

**Empirical Forecasting of Slow-Onset Disasters for Improved Emergency Response:
An Application to Kenya's Arid and Semi-Arid Lands**

Andrew Mude¹, Christopher Barrett, John McPeak, Robert Kaitho, Patti Kristjansen

**Incomplete Rough Draft:
Please do not cite without authors permission
Version: May, 2006**

The authors would like to thank Luc Christaensen, Allan Kute, Francesca Molinari, Nicholas Ndiwa, Kamau Ngamau, Abisalom Omolo, Maren Radney, Rob Rose, Julia Stone, Ben Watkins, Stephen Younger and seminar participants at Cornell University for all their helpful comments, insights and assistance. This research was partly funded by the USAID's Pastoral Risk Management (PARIMA) program, the International Livestock Research Institute, and the World Bank. We also acknowledge the contribution of the USAID Global Livestock Collaborative Research Support Programs (GL-CRSP), Livestock Early Warning System (LEWS), and its successor the Livestock Information Network and Knowledge System (LINKS).

¹ Corresponding author. Email: agm26@cornell.edu. Mail: Department of Economics, Cornell University, Ithaca NY, 14853

1 INTRODUCTION

The ability to forecast, with reasonable accuracy, the onset, duration and severity of droughts, floods, disease outbreaks, etc., especially in terms of their prospective human welfare effects, is critical to the design of rapid, timely and cost-effective emergency response systems that can minimize the suffering of populations adversely affected by disasters. As the consensus on climate change and its consequences grows, there is an increasing worry that the frequency of such disasters will continue to rise. Incidences of humanitarian catastrophe and the ensuing demand for emergency response are therefore widely expected to increase.

Given the finite resources allocated for emergency response initiatives, there is growing demand worldwide for the development of rigorous, efficient and practicable methods of emergency needs assessment. The World Food Program's emerging project on Strengthening Emergency Needs Assessment Capacity (SENAC), for example, is currently engaged in the identification of methods and tools for emergency assessments, and for evaluating the effects of different types and combinations of shocks on various livelihood groups. The famine early warning systems network (FEWS-NET), funded by USAID, is a comprehensive effort to provide timely assessments on the status of food insecurity threats in 20 African countries, as well as in Afghanistan and Haiti. They are also tasked with strengthening regional and national early warning and food security organizations through networking and capacity development. The USAID Global Livestock Collaborative Research Support Program (GL-CRSP) Livestock Early Warning System (LEWS) project and its successor, the Livestock Information Network and Knowledge System (LINKS) project, represent a complementary effort to provide pastoralists in East Africa with a tools to track climate and forage conditions so as to inform herd management decisions and mitigate the consequences of drought.

There could be great humanitarian and economic value to accurate forecasting of the needs of populations vulnerable to slow-onset disasters such as drought – as quite distinct from rapid-onset disasters such as those associated with many other natural disasters, such as earthquakes, hurricanes or tsunamis – and identification of where and when to intervene and at what scale. In this paper, we aim to develop such a method.

We focus on the arid lands of northern Kenya, largely populated by nomadic pastoralists and particularly vulnerable to covariate shocks in the form of droughts and floods.

Our primary objective is to make use of household data collected over several years by the Arid Lands Resource Management Project (ALRMP) in Kenya, and the spatially explicit data on forage conditions, rainfall and NDVI generated by the LEWS/LINKS team to develop an empirical forecasting model that can predict the expected human impact of covariate shocks and thereby provide a useful statistical method for early warning emergency needs assessment.

We aim to explore the village-level joint dynamics of herd size – pastoralists’ primary asset (Lybbert et al. 2004, McPeak 2004), lactation rates – pastoralists’ primary income source (McPeak 2004), climate and forage conditions - the key exogenous environmental determinants of productivity dynamics in pastoralist systems (Coppock et al., 1986; Ellis, 1994; Ellis and Swift, 1988), and child nutritional status, a key outcome variable of humanitarian and economic interest. The idea is to identify the dynamic and potentially nonlinear interrelationships between natural resource conditions, herd sizes, livestock productivity (as measured by lactation rates) and human well-being (as proxied by child nutritional indicators), and especially the threshold points that serve as leading indicators for imminent crisis (e.g., whether patterns in climate variables and/or herd sizes can predict the aggregate movement of child nutritional indicators).

The rest of the paper is structured as follows. Section 2 summarizes the nature and content of the data we use. In section 3, we briefly describe the methodology we apply to estimate our forecasting model. Descriptive statistics of key model variables are presented and discussed in section 4. In section 5 we present forecasting results and test the performance of our forecasts. Finally, section 6 concludes.

2 DATA

The World Bank-funded ALRMP seeks to address the vulnerability of the populations living in Kenya's arid lands and to improve their ability to manage the risks that commonly befall the region (World Bank Group, 2003). As part of the project, repeated cross-sectional data have been collected in various communities across Kenya's arid districts since 1996. Data relevant for our objective are sourced from household level surveys that contain detailed information on livestock such as herd sizes, mortality rates, lactation rates, and managed off-take rates. Critically, child nutritional data in the form of mid-upper arm circumference (MUAC) was also collected.

While data have been collected monthly from 1996 and across 10 different districts, poor organization and storage unfortunately resulted in significant quantities of lost data that rendered many sections too patchy for any rigorous analysis. Consequently, our effective data set, while still substantial, constitutes a mere subset of what was collected. Furthermore, no authoritative document on the collection procedures and sampling methodologies employed exists. As such, while we know that community sites were purposively selected to take into account population density and spatial distribution across a district, we are only aware that enumerators were asked to randomly select 30 households per community without being clear on the method of randomization used or whether it was enforced. Nevertheless, while this forces us to be cautious about (mis)representing our results as statistically representative of any single place, the data seem sufficiently rich to shed important light on dynamic processes that are as yet not well understood and for which quantitative evidence of any sort is distressingly scarce.

And though surveyed households were theoretically revisited each month for a year before a new sample was generated, insufficient record keeping prevents individual households from being linked across periods. As such, we create a community-level pseudo-panel by generating community-level summary statistics of the pertinent variables. Despite the unavoidable loss of information, such a pseudo-panel is well suited to investigating the impacts of covariate, community-level shocks whose dynamics are more relevant for external emergency interventions than are idiosyncratic, household-level shocks (Deaton, 1985; McKenzie, 2004).

