student t-distributions may be a better choice. Nonetheless, we choose the bi-variate normal distribution for subsequent simulation analyses as it appears to be a reasonable approximation.

[4]The relative importance is calculated following the Shapely value approach as discussed in Lundberg and Lee (2017).

Table 3: Selected variables and their importances for low famine risk states F^L

We present the selected variables for low famine risk states from a panel LASSO for model (4). In each cell, a "|" separates the refitted LASSO coefficients and their relative importances. The five most important variables are ranked using roman numerals [I]-[V]. Negative coefficient signs indicate worsening conditions, i.e. population shares in IPC1 transition into IPC2 (see Figure 9, Appendix). The results are grouped by variable categories, spatial- and temporal aggregations over monthly lags. In total, 29 variables were selected from 1670 candidate variables (Table 2). The dashes (-) are not selected. The autoregressive variables F^L_{t-1}, F^H_{t-1} and the constant are always selected. We use grouped 10-fold cross-validation to choose the regularization parameter.

	Aggregation		Monthly lags						
Variable	Spatial	Temporal	0	1	2	3	4	5	6
Agronomic stress and weather (cuml. importance 38.28%)									
ESI (%Δ)	std	-	-	-	-	0.000 | 0.79%	-	-	-
		avg_{3M}	-	-	0.000 | 0.69%	-	-	-	-
ET	avg	std_{3M}	-	-	0.009 | 0.81%	-	-	-	-
	std	avg_{3M}	-	-	-	-	-	-	**[III]** -0.045 | 5.38%
		avg_{6M}	-	-	-	-	-	-	-0.027 | 2.89%
ET (anom,%Δ)	std	-	-	-	-	-	0.001 | 1.48%	-	-
NDVI	avg	std_{3M}	-	-	0.743 | 0.92%	-	-	-	-
NDVI (%Δ)	avg	-	-	-	0.658 | 4.62%	-	-	-	-
NDVI (anom)	avg	std_{3M}	-0.013 | 1.54%	-	-	-	-	-	-
NDVI (anom,%Δ)	std	-	-0.022 | 0.57%	-	-	-	-	-	-
Rainfall	avg	avg_{3M}	-	-	-0.004 | 1.91%	-	-	-	-
		std_{6M}	-	-	0.017 | 3.77%	-	-	-	-0.008 | 1.91%
	std	avg_{3M}	-	-	-0.008 | 1.56%	-	-	-	-
		avg_{6M}	-	-	-	-	-	-	0.006 | 0.83%
Rainfall (anom)	avg	std_{6M}	-	-	-	-	-	-	-0.032 | 3.75%
	std	-	-	-	-	-	-	0.034 | 4.05%	-
Rainfall (anom,%Δ)	std	-	-	-	-	-	-	-	0.000 | 0.82%
Conflict (cuml. importance 10.74%)									
Events	std	-	-	-0.064 | 5.02%	-	-	-	-	-
Events (IDW)	std	-	0.011 | 0.74%	-	-	-	-	-	-
Fatalities (IDW)	avg	-	-	-1.335 | 2.17%	-	-	-	-	-
Fatalities (%Δ)	avg	avg_{3M}	-	-	-	-	-0.365 | 2.81%	-	-
Economic (cuml. importance 18.89%)									
Food price	avg	std_{3M}	2.111 | 1.69%	-	-	-	-	-	-
Food price (SA,%Δ)	std	-	-	-	4.388 | 2.82%	-	-	-	-
Food price (%Δ)	avg	-	**[V]** -4.089 | 5.11%	-2.161 | 2.48%	-	-	-	-	-
		avg_{3M}	-	-	-	-	-	0.519 | 0.51%	-
		avg_{6M}	-	-	-	-	0.830 | 0.62%	-	-
GDP	std	-	-	-	-	-	**[II]** 0.000 | 5.66%	-	-
Model (cuml. importance 32.08%)									
F^H_{t-1}			**[IV]** 0.097 | 5.18%	-	-	-	-	-	-
F^L_{t-1}			**[I]** 0.536 | 26.90%	-	-	-	-	-	-
Constant			0.264 | 0.00%	-	-	-	-	-	-

Abbreviations. ET=Evapotranspiration, CPI=Consumer Price Index, GDP=Gross domestic product, ESI=Evaporative Stress Index, NDVI=Normalized Difference Vegetation Index. IDW=Inverse-distance weighted interpolation of conflict events. SA=Seasonally-adjusted food price index. "anom" are anomalies to the mean. The percentage (%Δ) indicates monthly percentage changes. In the spatial aggregation column, "avg" and "std" are the average and standard deviation over all admin-1 districts to aggregate data to the country level (admin-0). Similarly in the temporal aggregation column, where subscripts reflect the rolling aggregation window.

Table 4: Selected variables and their importances for high famine risk states F^H

We present the selected variables for high famine risk states from a panel LASSO for model (4). In each cell, a "|" separates the refitted LASSO coefficients and their relative importances. The five most important variables are ranked using roman numerals [I]-[V]. Negative coefficient signs indicate worsening conditions, i.e. population shares in IPC1,2 transition into IPC3+ (see Figure 9, Appendix). The results are grouped by variable category, spatial- and temporal aggregations over monthly lags. In total, 31 variables were selected from 1670 candidate variables (Table 2). The dashes (-) are not selected. The autoregressive variables F^L_{t-1}, F^H_{t-1} and the constant are always selected, omitted here. We use grouped 10-fold cross-validation to choose the regularization parameter.

	Aggregation		Monthly lags						
Variable	Spatial	Temporal	0	1	2	3	4	5	6
Agronomic stress and weather (cuml. importance 33.21%)									
ET (anom)	avg	-	-	-	0.038 | 2.70%	-	-	-	-
		avg_{3M}	-	-	0.010 | 0.59%	-	-	-	-
	std	avg_{6M}	-	-	-	-	-	-	-0.107 | 3.19%
NDVI	avg	std_{6M}	-	-	-	-	-	-	-1.341 | 2.22%
NDVI (%Δ)	avg	-	-	-	-	-	-	0.482 | 3.50%	-
NDVI (anom)	avg	avg_{6M}	0.006 | 1.25%	-	-	-	-	-	-
		std_{6M}	-	-	-	0.015 | 2.09%	-	-	-
NDVI (anom,%Δ)	avg	avg_{3M}	-	-	-	-0.010 | 0.03%	-	-	-
	std	-	-	-	-	-	0.020 | 0.54%	-	-
Rainfall	avg	avg_{3M}	-	-	**[III]** 0.013 | 7.09%	-	-	-	-
		std_{3M}	-	-	-	-0.001 | 0.15%	-	-	-
		std_{6M}	-	-	0.013 | 3.44%	-	-	-	-
Rainfall (%Δ)	std	avg_{3M}	-	-0.000 | 0.93%	-	-	-	0.000 | 0.51%	-
		avg_{6M}	-	-	-	-	-	0.000 | 0.89%	-
Rainfall (anom)	avg	avg_{6M}	-	-	-0.010 | 0.91%	-	0.034 | 3.18%	-	-
Conflict (cuml. importance 13.81%)									
Events	avg	std_{3M}	-0.214 | 2.22%	-	-	-	-	-	-
Fatalities (IDW)	std	avg_{3M}	-1.377 | 2.71%	-	-	-	-	-	-
Fatalities (%Δ)	avg	avg_{3M}	**[IV]** -0.481 | 4.32%	-	-	-	-	-	-
		std_{6M}	-	-	-	-	-	**[V]** -0.702 | 3.96%	-
	std	-	-0.162 | 0.60%	-	-	-	-	-	-
Economic (cuml. importance 14.18%)									
CPI	avg	std_{3M}	-5.257 | 2.26%	-	-	-	-3.964 | 1.54%	-	-
		std_{6M}	2.062 | 1.19%	-	-	-	-	-	-
CPI (%Δ)	avg	avg_{3M}	-0.313 | 0.20%	-	-	-	-	-	-
		avg_{6M}	-2.247 | 1.24%	-	-	-	-	-	-
		std_{3M}	-	-	-	-	-2.326 | 1.84%	-	-
		std_{6M}	-	-	-	-	0.115 | 0.10%	-	-
Food price (%Δ)	avg	-	-1.553 | 2.19%	-2.490 | 3.22%	-	-	-	-	-
		avg_{3M}	-0.384 | 0.42%	-	-	-	-	-	-
Model (cuml. importance 38.78%)									
F^H_{t-1}			**[I]** 0.455 | 27.43%	-	-	-	-	-	-
F^L_{t-1}			**[II]** 0.201 | 11.36%	-	-	-	-	-	-
Constant			0.248 | 0.00%	-	-	-	-	-	-

Abbreviations. ET=Evapotranspiration, CPI=Consumer Price Index, GDP=Gross domestic product, ESI=Evaporative Stress Index, NDVI=Normalized Difference Vegetation Index. IDW=Inverse-distance weighted interpolation of conflict events. SA=Seasonally-adjusted food price index. "anom" are anomalies to the mean. The percentage (%Δ) indicates monthly percentage changes. In the spatial aggregation column, "avg" and "std" are the average and standard deviation over all admin-1 districts to aggregate data to the country level (admin-0). Similarly in the temporal aggregation column, where subscripts reflect the rolling aggregation window.

Figure 4: One-step ahead forecasting scheme

We use the following setup to make one-head ahead forecasts, in order to reflect the information available for decision makers at each point in time. The black lines and dots represent the observed data based on which the model is estimated. The red squares represent the one-step ahead forecasts. The estimation period is at least 15 time observations long which is about half of the total available data.

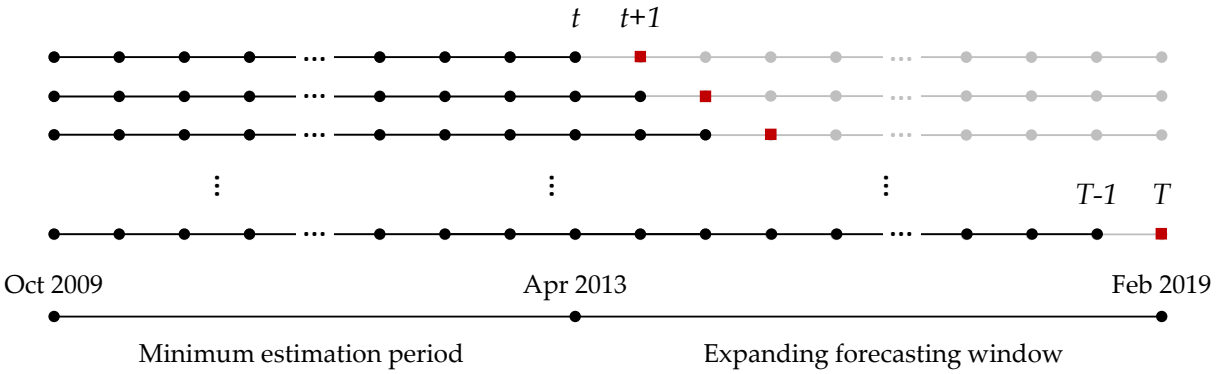

5 Forecasting

Using the selected variables and estimated model parameters we now demonstrate the model's prediction performance. At each point in time we calculate the one-step ahead prediction and compare it to the famine risk distribution of the next FEWS NET release. This imitates the situation of policy makers, who decide based on the available information. Let $\mathbb{I}_T$ denote all information since the beginning of recorded famine history in October 2009 until the hypothetical decision date T. Then the one-step ahead prediction is defined as

$$X_{T+1|T} := \mathbb{E}[X_{T+1}|\mathbb{I}_T]$$

In order to avoid incorporating future information in the one-step ahead predictions, which would not have been available, we estimate the model only using the information until T and then make a forecast for $X_{T+1|T}$. We then re-estimate the model with information until $T+1$ and make a forecast for $X_{T+2|T+1}$. The minimal observation period is at least $T = 15$ such that there is sufficient information to fit the model. We note, however, that the set of selected variables in Tables 3 and 4 are taken as given. Unquestionably, the selection relies on the full sample and therefore on future information. Nonetheless, we decided against re-selecting the variables at each point in time to avoid model instability.

The results of the rolling forecasts using 1,000 hypothetical draws are shown in Figure 5 and associated accuracy in Table 5. We can see that the model is more accurate for higher IPC phases for almost all countries. The predictions are within the 50% confidence interval in the majority of cases. Lower IPC stages are associated with lower precision and broader confidence intervals. Interestingly, the model performs well in terms of predicting the correct direction of IPC3+ dynamics, for instance, in Afghanistan, Chad, Malawi, Mozambique and Zimbabwe.

Table 5: One-step ahead forecast accuracy

We present the root-mean-square errors (RMSE) of the one-step ahead forecasts for each country and IPC phase. The last row (column) show the average over IPC phases (countries).

IPC Phase	Afghanistan	Chad	Ethiopia	Guatemala	Haiti	Kenya	Malawi	Mali	Mozambique	Niger	Nigeria	Somalia	South Sudan	Uganda	Zimbabwe	Average
IPC1	0.307	0.294	0.188	0.247	0.256	0.203	0.269	0.238	0.229	0.321	0.366	0.241	0.258	0.278	0.224	0.261
IPC2	0.274	0.191	0.120	0.191	0.179	0.110	0.241	0.169	0.180	0.215	0.227	0.348	0.174	0.204	0.179	0.200
IPC3+	0.107	0.125	0.123	0.073	0.108	0.109	0.249	0.079	0.085	0.126	0.147	0.188	0.258	0.079	0.198	0.137
Average	0.229	0.203	0.144	0.170	0.181	0.141	0.253	0.162	0.165	0.221	0.246	0.259	0.230	0.187	0.200	0.199

Figure 5: One-step ahead forecasts

The forecasts below are produced using the scheme in Figure 4. The black lines with round markers are the observed values. The colored squares are the average one-step ahead forecasts from 1,000 simulations. The 90% and 50% (asymmetric) confidence intervals are based on the same draws.