We supplement the ALRMP data with a rich source of climate and forage availability data collected and produced by LEWS/LINKS researchers. LEWS/LINKS has developed a set of technologies and models that provide high-resolution, high-frequency estimates of livestock forage availability in the pastoralist-dense regions of East Africa (Kaitho et al. 2003, Stuth et al. 2003). These data provide key variables for our model. Most major covariate shocks that hit pastoralist communities are a function of adverse forage and water conditions associated with climate fluctuations (Ellis and Swift 1988, Ellis 1994, Galvin et al. 2001). As changes in livestock fertility and mortality are closely related to forage and water quality and availability, access to the detailed dynamics of these variables would greatly enhance the value and precision of models generated to investigate and predict the human impact of the climate shocks that frequently destabilize pastoralist communities.

3 METHODOLOGY

As adverse covariate shocks come in various forms and have varied effects on the affected populations, emergency assessments often focus on the response, or predicted response, of critical indicator variables. Proxies of food insecurity such as the real price of key staples or the level of food production, availability, or expenditure can yield estimates of the population's vulnerability to hunger and starvation. Fortunately, ALRMP data include sample readings of the Mid-Upper Arm Circumference (MUAC) for selected community children. Among indicators of nutritional wellbeing, MUAC is particularly well-suited for our purposes. As a measure of wasting, MUAC is capable of capturing short-term fluctuations in the presence of nutritional stress and can thus serve as a gauge of the human impact of various shocks. Furthermore, MUAC is easier and less costly to collect than weight-for-height (W/H), the most commonly used and most documented anthropometric measure for wasting. Indeed, several studies have shown MUAC to be a far better predictor of child mortality than W/H (Chen et al. 1980, Alam et al. 1989, Vella et al. 1994). As such, MUAC is a particularly appropriate indicator of the welfare impacts of a humanitarian crisis.

As we are primarily concerned with generating the most accurate forecasts possible, our objective is to maximize the efficiency of the model. While estimating

unbiased coefficients would allow us to reasonably infer the relationship between key model variables and MUAC, attempting to correct for endogeneity by using instruments, or imposing limits on the error structure would necessarily reduce efficiency and thus the accuracy of forecasts. Consequently, we leave the potentially interesting inference-based estimation for further study and focus exclusively on optimizing the forecast capabilities of our model.

As MUAC trends are likely to display a high degree of persistence, it would be logical to include lagged values of MUAC as an explanatory variable to capture the inevitable dynamic component. We therefore estimate a dynamic model for our unbalanced panel that takes the following form:

$$y_{it} = \lambda y_{i,t-1} + X'_{it} \beta_t + X'_{i,t-1} \beta_{t-1} + u_i + \varphi_t + v_{it} \quad \left(\begin{array}{l} i = 1, \dots, N \\ t = 1, \dots, T \end{array} \right)^2,$$

where y_{it} denotes is an observation on the dependent variable for unit i (community level MUAC means in our case), with $y_{i,t-1}$ the lagged MUAC means that captures the dynamic nature of the relationship via the unknown coefficient λ . X_{it} is a vector of contemporary explanatory variables that capture observed heterogeneity. In this case, X_{it} would be all those variables thought to have some effect on MUAC. These include variables such as herd size, herd mortality, lactation and livestock sales and slaughters that many affect MUAC directly via the provision of nutrients and calories in milk and meat, or indirectly through changes in liquidity due to managed offtake. Policy variables such as the provision of food aid, which occurs with some regularity across our sample, are also to be included in X_{it} as are biophysical variables such as rain, NDVI and forage conditions that indirectly affect MUAC via herd compositions, food prices etc.

u_i denotes the community effect capturing the unobserved heterogeneity, φ_t are time dummies that capture additional seasonality not controlled for by temporal variation in $X_{i,t}$, and v_{it} is the unobserved error term. The dynamic model characterized above is quite general and, as implied by the inclusion of $X_{i,t-1}$, can accommodate multiple lags of both the dependent and independent variables. Lag structure is generally chosen based

² Recall that the panel is unbalanced. As such, some observations will be missing in the interval [0,T]

on the assumed data generating process, the estimators chosen and the researcher's objective. As our objective is to maximize efficiency, we will use the Root Mean Square Error (RMSE) to select for optimal lag structure.

The unique focus on efficiency considerably simplifies our estimation. One key problem with estimating a dynamic model is that the presence of the lagged dependent variable introduces endogeneity that causes the least squared dummy variable (LSDV) estimator to become biased and inconsistent.³ While several estimators have been developed that deal with this problem (Anderson and Hsiao, 1982; Arellano and Bond, 1991; Kiviet, 1995; Hansen, 2001), they trade-off efficiency for unbiasedness and result in higher variances than the LSDV estimator. The inclusion of other predetermined or endogenous explanatory variables, such as food aid, would also result in biased coefficients if not appropriately instrumented. As we are not interested in coefficient estimates, we can abstract away from these concerns and simply use the LSDV estimator which is well known to be of least variance.

4 DESCRIPTIVE STATISTICS

In what follows we briefly describe the data used to estimate our models and present graphical representations of the trends of key variables. The data structure, shown in Table 1, lists the number of observations per district for which we have the full set of pertinent variables in the specified time periods. Recall that the observations are all community level aggregates, largely community means or functions of means. Though each survey round was to sample 30 households per community, smaller samples of much less are frequently observed in the data. To strike a balance between throwing out observations and taking averages over unreasonably small sample sizes, we only include the 55 communities for which 15 or more households were sampled in a given time period.

³ The LSDV estimator is simply an Ordinary Least Squared, Fixed-Effects Panel estimator.

Table 1: Available Observations by District, Across Time.