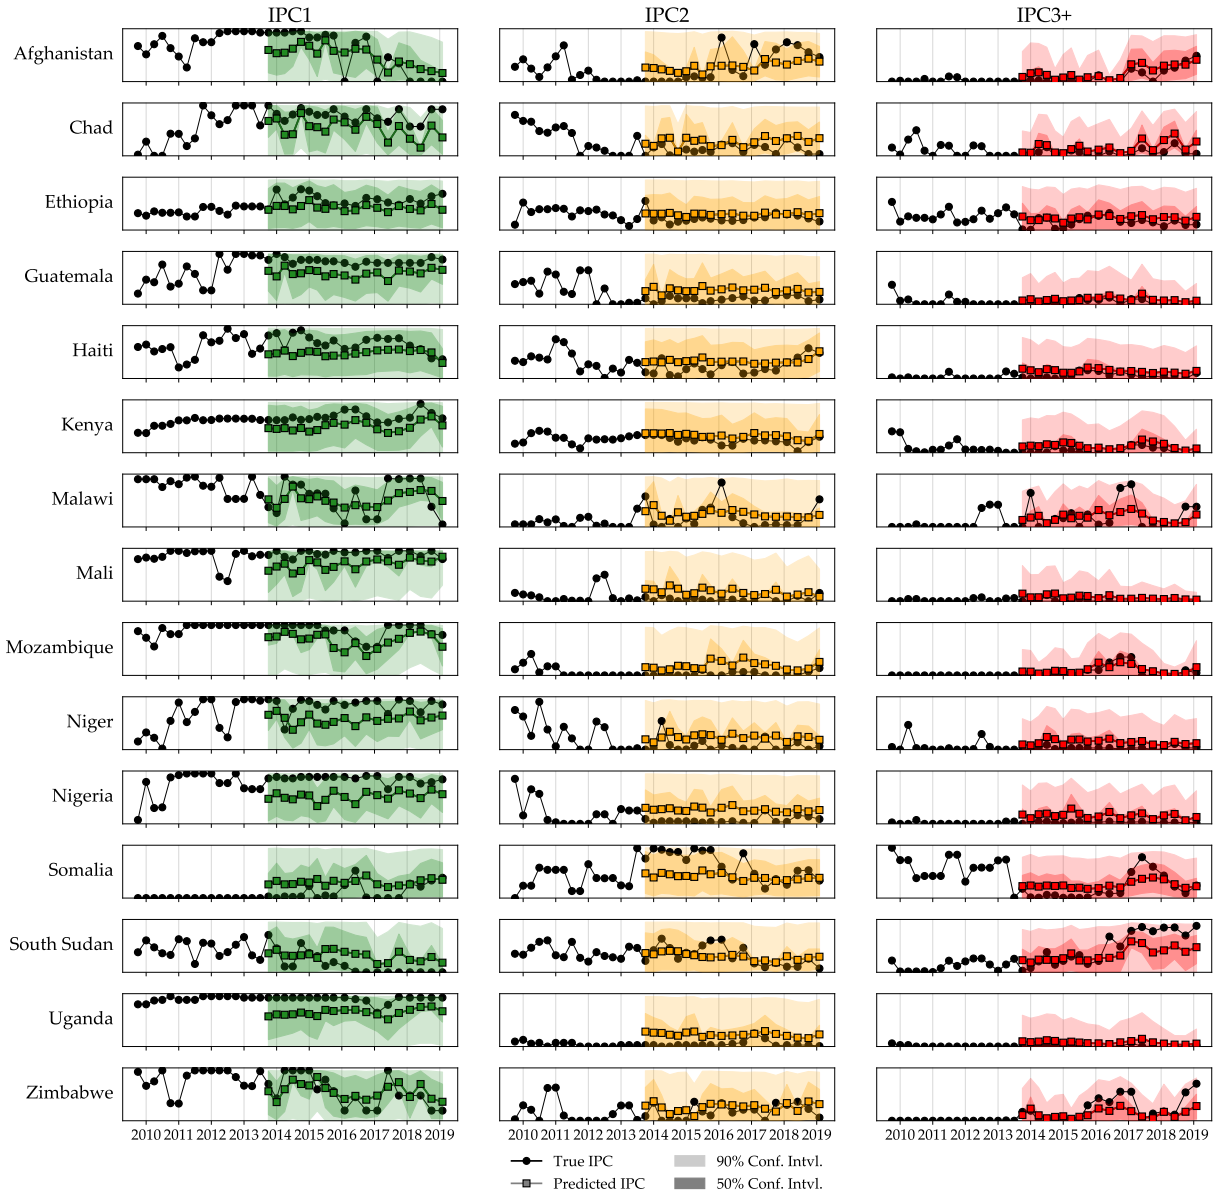

6 Expert opinion

One of the main challenges in this exercise lies in the scarcity of data. The short time series imposes difficulties for the identification of model parameters, in particular for the autoregression matrix A in (3). Several remedy options are possible. First, one can impose additional model structure, such as moving averages, additional lags, nonlinear dependencies or fat-tailed distributions. However, this quickly leads to over-parameterization and comes at the cost of tractability. Second, one can engineer additional features and expand the set of exogenous variables. While there is merit to this approach, we believe that the current set of 1,670 explanatory variables is sufficiently complete. Furthermore, more regressors only help marginally in identifying the autocorrelation matrix A. We follow a third approach involving expert opinion.

A similar problem has plagued policy makers and regulatory institutions when making macroeconomic forecasts. GDP figures and other macroeconomic indicators are usually available only on annual or quarterly frequencies. Even in advanced economies the lack of data presents a major bottleneck for empirical analyses. Major policy institutions such as the International Monetary Fund (Ciccarelli & Rebucci, 2003), the Federal Reserve (Carriero et al., 2011) or the European Central Bank (Bańbura et al., 2008) have dealt with this problem by augmenting their forecasting models with expert opinion (Giannone et al., 2019; Litterman, 1986; Sims & Zha, 1998). Their models are comparable to our VARX model in (3).

Expert opinion can be used to make direct forecasts through surveys and expert panels about the development or trend in subsequent periods. Yet, decision-makers prefer to understand the empirical evidence leading up to a certain prediction or outcome. Hence, we advocate a framework where expert opinion is not a substitute but a complement to the statistical model. Similar to the approach of the aforementioned policy institutions, we inject the subjective judgments into the model itself. This is different from a purely data-driven approach which starts off agnostic and estimates its parameters entirely from the data. While this can be argued to be the most objective way to conduct empirical investigations, we also run the risk of over-fitting to the short data set and are effectively handicapping ourselves. Specifically, certain country groups are structurally different from others, for instance, South Sudan, Somalia or the Republic of Yemen are known to be fragile while countries like Uganda or Mali have more stable food security situations (Figure 1). Moreover, low food insecurity stages tend to be persistent while high food insecurity stages tend to be more volatile (Table 6a and 6b). This information can ideally be learned by the model itself but requires sufficient historical cases to learn from. Not only is this condition not met in our case, it is also important to have representative and relevant data: While the history of food scarcity is long, Devereux (2000, 2006) makes a convincing case that food insecurity in the 21st century is of a different nature than its predecessors.

Adopting a Bayesian view is fully consistent with our stochastic framework and its targeted end use. One might even argue that it is more suitable than the classical view since parameters are inherently stochastic, not only due to estimation uncertainty. This allows us to construct stochastic catalogs that contain expert opinion without losing any of the attractive distributional properties. The priors also alleviate limitations imposed by the short data history, rendering single-country models feasible.

Concretely, instead of letting the model learn everything from the data, we "warm start" it with prior expert beliefs. For example, we tell the model a priori that low IPC states tend to remain in low IPC states by setting a^0_{LL} close to 1.0 in equation (3).[5] The exact value for a^0_{LL} can certainly be tailored to specific countries. Most importantly, these beliefs are not calibrations. The key principle in Bayesian statistics is that prior knowledge is informative but should leave sufficient room for data to influence the posterior estimate of $\hat{a}_{LL}$. The result of this Bayesian updating is therefore an amalgamation between prior beliefs and empirical observations. How much we allow the model to learn from the data, or how much the posterior can possibly differ from the prior beliefs, depends on how confident we are about our beliefs. If we were to impose a^0_{LL} with almost total confidence, there will be little room for the data to move the posterior away from the prior. In contrast, if we were to impose a^0_{LL} with close to no confidence, we would effectively revert back to a "cold started" model.