YYYYMM	BARINGO	MARSABIT	SAMBURU	TURKANA	YYYYMM	BARINGO	MARSABIT	SAMBURU	TURKANA
200002		9	13		200210	12	9	13	19
200003		7	13		200211	12	8	12	19
200004		9	14	16	200212	11	9	12	19
200005			14	19	200301	11	9	13	19
200006		8	14	17	200302		9	13	18
200007		9	14	8	200303	11	9	14	18
200008			14	19	200304	12	9	12	19
200009		7	13	19	200305	11	9	11	19
200010		9	13	19	200306	10	9	14	19
200011		8	14	18	200307	10	9	14	18
200012		7	14	18	200308	10	9	13	19
200101	12	8	14	20	200309	7	9	13	19
200102	11	5	14	19	200310	7	9	13	19
200103	12	6	13	20	200311		7	11	19
200104	12	8	12	18	200312	1	8	13	19
200105	12	8	13	19	200401	12	9	13	19
200106	11	8	13	17	200402	12	9	13	20
200107	12	8	12	19	200403	12	9	14	19
200108	12	8	11	17	200404	12	9	12	18
200109	12	8	12	18	200405	12	9	8	18
200110	11	9	13	18	200406	12	9	13	19
200111	12	9	13	19	200407	10	9	11	18
200112	12	9	11	19	200408	12	9	14	20
200201	12	9	14	19	200409	10	9	13	19
200202	12	8	13	18	200410	11	9	14	19
200203	12	9	13	20	200411	11	9	10	18
200204	12	9	14	19	200412	12	8	12	18
200205	12	9	13	18	200501	11	9	14	16
200206	12	9	14	19	200502	12	7	13	18
200207	12	8	13	19	200503	12	9	14	19
200208	12	9	13	18	200504	11	9		18
200209	12	9	13	19	200505	12	9		18

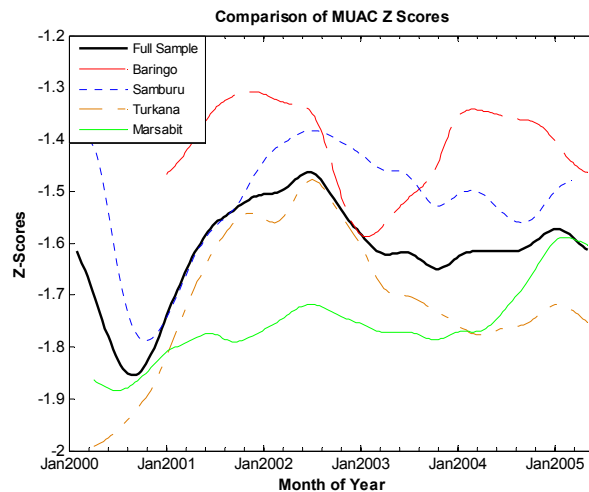
As previously noted, poor data entry and storage protocols render the ALRMP panel data quite unbalanced. All time periods suffer the loss of several observations, often across multiple districts. In certain cases, all observations from a particular district are missing. No district has observations in each time period under consideration, nor can any districts claim a contiguous set of observations for the time periods for which observations are available. As we would lose too many observations if we conditioned our estimates on a balanced sub-sample, we use the unbalanced panel as is.⁴ Fortunately,

⁴ We attempted to artificially balance the panel by imputing the missing variables. Using the *hotdeck* command in STATA, which applies the Bayesian bootstrapping method of Rubin and Shenker (1986) to

most of the oft-used dynamic panel estimators are capable of estimating unbalanced panels. Moreover, it has been shown that given the same number of effective lagged observations, a highly unbalanced panel does not negatively affect the bias or Root Mean Squared Error (RMSE) of estimators (Bruno, 2005).

To get a sense of the dynamics of key variables in our model, we estimated kernel regressions of the variables across time. Figure 1 captures the trend of our dependent variable, MUAC, across our sample districts. As MUAC was collected for children aged 6-59 months residing within sample households or in close proximity (for a maximum of 5 children per household), we transformed our MUAC data into standardized Z-scores.⁵ We used the internationally recognized 1978 CDC/WHO growth chart whose reference population is American children sampled in the 1977 National Center for Health Statistics survey.

Figure 1



The generally low Z-scores show that even during relatively good periods, the average MUAC value of the sample population is much less than the reference population. Nonetheless, within our sample, significant temporal fluctuations in the Z-


estimate the missing values, we created multiple imputed data sets. Estimates based on the imputed data sets were significantly inferior to those on the unbalanced panel across all specifications that we tried. This is likely because the number of missing observations is quite large and imputing all the variables of the missing observations adds a considerable amount of error in the variables limiting the accuracy and consistency of estimates based on them.

⁵ While standardization is condition on both age and sex, we only had age data. As such, we generated Z-scores for both males and females and took the average.

scores are suggestive of the influence exogenous shocks exert on nutritional wellbeing. 2000 looks to have been a particularly bad year with MUAC levels falling drastically to their lowest levels in the sample period. Rapid gains in MUAC were registered early in 2001 with sustained improvement until mid-2002 when a gradual fall set in but generally flattened out by early 2003. While the trend in sample averages are somewhat mirrored by the district means, there are also clear inter-district differences, both in levels and trends.

Trends in rainfall and forages (Figure 2) rates may partly explain the MUAC dynamics. As can be inferred from the figure, 2000 was a year of drought, with the main rains, often coming in March and April, failing. Lack of rain results in falling forage availability, and thus the shrinking of available grazing for livestock. Largely habited by pastoralists who live off their animals, the drought may account for the low MUAC levels witnessed in 2000. However, strict agroecological determinism of child nutrition status seems implausible. Other factors must also influence MUAC as is evident with falling MUAC rates around mid-2002 despite relatively high forage.


Figure 2.



The increasing incidence of drought and subsequent famine in Kenya’s arid north in the past decade or so has catalyzed a significant food aid response whose intensity varies across time depending on the degree of the crisis affecting recipient populations as well as other socioeconomic and political factors. Figure 3 shows the trends in the fraction of households that receive Unimix (a micronutrient fortified corn soy blend

specifically formulated for undernourished children) and the fraction receiving food aid in the form of cereals. While Unimix is not as widely distributed as cereals, their pattern of distribution across time is generally similar despite a notable exception in the early months of 2000 where Unimix distribution took longer to respond to the MUAC crisis than did cereal provision. The rapid scaling back of both Unimix and cereal support in 2002 may explain why, despite better than average forage conditions, MUAC Z-scores began to fall from mid-2002.

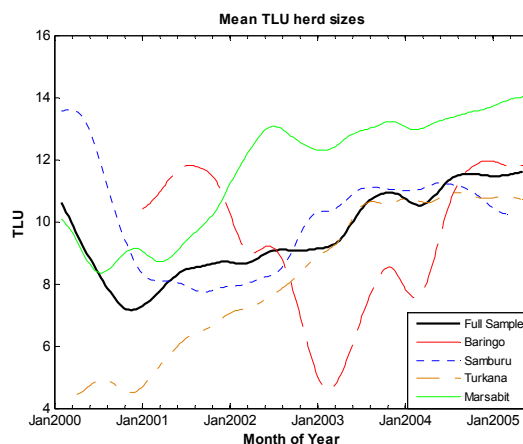
Figure 3



Livestock represent the main stock of assets and a significant source of nutrition for pastoralists. As such, herd dynamics may also contain an important source of information pertaining to the health and wellbeing of pastoralist populations. For example, declining lactation rates would limit the availability of a key source of nutrition among pastoralists which, if severe enough, would translate to lower MUAC. Increasing incidence of sales or slaughter on the other hand, may signal coping strategy behavior triggered by negative shocks. In Figure 4, we present herd sizes aggregated into Total Livestock Units (TLU) and compare their dynamics across districts.⁶

⁶ TLUs allow for a comparison of livestock quantities across species. One TLU is equivalent to 1 head of cattle, 0.7 camels, 10 goats or 11 sheep.