6.1 Improving forecasts with expert opinion: Afghanistan

Country profiles on Afghanistan from various organizations, such as the FAO, WFP or Office for the Coordination of Humanitarian Affairs (OCHA) highlight the growing threat of urban poverty and rapidly increasing food insecurity levels, coupled with persistent conflict and regional inequalities (FAO, 2012). This causes Afghanistan to be one of the most relevant cases while also being one of the most challenging ones due to the few crisis situations with which to inform the model. On these grounds, publicly released assessments by humanitarian organizations are highly relevant for "warm-starting" the model. In this section we illustrate the Bayesian approach for Afghanistan, summarized in Figure 6.[6] We estimate an entirely data-driven model (3) before including weak and strong prior beliefs on the persistence of food insecurity developments.[7] For details on the statistical theory and technical implementation we refer to Bańbura et al. (2008) and Christoffel et al. (2011).

Case 1 (no prior) In Figure 6a we show the posterior distributions of the autocorrelation matrix $\hat{A}^{MLE}$ and of the constant $\hat{\mu}^{MLE}$ with 1,000 draws. The blue triangle (▼) locates the maximum likelihood estimate (MLE) that only relies on observed data. We can see that $\hat{a}^{MLE}_{HH}$ is close to zero, implying strong mean reversion and no persistence. Consequently, the forecasted IPC3+ phases do not continue on the previous trajectory. In contrast, a^{MLE}_{LL} is very high, resulting in IPC1,2 phases to largely continue their previous paths.

Case 2 (weak prior) Figure 6b includes prior information $a^0_{LL} = 1$ and $a^0_{HH} = 1$ with weak confidence. The location of the prior is demarcated with the red triangles (▼). The posterior distributions of $\hat{a}_{LL}$ have means that are always located between prior and a^0_{LL} and a^{MLE}_{LL}. The weak prior has little impact on $\hat{a}_{LL}$, since its MLE was already at 0.9. Conversely, we can see that a weak prior is already informative for $\hat{a}_{HH}$, since the MLE was rather uninformative. As a result, the forecasted values for all IPC phases tend to remain close to their previous trajectories.

[5] We use the zero superscript, A^0, for prior means, the hat, $\hat{A}$, for posterior means, and the MLE superscript, A^{MLE}, for the maximum likelihood estimate without any prior information.

[6] The results in Figure 6 are from the single-country model, equation (3), and are therefore not directly comparable with the results for Afghanistan in Figure 5, which are based on the multi-country model, equation (6).

[7] We are implementing a version of the Minnesota prior (Litterman, 1986; Sims & Zha, 1998).

Case 3 (strong prior) Figure 6c impose the same prior as before, but with higher confidence. In this case, the MLE values have only little influence on the posterior estimate. As a result we are effectively imposing our belief that the situation will almost certainly continue its recent trend. Clearly, this case is not far away from model calibration and should only be reserved for extreme, unprecedented circumstances with agreement among experts.

Note that the priors used here only reflect lessons learned by the authors throughout the modeling process. In a real-world application, these priors should be informed by field experts, researchers and specialists through panels and surveys. Obtaining such priors and implementing them is outside the scope of this study.

6.2 Exceedance probability (EP) curves

For financial applications, particularly for risk management and cost-benefit analyses, EP curves allow us to derive useful risk metrics. We leverage our stochastic framework to calculate the probability of exceeding specific population (share) thresholds under distress. Figure 7 illustrates this concept for total population and population share in IPC3+ in each country. It shows the results of 10,000 simulations from the stochastic model for 2020, 2021 and 2022.

We can see how EP curves help us distinguish between risk profiles across countries and forecast horizons. For instance, in the short-term, Zimbabwe is at highest risk among all countries considered. At the other end of the spectrum, Uganda and Kenya are comparatively unaffected by food insecurity risks. For longer-term predictions, these extreme curves tend to vanish. This is due to the uncertainty surrounding the dynamic predictions, which increases with the forecast horizon.

Figure 6: Effect of prior beliefs on forecasting in Afghanistan

We illustrate the effect of imposing prior information with increasing degrees of confidence across the three panels. In each panel, the first row depicts the coefficients a_{LL}, a_{LH} and the constant μ^1 from equation (1), the second row, analogously, the coefficients a_{HL}, a_{HH} and the constant μ^2 of equation (2). The blue triangles locate the OLS estimators, the red triangles the priors and the black vertical bars show the posterior estimates. The third row shows the associated forecasting performance of each IPC component, where we use observations after February 2018 as the holdout sample. The light and dark shaded areas demarcate the 95% and 50% confidence bands.

(a) No prior

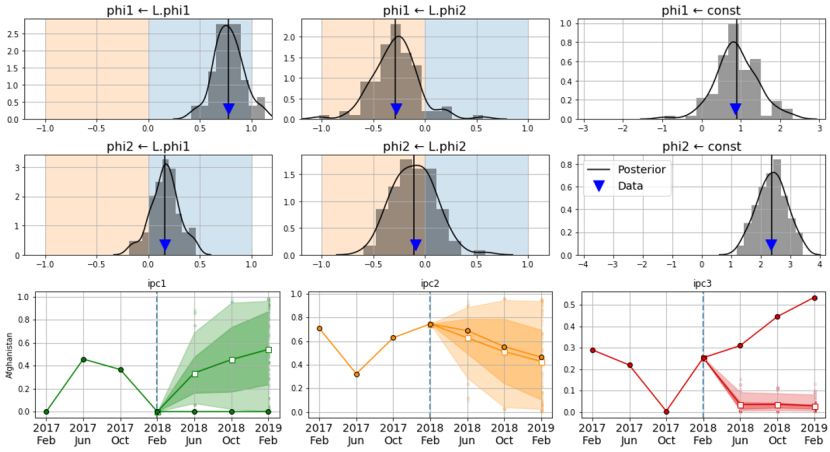

(b) Weak prior for persistency

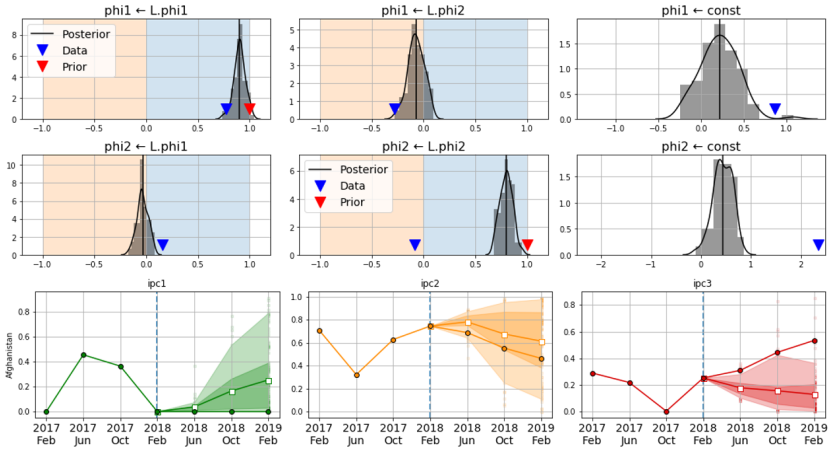

(c) Strong prior for persistency

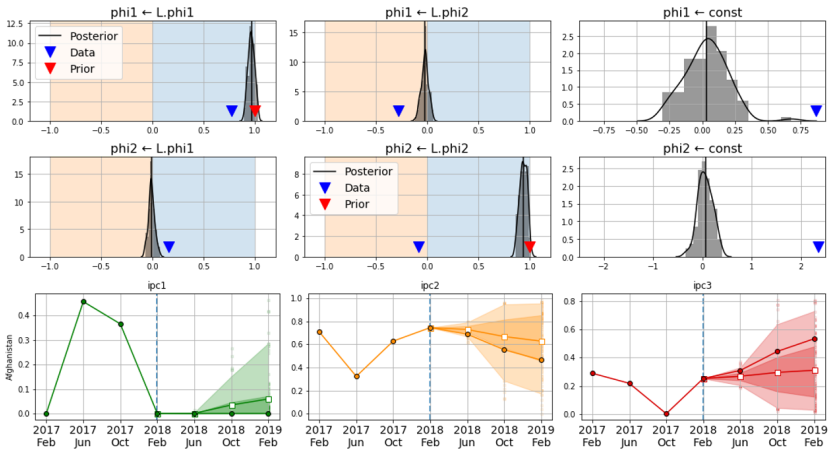

Figure 7: Exceedance probability (EP) curves

This figure depicts the outcome of 10,000 simulations based on the stochastic model. It relates the probability of exceeding (x-axis) a specific population or population share in distress (y-axis), that is in IPC3 or higher. To illustrate, the left end of the x-axis shows that there is a 100% probability that zero or more people are affected while the right end shows that with 0% probability the entire population is affected. The underlying values are based on dynamic forecasts starting after Feb 2019, with 12-, 24- and 36-month horizons, respectively.