Figure 4

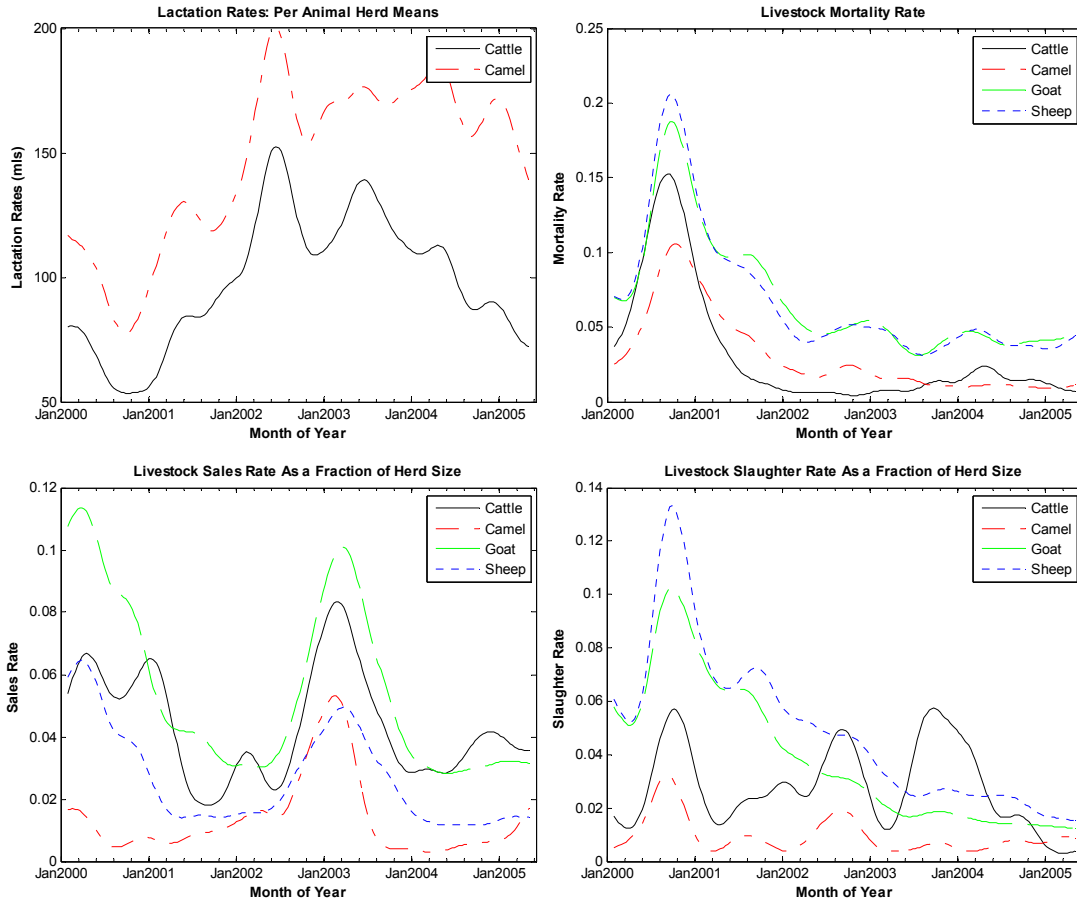


While there is a general upward trend in herd sizes across the sample period, the drought of 2000 resulted in a noticeable decline as can be clearly seen in the TLU trend. TLU trends follow a similar pattern across districts with the exception of Baringo where a significant fall in herd sizes occurs toward the end of 2002. Much of this decline can be explained by the rate of managed off-take (sales and slaughter) in Baringo, which was significantly higher than the rest of the sample at this time. In Figure 5, we graph the trends in lactation, mortality, sales and slaughter rates. As expected, the drought period in 2000 is characterized by a high degree of mortality across the four animals and significantly lower lactation rates among cattle and camels.⁷ Trends in sales and slaughter rates do not seem to follow a pattern that can be readily interpreted as either post-shock coping behavior or ex-ante risk mitigation. Nonetheless, sales and slaughter, especially among small stock are noticeably higher during the drought period⁸.

⁷ Lactation rates are calculated as daily means per herd and thus include both female and male of the species as well as young and old. This could partly explain the generally low rates of lactation posted.

⁸ As previously mentioned, the large increase in sales rates starting from the end of 2002 are largely driven by sales in Baringo. It is unclear, however, why sales in Baringo should be so high during this period. A possible explanation is that the exceptionally favorable forage availability in Baringo in early to mid 2002 created a particularly healthy stock of animals that fetched relatively high prices at the market. *El nino* floods that hit Baringo especially hard in late 2002, may also explain the spike in sales as the destruction caused by floods increased the demand for cash to cope with the unexpected catastrophe.

Figure 5



4 ESTIMATION RESULTS

4.1 Estimating the Forecasting Model

Armed with a sense of the dynamics of key model variables, we now move on to estimate our model. Recall that our primary objective is to estimate a model that generates the most accurate forecasts of MUAC levels possible. As we are not interested in making inferences on the relationship between our explanatory variables and MUAC levels, we do not concern ourselves with estimating unbiased, consistent coefficients. Instead, we focus exclusively on trying to maximize the efficiency of our model. As such, we include as many explanatory variables as possible, the only condition being that they contribute to an increase in the degrees of freedom adjusted, Root Mean Squared

Error (RMSE).⁹ We use the LSDV estimator which, as discussed in section 3, is the most efficient estimator for the circumstance.

After estimating numerous specifications of various combinations of the variables, including squared terms, cross-products, and varying lag lengths, the specification described in Table 2, offered the best RMSE.