(a) Population

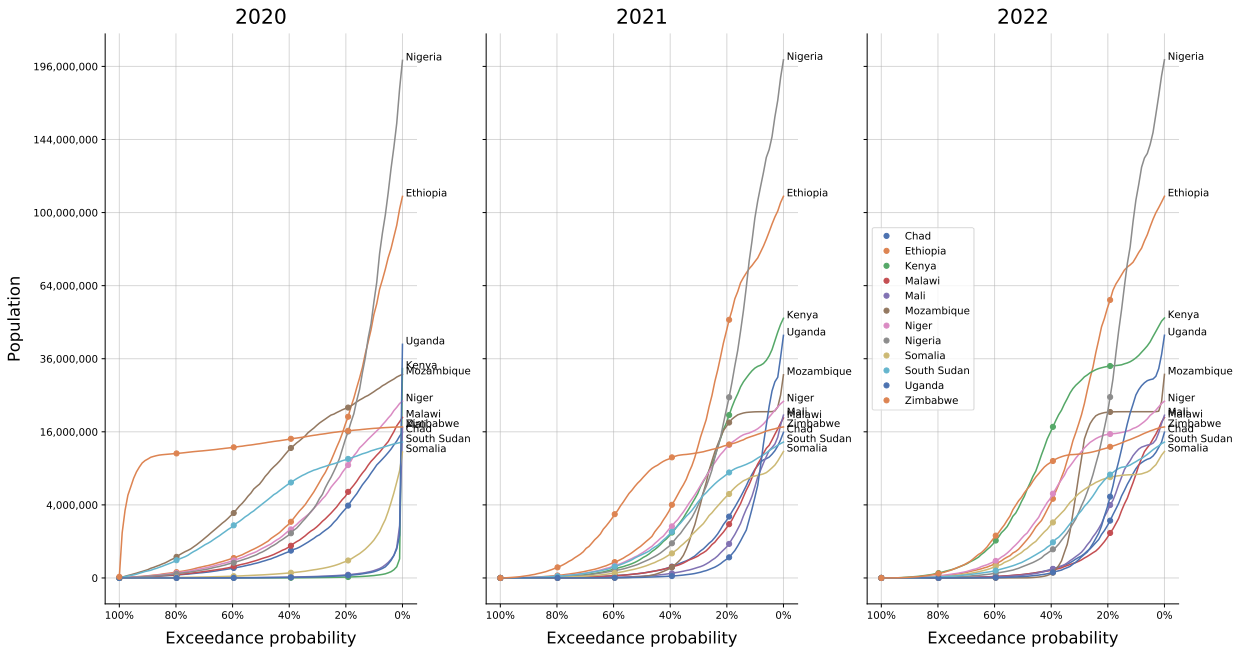

(b) Population share

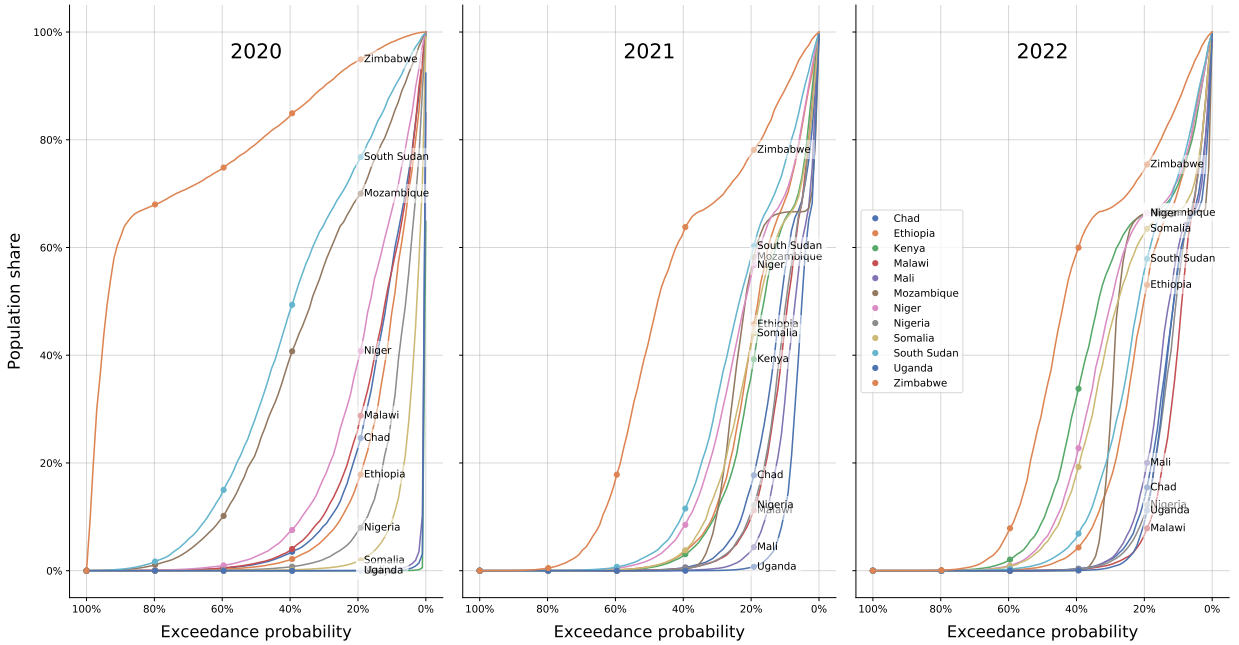

7 Conclusion

In this paper we presented a fully stochastic framework to model and predict food crisis dynamics. We build on the recent advancements in food security analysis and measurement, most notably the IPC framework. A key novelty in our approach is to leverage food insecurity information of an entire population distribution, rather than calculating a single food insecurity metric. Consequently, the model provides a holistic picture of food crisis risk to decision-makers. It also lends itself to construct stochastic catalogues and exceedance probability curves that can serve as hazard modules for insurance and risk management products.

In order to model entire food insecurity distributions with conventional statistical methods, we introduce a convenient transformation, namely the compositional transformation. With it we can fit our model to the data, identify drivers of food insecurity, simulate alternative scenarios, and make forecasts into the near future. Due to the brevity of recorded food insecurity's historic data, it is difficult to estimate a separate model for each country. We therefore employ a panel approach where we pool the data of 15 countries and estimate a joint multi-country model.

Food insecurity research has shown that agronomic, weather, conflict and economic variables are major drivers for food crises. However, which exact variables or transformations thereof should be included to predict either high or low states of food insecurity is not immediately clear. We therefore construct a total set of 1,670 candidate variables for each of the above categories. Using a statistical variable selection method, we identify the 30 most important drivers. We find that food insecurity states are indeed past-dependent and that states of low food insecurity are more likely to transition to high states than vice versa. Furthermore, conflict variables in the recent past are more predictive for high food insecurity levels while agricultural and weather-related variables are more important for lower levels. Food prices are predictive for both cases.