Table 2: Model Specification

<u>Dependent Variable</u>		
MUAC Means (Z-Score)		
<u>Explanatory Variables</u>		
<u>Group</u>	<u>Lag Length</u>	<u>Variables</u>
MUAC Moments	L(1,2,3)	MUAC Mean, MUAC std. dev., MUAC Skewness, MUAC Kurtosis
Herd Dynamics (Large Stock)	L(1,2)	Size, Mortality, Sales, Slaughters
Herd Dynamics (Small Stock)	L(1,2)	Size, Mortality, Sales, Slaughters
Food Aid	L(1,2)	UNIMIX (Kgs and Fraction), Regular (Kgs and Fraction)
Biophysical	L(2,3,4,5,6)	Rain, Rain Squared, Forage, Forage Squared, NDVI, NDVI Squared
Seasonality		Monthly Dummies

Table 2 breaks the specification into groupings of the explanatory variables for ease of presentation. The first set of explanatory variables, MUAC Moments, control for the distribution of MUAC within a community. By more accurately describing the distribution of MUAC Z-Scores across a community, including higher moments of MUAC affords more precise forecasts of the means. Each MUAC moment variable, including values of the lagged dependent variable, are introduced with one, two, and three month lags. Including a fourth lag reduced the resulting RMSE. This most likely due to the unbalanced nature of the ALRMP panel data which causes an increasing number of observations to be dropped as additional lags are introduced.

The herd dynamics variables were disaggregated into two different categories: one for large stock (cattle and camels), and the other for small stock (goats and sheep).¹⁰

⁹ RMSE is the most commonly used criterion for comparing the performance of forecasting models. Other popular criteria include the asymptotically efficient Akaike Information Criterion (AIC), and the asymptotically consistent Schwarz Information Criterion (SIC). The main difference between these three model selection criteria is that they penalize degrees of freedom differently. For the various specifications we estimated we also calculated the resulting AIC and SIC. In no situation did any of the three criteria disagree as to which model was best, probably because degrees of freedom represent a large fraction of the total available observations. For simplicity, we thus only present the RMSE.

¹⁰ Including separate variables for each of the four animals reduced the RMSE.

As Figure 5 shows, large stock trends generally follow a similar pattern while small stock trends often move in tandem. These variables are also calculated in TLUs. As food aid receipts within sample households are quite considerable, and as they likely have an impact on MUAC levels, we include four food aid related variables; the fraction of households receiving UNIMIX, the fraction receiving regular cereals, the amount (in kilograms) of UNIMIX received per recipient household, and the amount of cereal received. Both the herd dynamic variables and the food aid variables were introduced with one and two month lags.

The biophysical variables, mean monthly rain (mm), forage (kg/ha) and NDVI were the only variables which were introduced with squared terms. In addition, each variable was included with five lags (from a two through six-month lag).¹¹ As the LEWS/LINKS data set from which these variables were sourced is complete, introducing multiple lags does not decrease the number of estimable observations. To account for other sources of seasonal variation not captured by these biophysical variables, we include monthly dummies.

Note that we do not introduce any of the explanatory variables contemporaneously. Despite the likelihood that this would increase RMSE, we leave out contemporary variables in order to allow for a legitimate forecast. This specification enables one-month ahead forecasts of MUAC means. Given the lags in emergency food aid response caused by bureaucratic and other constraints (Barrett and Maxwell, 2006), a one-month forecast is admittedly short, leaving little leeway for aid workers to make effective use of the forecasts. As you will note, the inverse relationship that exists between forecasting horizon and forecast precision involves a delicate tradeoff between these two parameters, both of which determine the usefulness of the forecasting model. As an alternative, we also estimate a three-month ahead forecasting model. We utilize the same specification presented in Table 2, the difference being that each variable is lagged by an additional two months in this case.¹²

¹¹ Starting from a one-month lag, or extending to a seven-month lag does not improve results. As the effect of biophysical variables on MUAC likely occurs with a lag, the fact that including a one-month lag did not improve the forecast accuracy suggests that the effect of these variables on MUAC is felt with more than a month lag. Conversely, the effect is likely to wear off after six-months.

¹² So, for example, the lag structure for the biophysical variables now becomes, L(4,5,6,7,8)

We present key statistics from both the one-month ahead and the three-month ahead forecast models in Table 3. We also include few statistics on the dependent variable. However, as both models introduce a considerable number of explanatory variables, and as we are not concerned about inference or the precision of estimated coefficients, we do not include coefficient estimates or t-statistics.¹³

Table 3: Forecasting MUAC Means: Model Statistics

		One-Month Ahead	Three-Month Ahead
R-squared:	within	0.5297	0.3430
	between	0.9800	0.8806
	overall	0.7848	0.6195
RMSE		0.1666	0.1986
# of Observations		2290	2185
# of Groups		53	53
Observations per group:	Min	19	17
	Avg	43.2	41.2
	Max	58	56
Dependent Variable Statistics			
Mean: -1.59	Std. dev: 0.40	Min: -2.73	Max: 0.27

The results are quite encouraging and suggest that forecasts generated from this model are likely to be quite accurate. At 0.1666 and 0.1986 for the one and three month models respectively, the RMSE, which defines the mean error deviation of the forecast from the actual, reveals an agreeable forecast performance, especially given that the variable being forecast has a sample standard deviation of 0.40. Furthermore, the models fit the data surprisingly well, generating high r-squared values for both within community temporal variation and between community spatial variation. The between r-squared are strikingly high at 0.98 and 0.88 for the one and three month models respectively. While the within r-squares, which are also relatively impressive at 0.53 and 0.34, are more relevant for forecasting, the extremely high between r-squares suggest that our explanatory variable explains the bulk of spatial variation in community level MUAC means across the sample area. This may also explain the good model fit despite the highly unbalanced nature of the panel. If spatial variation is so well captured by the explanatory variables, the absence of observations for certain communities will not have

¹³ Interested readers can obtain the full set of results from the authors.

as much effect of a negative effect on model fit and forecasting performance as missing a full set of observations for particular time periods. As we can see from Table 1, the unbalanced nature is fortunately more a product of few communities missing across numerous time periods rather than whole sets of communities missing at similar time.


Model in hand, the next step is to operationalize the resulting forecast estimates. How can food security managers or the relevant policy makers make the most use of the forecasts? Note that the forecast is a point estimate of predicted community MUAC means. Given a forecast estimate, one might want to know, for example, what the resulting 90% confidence interval for the true value is. Alternatively, one might want to know the degree of confidence with which he can claim that the true value lies below a certain critical threshold. In either case, if the objective is to gauge the intensity or magnitude of food insecurity, or to estimate the severity of child malnutrition, a potential weakness in our forecast is revealed.