Finally, our model was designed to complement expert opinion. It leaves ample room to incorporate subjective judgements and tailor the model to specific applications. This is essential for unprecedented situations that are not part of the recorded history, which makes them difficult for the model to anticipate. At the same time, consequences of natural disasters or civil wars are obvious to human observers. Using a Bayesian extension, we demonstrate how prior information can be incorporated into our model and how it can significantly improve model performance. This extension is particularly advantageous in a data scarce environment.

This framework may be a useful addition to the analytical toolkit of policymakers and humanitarian organizations as well as financial institutions. The stochastic nature could be particularly attractive for financial triggering decisions, scenario analyses or other risk assessments. While it is well-suited for short- to medium-term predictions, application-specific adjustments and extensions are necessary for forecasts further into the future. We especially advise to pay careful attention to the variable selection step.

8 Appendix

Table 6: Temporal dependency of food insecurity

The subtables (a) and (b) examine the autoregressive character of X_{ct} and F_{ct}, respectively, for each country. In the first column group, we present the autocorrelation coefficient $\hat{\rho}$ of a AR(1) regression $y_t = \mu + \rho y_{t-1} + e_t$, where $y_t \in \{X_{ct}, F_{ct}\}$, with the corresponding p-values. The second column group shows the test statistic and MacKinnon approximate p-value for the augmented Dickey-Fuller test for unit roots.
We can see that while unit roots are rarely present, the autoregressive character is statistically highly significant. In particular, the significance for ρ is highest for low IPC states and progressively decreases with deteriorating levels of food insecurity. We note that the results are obtained for a short time series of 34 periods.

(a) IPC distributions X_{ct}

	AR(1) regression						Augmented Dickey-Fuller					
	IPC1		IPC2		IPC3+		IPC1		IPC2		IPC3+	
Country	coef.	p-val	coef.	p-val	coef.	p-val	stat.	p-val	stat.	p-val	stat.	p-val
Afghanistan	0.660	• 0.000	0.529	• 0.002	0.945	• 0.000	-1.254	0.650	-3.097	• 0.027	-0.369	0.915
Chad	0.626	• 0.000	0.635	• 0.000	0.257	0.141	-3.063	• 0.029	-2.697	◦ 0.074	-4.622	• 0.000
Ethiopia	0.674	• 0.000	0.278	0.104	0.286	◦ 0.066	-1.783	0.389	-4.346	• 0.000	-4.758	• 0.000
Guatemala	0.509	• 0.001	0.503	• 0.002	0.168	0.140	-2.501	0.115	-3.304	• 0.015	-7.510	• 0.000
Haiti	0.516	• 0.003	0.571	• 0.001	0.412	• 0.018	-3.045	• 0.031	-2.868	• 0.049	-3.559	• 0.007
Kenya	0.640	• 0.000	0.505	• 0.002	0.520	• 0.000	-3.203	• 0.020	-3.545	• 0.007	-4.304	• 0.000
Malawi	0.428	• 0.019	0.360	• 0.050	0.279	0.113	-3.305	• 0.015	-3.618	• 0.005	-4.296	• 0.000
Mali	0.363	• 0.035	0.371	• 0.031	0.079	0.658	-4.549	• 0.000	-4.251	• 0.001	-4.941	• 0.000
Mozambique	0.594	• 0.000	0.528	• 0.001	0.694	• 0.000	-2.820	◦ 0.055	-3.111	• 0.026	-2.775	◦ 0.062
Niger	0.499	• 0.001	0.350	• 0.026	-0.062	0.729	-3.595	• 0.006	-4.345	• 0.000	-6.016	• 0.000
Nigeria	0.280	• 0.036	0.303	• 0.021	0.264	0.131	-5.637	• 0.000	-2.564	0.101	-4.323	• 0.000
Somalia	0.533	• 0.003	0.552	• 0.000	0.678	• 0.000	-2.765	◦ 0.063	-3.262	• 0.017	-2.787	◦ 0.060
South Sudan	0.618	• 0.000	0.453	• 0.011	0.918	• 0.000	-2.655	◦ 0.082	-3.275	• 0.016	-0.430	0.905
Uganda	0.511	• 0.001	0.523	• 0.001	0.271	◦ 0.060	-3.407	• 0.011	-3.204	• 0.020	-5.242	• 0.000
Zimbabwe	0.493	• 0.004	0.226	0.197	0.812	• 0.000	-3.163	• 0.022	-4.515	• 0.000	-1.310	0.625

• : significant on 5% level, ◦ : significant on 10% level.

(b) Famine risk state F_{ct}

	AR(1) regression				Augmented Dickey-Fuller			
	F^L ("low" IPC)		F^H ("high" IPC)		F^L ("low" IPC)		F^H ("high" IPC)	
Country	coef.	p-value	coef.	p-value	stat.	p-value	stat.	p-value
Afghanistan	0.749	• 0.000	-2.032	0.273	0.349	◦ 0.057	-3.689	• 0.004
Chad	0.387	• 0.018	-3.932	• 0.002	0.267	0.121	-4.468	• 0.000
Ethiopia	0.545	• 0.001	-1.808	0.377	0.362	• 0.027	-4.422	• 0.000
Guatemala	0.292	◦ 0.087	-2.616	◦ 0.090	0.384	• 0.017	-4.054	• 0.001
Haiti	0.403	• 0.020	-3.618	• 0.005	0.380	• 0.030	-3.713	• 0.004
Kenya	0.389	• 0.022	-3.785	• 0.003	0.646	• 0.000	-2.729	◦ 0.069
Malawi	0.167	0.366	-4.571	• 0.000	0.317	◦ 0.071	-4.021	• 0.001
Mali	-0.028	0.874	-5.448	• 0.000	0.153	0.388	-5.130	• 0.000
Mozambique	0.713	• 0.000	-2.275	0.180	0.734	• 0.000	-2.123	0.236
Niger	0.338	• 0.042	-4.142	• 0.001	0.319	◦ 0.065	-4.077	• 0.001
Nigeria	0.446	• 0.004	-3.860	• 0.002	0.569	• 0.000	-2.980	• 0.037
Somalia	0.212	0.231	-2.356	0.155	0.596	• 0.000	-3.828	• 0.003
South Sudan	0.546	• 0.001	-3.057	• 0.030	0.903	• 0.000	-0.995	0.755
Uganda	0.552	• 0.001	-3.132	• 0.024	0.439	• 0.007	-3.682	• 0.004
Zimbabwe	0.422	• 0.014	-3.573	• 0.006	0.720	• 0.000	-2.052	0.264

• : significant on 5% level, ◦ : significant on 10% level.

Figure 8: Residual distributions of regularized panel model

We examine the model fit through the residuals from the regularized PVARX model with the selected variables described in Tables 3 and 4. In the scatter plot, we see the residuals from the predictions $\hat{F}_t^L, \hat{F}_t^H$ against each other and overlay a bi-variate kernel density. The two adjacent distribution plots show the marginal densities of the residuals and compare it to standard normal distributions in dashed lines. The residual distributions exhibit excess kurtosis versus the standard normal, i.e. more probability mass in the center and fatter tails.