Our model forecasts the mean MUAC Z-score for sample community children. The distribution of MUAC is likely to reflect within and between community inequalities in calorie and nutrient intake and may therefore vary between and within communities. MUAC distributions could also vary across time due to differential capacities of households to support their children's nutritional wellbeing during times of food stress. Children of particularly vulnerable households are likely to suffer higher rates of nutritional deficiencies and wasting in the face of shocks. This would show up in the data as a lengthening in the left-tail of the MUAC distribution, quite apart from a general leftward shift in the entire distribution that would result if the adverse nutrition effects of the shock were equally felt by all community children. As emergency response to widespread hunger is largely conditioned by the degree and prevalence of gross malnutrition, often defined by the proportion of children whose anthropometric conditions fall below a certain 'unacceptable' threshold (Howe and Devereux, 2004; UN Sub-Committee on Nutrition, 1999; World Food Programme, 2000), forecasting MUAC means may not be the most suitable strategy.

To better answer the appropriate question in this setting, we re-estimate our model with a dependent variable that more accurately specifies the left-tail of the MUAC distribution. We define the dependent variable as the fraction of children in each

community whose MUAC Z-score is less than -2. We set -2 as our threshold to be consistent with the benchmark often employed by key emergency relief agencies to define various levels of food stress and famine (Howe and Devereux, 2004; UN Sub-Committee on Nutrition, 1999; World Food Programme, 2000).¹⁴ Figure 6 presents full-sample and district level kernel regression of the fraction of children per community with MUAC Z-scores below -2 across time.

Figure 6



Not surprisingly, the proportion of children with MUAC Z-scores below -2 tells a largely similar story of the degree of food stress as that suggested by the results on MUAC means (Figure 1). 2000 was clearly a time of great deprivation as witnessed by the considerable proportion of children so malnourished that their MUAC Z-scores fall below 2 standard deviations of the reference mean. According to the scale of famine intensity devised by Howe and Devereux (2004), the vast majority of our sample would be experiencing ‘famine conditions’ which they define as 20% or more of children with Z-scores<-2. Indeed, in some areas of Turkana, the condition could be classified as a ‘severe famine’ (40% or more with Z-scores<-2).¹⁵ Differences across districts also mirror the pattern revealed by MUAC means, with Baringo and Samburu relatively better

¹⁴ Howe and Devereux (2004) offer a useful framework by which to define the intensity and magnitude of famine. The scaling system they develop uses various levels of wasting, defined by the proportion of children with Z-Score less than -2, as a key indicator variable (though the anthropometric measure they use is weight-for-height, standardizing into Z-Scores should make use of MUAC a valid substitute). Such a scaling system would be useful in conjunction with forecast estimates as a means to classify the estimated severity of food insecurity.

¹⁵ Note that the figures display averages over communities. During the year 2000, 35 out of 56 of our sample communities could have been classified as experiencing severe famine conditions for one or more months. Fourteen of these suffered severe famines for 6 or more months of 2000.

off. Under the Howe-Devereux classification scale, Marsabit, where children are generally more malnourished, sample communities could be said to have experienced ‘famine conditions’ from early 2000 to late 2004. By offering a universal scale by which we can interpret our results in relation to the severity and intensity of food insecurity faced by the target population, transforming our data from MUAC means into a proportion of the sample who fall below a critically defined threshold seems to be of greater use.

Table 4 presents key statistics for the forecasting model with proportion MUAC $Z\text{-score} < -2$ as its dependent variable. Other than simply substituting the relevant dependent variables and their lags, no other difference exists between the specification we employed in this case and that presented in Table 2. Again, we employed an LSDV estimator for both a one-month ahead and a three-month ahead forecast.¹⁶ While forecasting the left tail of the distribution offers a more accurate description of the condition of the most affected sub-set of the population, the results indicate MUAC mean forecasts are more accurate. While the overall R-squared of both the one-month and the three-month forecasts are slightly lower for the proportion forecast, RMSE, the more relevant performance indicator, favor the mean forecast. Though absolute RMSE, at 0.1205 and 0.1395 for the one-month and three-month models respectively, are lower for the proportion forecast, we must keep in mind that RMSE are presented in units of the dependent variable. As such with MUAC means ranging from -2.73 to 0.27 across the sample period as opposed to the zero to one range of MUAC proportions, it would seem that the forecasting performance for the means model is superior.

¹⁶ Note that in theory the dependent variable as a proportion is doubly centered at zero from below and one from above. In practice we find that left censoring is particularly prevalent with 27% zero-observations (Less than 1% of the observations are right-censored). This suggests using a panel tobit estimator. We fit a random-effects tobit model (a sufficient statistic allowing for the fixed effects to be conditioned out of likelihood does not exist) to our data but the resulting RMSE was larger. We thus opted for the more efficient LSDV estimator.

Table 4: Forecasting proportion with MUAC Z-score < -2: Model Statistics

		One-Month Ahead	Three-Month Ahead
R-squared:	within	0.5095	0.3391
	between	0.981	0.9022
	overall	0.7489	0.6054
RMSE		0.1205	0.1395
# of Observations		2290	2185
# of Groups		53	53
Observations per group:	Min	19	17
	Avg	43.2	41.2
	Max	58	56
Dependent Variable Statistics			
Mean:	-1.59	Std. dev:	0.40
		Min:	-2.73
		Max:	0.27

4.2 Testing Forecast Performance

To compare the forecasting performance of the two models, we move on to generate a series of rolling one and three month forecasts for both the mean and proportions models. We begin by generating forecasts for January 2004. As such, one-month ahead forecasts are based on model coefficients estimated on the subset of the data ending at December 2003 and the three-month forecasts are estimated on data up to October 2003. We then fold in an extra month of data, estimating one-month forecasts for February 2004 using data truncated at January 2004 and three-month forecasts on the subset of data ending on November 2004. We continue in this manner, generating monthly forecasts until May 2005. In Figure 7, we show the resulting one and three month forecasts for MUAC means in conjunction with the actual means. We first present the forecasts superimposed on the full period smoothed actual means and then highlight the forecasting errors by focusing only on the forecasted period and the unsmoothed actual means. Figure 8 does the same for the proportions model.