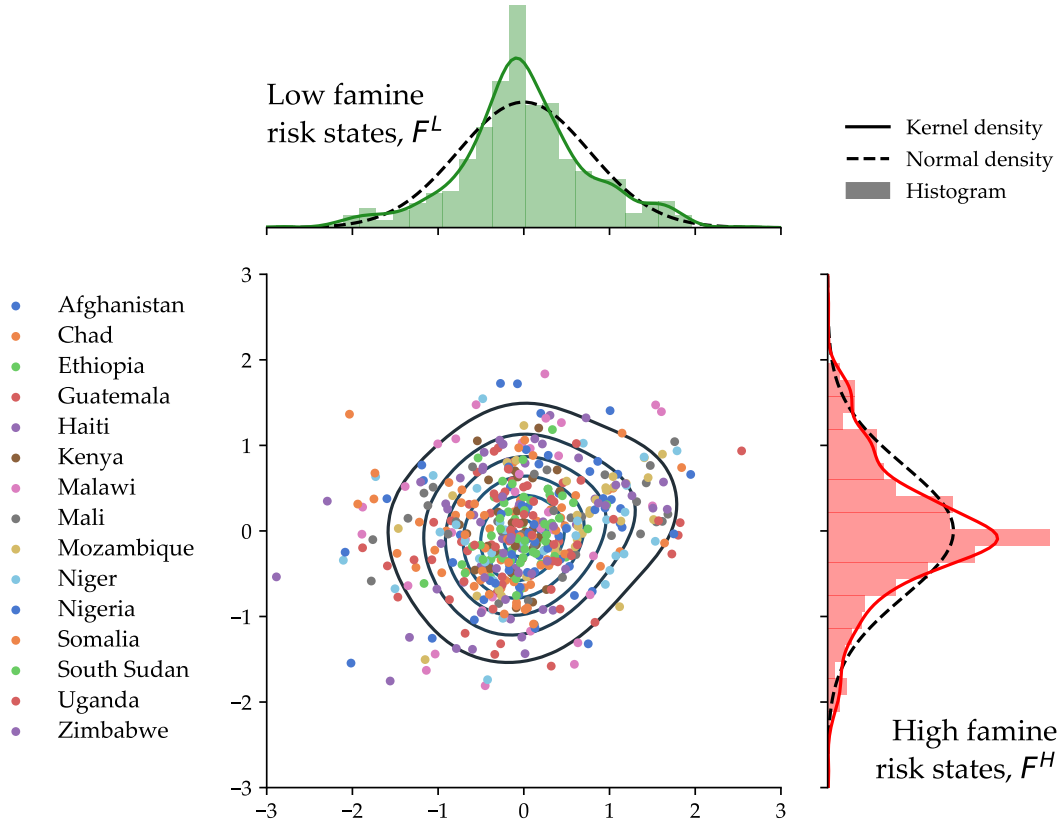

8.1 Famine risk state transformation

The famine transformation function G consists of two steps. For clarity of notation we abstract from time and country subscripts. First, a compositional transformation C that maps the food insecurity distribution vector X to a intermediate vector $Y = C(X)$. Second, an inverse logistic transformation that maps the intermediate vector Y into the famine risk state vector $F = \text{logit}^{-1}(Y)$. The famine transformation function G is therefore a composite of both, that is $F = G(X) = (\text{logit}^{-1} \circ C)(X)$. The the subsequent sections we discuss the two components in detail and how to recover food insecurity distributions from famine risk state vectors.

8.1.1 Compositional transformation

The function C can take two forms. Either, the hyper-spherical transformation (Wang et al., 2007) or the alpha-transformation (Tsagris et al., 2016). Both transformations remove the redundant dimension of X, which lies within the K-dimensional unit simplex Δ^K, that is $C : \Delta^K \rightarrow S$ where S is a bounded subset of $\mathbb{R}^{K-1}$. The hyper-spherical transformation maps X to polar coordinates while the alpha-transformation uses a Helmert transformation. In this article, we employ the latter with the following implementation details

$$Y = C(X; \alpha, H) = H\frac{1}{\alpha}(Kz - \iota)$$

with the vector of ones ι, a scalar $\alpha = 0.5$ and the compositional power transformation

$$z = u(X, \alpha) = \frac{1}{\sum_{k=1}^{K} x_k^\alpha} \left[x_1^\alpha, ..., x_K^\alpha\right]^\top$$

where H is a Helmert matrix, which is an isometric and bijection linear projection from Δ^K to $\mathbb{R}^{K-1}$. In our case with $K = 3$ we use

$$H = \begin{bmatrix} 0.707 & -0.707 & 0.000 \\ 0.408 & 0.408 & -0.816 \end{bmatrix}$$

For further details on the general case $K \neq 3$, we refer to Tsagris et al. (2016).

8.1.2 Logistic transformation

This transformation is a technical detail that ensures that the predicted values $\hat{Y}$ are in the simplex S and can therefore be safely transformed back into famine risk distribution vectors $X = C^{-1}(Y)$. The leads us to include an inverse logit transformation after the compositional transformation, such that the statistical model is fitted to F rather than Y. The logistic transformation is defined as

$$\text{logit}(F) = (b - a)\frac{\exp(F)}{1 + \exp(F)} + a$$

and the boundaries a, b depend on the α value used in the compositional transformation. For $\alpha = 0.5$ the boundaries for F^L are $[-4.243, 4.243]$ and for F^H they are $[-4.899, 2.449]$.

Figure 9: Compositional transformation

The two panels depict the translation between bounded IPC population distributions X (y-axis) and the unbounded components of the famine risk state vector F (x-axes). The transformation is bijective and was chosen such that *decreasing* famine risk states indicate *deteriorating* food insecurity situations. E.g. as F^L decreases, the share of population in IPC1 (none/minimal) transitions to IPC2 (stressed). Similarly, as F^H decreases, the population shares in IPC1,2 transition in equal parts to IPC3+ where urgent action is needed.

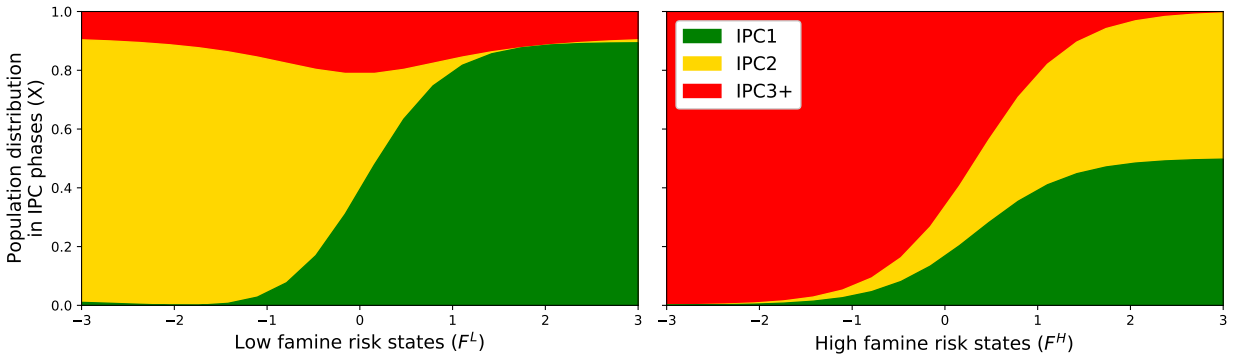

8.1.3 Recovering the full food insecurity distribution

After the model has been estimated on $F = G(X) = \text{logit}^{-1}(C(X))$, we can recover X from F by using the inverse of G, that is $X = G^{-1}(F) = C^{-1}(\text{logit}(F))$. The inverse of C is

$$C^{-1}(F; \alpha, H) = u^{-1}(\alpha H^\top F + \iota)$$

and accordingly

$$u^{-1}(z; \alpha) = \frac{1}{\sum_{k=1}^{K} z_k^{1/\alpha}} \left[z_1^{1/\alpha}, ..., z_K^{1/\alpha} \right]^\top .$$

References

Andrée, B. P. J. (2020). *Theory and Application of Dynamic Spatial Time Series Models*. Amsterdam, Rozenberg Publishers; the Tinbergen Institute.