Figure 7


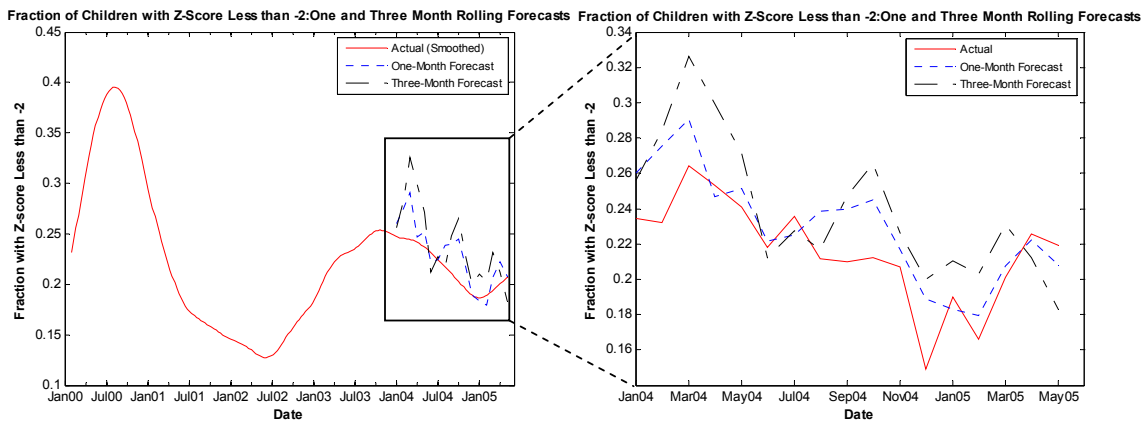


Figure 8



Two key points emerge from both Figure 7 and 8. First, it seems that there is a tendency for the model to systematically overstate the degree of malnutrition. For the means model, the vast majority of both one and three month forecasts indicate lower than actual means. Similarly, the proportions forecast consistently predict a larger fraction of children with under -2 MUAC Z-scores. As the forecast residuals for both models have roughly zero means, this suggests that forecasts are being understated along earlier ranges of the sample period. One explanation could be that the high degree of persistence in the model, implied by the modest deterioration in forecast accuracy as the horizon increases from one to three months, causes forecasts to respond slower to a marked turnaround in

MUAC levels.¹⁷ As such, the consistent decline in MUAC means witnessed from mid 2002 until late 2003 that tapers off just as we begin forecasting in January 2004 could explain the initially large overshooting of the forecasts which continue in the trajectory of past trends for a few months before correcting. This is of significance as it suggests that forecasting accuracy may diminish at particularly critical periods such as, for example, when a severe shock hits that causes a rapid deterioration in food security as was the case in early 2000. For emergency response and food security managers, such a time is when accurate information would be most valuable.

Nevertheless, as there have only been a few abrupt changes in MUAC trends over the sample period, the model may 'learn' better as more data are folded into it. Indeed, the second salient point revealed in Figure 7 and 8 is that the model seems to learn quite fast as more data is folded into this. In both cases, and for both forecast horizons, it is clear that forecasts track actual trends closer as we roll forward in time. While this could be a feature of the reasonably stable trends across the latter part of the forecasting period, a portion of the improved forecasts may be explained by learning. This highlights the value, not only of adding more data to the model, but also of assuring that data are contiguous as the highly unbalanced nature of the ALRMP panel as it currently stands undoubtedly reduces the models performance.

Another revealing feature of these forecasts is that despite the fact that the RMSE (when normalized for units) suggests that the mean model has superior forecast performance than the proportions model, this does not come across in comparing Figures 7 and 8. From Figure 7, the largest differences between actual and forecasted values, which seem to occur in March 2004, are about 0.07 and 0.14 Z-scores for the one and three month forecasts respectively. From Figure 7, forecast errors, which also appear to occur in March 2004, are approximately 0.02 and 0.07 percentage points for the one and three month forecasts respectively.

While this does not provide conclusive support for the relative performance of either model, especially as these full sample average results mask much of the structure of community level forecast errors, it nevertheless calls into question the superiority of

¹⁷ Moreover, the three lagged dependent variables for all four specifications were both statistically significant and had relatively large coefficients, especially for the earlier lags. This supports the hypothesis of considerable persistence.

the means model over the proportion model and begs for further investigation. Consequently, we further test the models by gauging how they perform based on the likelihood of correct responses that the models provide to questions they are likely to be called upon to answer.

For practical purposes, suppose that a food security manager wants to know the likelihood that a condition of 'famine' will prevail in a particular area. Given a particular forecast, he would then need to know the confidence with which the actual proportion of children with MUAC Z-score ≤ -2 was greater or equal to 20%.¹⁸ We generate such confidence levels for the one and three month forecasts of both the proportion and the means models. As none of the forecasting residuals for any of the models were normally distributed¹⁹ we employed percentile counts of the in-sample distribution of forecast residuals to generate our confidence levels. So, for example, if 10% of the sample residuals for the one-month proportion forecast had an error of greater than 0.15 percentage points, and the forecasted one-month ahead proportion was 0.05, then the confidence with which one could claim that in one month's time, 20% or more children will register MUAC Z-scores ≤ -2 , would only be 10%.

Extracting confidence levels allows for the defining of trigger points that when crossed would set in motion a series of predetermined response mechanisms (Barrett, 1997). Thus, for example, one could imagine a set of rules such that if famine is expected to arise with 66% confidence, food aid and other forms of emergency support are immediately deployed to the affected areas. Since such a coordinated response inevitably takes time as logistical, financial and bureaucratic concerns are sorted out, a dual-trigger system could be devised in which, at the crossing of a lower confidence band (say 33%), readiness measures are taken in preparation of a deployment should the second trigger be crossed.

We use such a decision making system to gauge the performance of the forecasts. First, we generate confidence levels for each forecasting horizon for both the means and

¹⁸ Following Howe and Devereux (2004), we define famine to be the condition at which twenty percent or more children have MUAC Z-scores ≤ -2 . For purposes of testing the performance of the means model, we employ an alternative MUAC means-based definition of famine that we arbitrarily set as the condition at which the mean level of MUAC Z-scores is less than -1.8. The model's performance should be robust to threshold chosen to define famine.

¹⁹ Shapiro-Wilks tests for normality of a distribution rejected normality of the forecasting residuals in each of the four cases.

proportions models. We then arbitrarily set 3 trigger points, one each at the 75%, 66% and 50% confidence levels. The trigger points can be thought of as the minimum confidence level that a policy maker requires to initiate a famine emergency response. Defining emergency response when there is actually a famine, or no response when there is no famine, as a ‘correct’ decision, we calculate the fraction of correct decisions that would result from utilizing the different models. Table 10 presents the results. We also calculate the fraction of wrong decisions that are Type 1 errors - the proportion of total wrong decisions that result from failing to respond when a famine actually occurs.