Andrée, B. P. J., Chamorro, A., Spencer, P., Koomen, E., & Dogo, H. (2019). Revisiting the relation between economic growth and the environment; a global assessment of deforestation, pollution and carbon emission. *Renewable and Sustainable Energy Reviews*, *114*, 109221.

Andrée, B. P. J., Kraay, A., Chamorro, A., Spencer, P., & Wang, D. (2020). Predicting Food Crisis. *World Bank Policy Research Working Papers*.

Bańbura, M., Giannone, D., & Reichlin, L. (2008). *Large Bayesian VARs* (ECB Working Paper Series No. 966).

Barbier, E. B., & Hochard, J. P. (2018). Land degradation and poverty. *Nature Sustainability*, *1*(11), 623–631.

Barrett, C. B., & Bevis, L. E. (2015). The self-reinforcing feedback between low soil fertility and chronic poverty. *Nature Geoscience*, *8*(12), 907–912.

Bergmeir, C., & Benítez, J. M. (2012). On the use of cross-validation for time series predictor evaluation. *Information Sciences*, *191*, 192–213.

Bergmeir, C., Hyndman, R. J., & Koo, B. (2018). A Note on the Validity of Cross-Validation for Evaluating Autoregressive Time Series Prediction. *Computational Statistics and Data Analysis*.

Berk, R., Brown, L., Buja, A., Zhang, K., & Zhao, L. (2013). Valid post-selection inference. *The Annals of Statistics*, *41*(2), 802–837.

Blasques, F., & Duplinskiy, A. (2018). Penalized indirect inference. *Journal of Econometrics*, *205*(1), 34–54.

Carriero, A., Clark, T. E., & Marcellino, M. (2011). *Bayesian VARs Specification Choices and Forecast Accuracy* (Federal Reserve Bank of Cleveland, Working Paper no. 11-12).

Christoffel, K., Coenen, G., & Warne, A. (2011). Forecasting With DSGE Models, In *The Oxford Handbook of Economic Forecasting*. Oxford University Press.

Ciccarelli, M., & Rebucci, A. (2003). *Bayesian Vars : A Survey of the Recent Literature with An Application to the European Monetary System* (IMF Working Paper WP/03/102). International Monetary Fund.

Devereux, S. (2000). *Famine in the Twentieth Century* (IDS Working Paper No. 105). Institute of Development Studies.

Devereux, S. (2006). *The New Famines: Why Famines Persist in an Era of Globalization*. Routledge.

Diogo, V., Reidsma, P., Schaap, B., Andree, B. P. J., & Koomen, E. (2017). Assessing local and regional economic impacts of climatic extremes and feasibility of adaptation measures in Dutch arable farming systems. *Agricultural Systems*, *157*, 216–229.

D'Souza, A., & Jolliffe, D. (2013). Conflict, food price shocks, and food insecurity: The experience of Afghan households. *Food Policy*.

Duraiappah, A. K. (1998). Poverty and environmental degradation: A review and analysis of the nexus. *World Development*, *26*(12), 2169–2179.

FAO. (2012). *Country Programming Framework (CPF) 2012-2015 for Afghanistan* (tech. rep.). Food and Agriculture Organization of the United Nations.

FAO. (2019). *The state of food security and nutrition in the world: Safeguarding against economic slowdowns and downturns.* (tech. rep.). Food and Agriculture Organization of the United Nations.

FAO. (2020). *Desert Locust Bulletin* (Report No. 498). Food and Agriculture Organization of the United Nations.

Giannone, D., Lenza, M., & Primiceri, G. E. (2019). Priors for the Long Run. *Journal of the American Statistical Association*, *114*(526), 565–580.

Grainger, A. (1990). *The threatening desert: Controlling desertification*. London, John Wiley & Sons, Ltd.

Headey, D. (2011). Rethinking the global food crisis: The role of trade shocks. *Food Policy*.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*, *12*(1), 55–67.

Ingram, J., Ericksen, P. J., & Liverman, D. (2010). *Food security and global environmental change*. London, Earthscan.

IPC Global Partners. (2019). *Integrated Food Security Phase Classification Technical Manual Version 3.0. Evidence and Standards for Better Food Security and Nutrition Decisions*. Rome, IPC Global Partners.

Lee, J. D., Sun, D. L., Sun, Y., & Taylor, J. E. (2016). Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, *44*(3), 907–927.

Lentz, E. C., Michelson, H., Baylis, K., & Zhou, Y. (2019). A data-driven approach improves food insecurity crisis prediction. *World Development*, *122*, 399–409.

Litterman, R. B. (1986). Forecasting with Bayesian Vector Autoregressions: Five Years of Experience. *Journal of Business & Economic Statistics*, *4*(1), 25.

Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett, Eds.). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30*, Curran Associates, Inc.

Maxwell, D., Khalif, A., Hailey, P., & Checchi, F. (2020). Viewpoint: Determining famine: Multi-dimensional analysis for the twenty-first century. *Food Policy*, *92*, 101832.

Misselhorn, A. A. (2005). What drives food insecurity in southern Africa? a meta-analysis of household economy studies. *Global Environmental Change*.

Myers, S. S., Smith, M. R., Guth, S., Golden, C. D., Vaitla, B., Mueller, N. D., Dangour, A. D., & Huybers, P. (2017). Climate Change and Global Food Systems: Potential Impacts on Food Security and Undernutrition. *Annual Review of Public Health*, *38*(1), 259–277.

Pape, U. J., Parisotto, L., Phipps-Ebeler, V., Mueller, A. J. M., Ralston, L. R., Nezam, T., & Sharma, A. (2018). *Impact of Conflict and Shocks on Poverty : South Sudan Poverty Assessment 2017* (Report No: AUS0000204). The World Bank.

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., & Dormann, C. F. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, *40*(8), 913–929.

Salama, P., Moloney, G., Bilukha, O. O., Talley, L., Maxwell, D., Hailey, P., Hillbruner, C., Masese-Mwirigi, L., Odundo, E., & Golden, M. H. (2012). Famine in Somalia: Evidence for a declaration. *Global Food Security*, *1*(1), 13–19.

Sims, C. A., & Zha, T. (1998). Bayesian Methods for Dynamic Multivariate Models. *International Economic Review*, *39*(4), 949.

Singh, R. B. (2012). Climate Change and Food Security, In *Improving crop productivity in sustainable agriculture*.

Stern, D. I., Common, M. S., & Barbier, E. B. (1996). Economic growth and environmental degradation: The environmental Kuznets curve and sustainable development. *World Development*, *24*(7), 1151–1160.

Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, *58*(1), 267–288.

Tsagris, M., Preston, S., & Wood, A. T. A. (2016). Improved Classification for Compositional Data Using the alpha-transformation. *Journal of Classification*, *33*(2), 243–261.

Wang, H., Liu, Q., Mok, H. M., Fu, L., & Tse, W. M. (2007). A hyperspherical transformation forecasting model for compositional data. *European Journal of Operational Research*, *179*(2), 459–468.

WFP. (2020). *2020 - Global Report on Food Crises* (WFP Reports). World Food Programme.

Zhang, Y., & Yang, Y. (2015). Cross-validation for selecting a model selection procedure. *Journal of Econometrics*, *187*(1), 95–112.

Zou, H. (2006). The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*.

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B (Methodological)*, *67*(2), 301–320.