Table 10: Model Performance in Generating Correct Decision for Famine Response

Model Type	Forecast Horizon	Confidence Threshold		
		75%	66%	50%
		Fraction of Correct Decisions		
Proportions	One Month	0.777	0.786	0.785
	Three Month	0.753	0.756	0.758
Means	One Month	0.626	0.631	0.638
	Three Month	0.600	0.608	0.612
		Fraction of Errors that are Type 1		
Proportions	One Month	0.328	0.275	0.214
	Three Month	0.341	0.288	0.212
Means	One Month	0.215	0.180	0.120
	Three Month	0.227	0.194	0.134

The results clearly reveal the superiority of the proportions model in generating the correct decisions. The magnitude of differences in performance are quite striking. For all confidence thresholds, and both forecasting horizons, decisions based on the proportions model are about 15% more likely to be correct. Interestingly enough, the three-month proportions forecast also outperforms the one-month means forecast by around 8% under each confidence threshold. Further underscoring the results from Figures 7 and 8, the fairly small depreciation in performance as we increase forecasting horizon is promising and shows that the models can be used fairly accurately to give policy makers a sufficiently long three-month window to mitigate the consequences of impending shocks. The two-month gain in foresight is arguably worth the modest loss of accuracy.

Requiring higher degrees of confidence to guide the release of emergency resources for famine support does not necessarily result in more correct decisions. Though higher confidence thresholds do reduce the likelihood of taking costly measures to mitigate the consequences of a famine that does not materialize, it increases the frequency of making the arguably more costly mistake of failing to initiate the requisite response when famine actually occurs. As Table 10 presents the fraction of total error (wrong decisions) that are Type 1, all four forecasting regimes err on the side of responding to a wrongly declared famine than committing the mistake of failing to act when famine is real. Deciding how to distribute incorrect decisions between Type 1 and Type 2 errors is a normative judgment call best left to the policy arena. We nonetheless point out that lower confidence thresholds for decision making result in a smaller fraction of Type 1 errors. While the 50% confidence threshold results in the most correct decisions for all but the one-month proportions model, it does not follow that lower confidence thresholds always result in more correct decisions. The extreme example of a zero percent confidence level (and thus no Type 1 errors), in which it is evident that the fraction of correct decisions will fall considerably, suggests the existence of a strictly positive confidence threshold that optimizes the number of correct decisions.

Along with providing a more concrete way of defining nutritional deprivation that can easily be compared to a universally accepted classification scheme for the intensity and magnitude of famine, the superior performance of the proportions model in yielding correct decisions favors its adoption over the means model.

6 CONCLUSION AND POLICY IMPLICATIONS

Using a pseudo-panel of community level cohorts collected from primarily pastoralist communities selected across four districts in Kenya's arid north, we set out to develop an empirical forecasting model that could predict, with reasonable accuracy, the expected welfare impact of covariate shocks. We find that the joint dynamics of herd composition and herd management, climate and forage availability and food aid flows are able to forecast MUAC dynamics with impressive precision. Forecasting the proportion of children that fall below a critical MUAC threshold, the relevant MUAC parameter for assessments of the intensity and severity of a humanitarian crisis, yielded the particularly

accurate predictions. Moreover, offering policy-makers more response leeway by forecasting further into the future only marginally reduced forecast performance. On tests of comparative performance based on the likelihood of recommending the action that would have led to a ‘correct’ policy decision, three-month proportions forecasts, which consistently posted success rates of greater than 75% under three different decision-making regimes, was never more than 3% below the success rate of the one-month forecasts.

The policy implications are immediately clear. We show that it is possible to generate sufficiently accurate forecasts of a welfare indicator that is particularly sensitive to the covariate shocks that occasionally affect the target population. Moreover, the forecasts were generated from a relatively small set of variables that are not overly restrictive or costly to collect. In addition, reasonable accuracy is achieved despite deficiencies in the principle dataset used. If the ALRMP data were not so highly unbalanced, followed a more systematic sampling procedure, and provided a way to identify the claimed annual panel segments at the household level, forecast performance is likely to improve substantially. This suggests placing a premium on developing standardized collection procedures and failsafe methods for entering, identifying and storing data.

While several early warning and emergency assessment guides exist, our empirical forecasting method has the advantage of statistical rigor. Expected forecast accuracy can thus be calculated, allowing policy-makers to make more informed, confidence-guided decisions. Furthermore, once the model has been developed, it can easily and regularly be updated with new information, each time re-estimating the relevant parameters in a learning process that results in improved performance. Such a forecasting model is an invaluable tool for emergency awareness and response needs, offering rigorous, cost-effective and practical early warning capacity.

REFERENCES

- Alam A., Wojtyniac B., and Rahaman M. (1989), "Anthropometric indicators and risk of death", *American Journal of Clinical Nutrition* 49, 884-888
- Anderson, T. and C. Hsiao (1982), "Formulation and Estimation of Dynamic Models Using Panel Data", *Journal of Econometrics* 18, 47-82
- Arellano, M. and S. Bond (1991), "Some Tests of Specification for Panel Data: Monte Carlo Evidence and An Application to Employment Equations", *Review of Economic Studies* 58, 277-294
- Barrett, C. (1997), "Heteroscedastic Price Forecasting for Food Security Management in Developing Countries", *Oxford Development Studies* 25 (2), 225-236
- Barrett, C. and D. Maxwell, (2005), *Food Aid After Fifty Years: Recasting its role*.
Routledge: London, New York
- Chen L., Chowdhury, A. and S. Huffman (1980), "Anthropometric assessment of energy protein malnutrition and subsequent risk of mortality among preschool aged children", *American Journal of Clinical Nutrition*. 33: 1806-1845
- Coppock, D.L., Ellis, J., and D. Swift (1986), "Livestock feeding ecology and resource utilization in a nomadic pastoral ecosystem", *Journal of Applied Ecology* 23(2), 573-583.
- Deaton, A. (1985), "Panel data from time series cross-sections", *Journal of Econometrics* 30, 109-126.
- Ellis, J.E. (1994), "Climate variability and complex ecosystem dynamics: Implications for pastoral development" in Scoones I., editor, *Living with Uncertainty: New directions in pastoral development in Africa* (London: Intermediate Technology Publication, 1994)
- Ellis, J.E., and D. Swift (1988), "Stability of African pastoral ecosystems: Alternate paradigms and implications for development," *Journal of Range Management* 41(6): 450-469
- Funk, C., J. Michaelsen, J. Verdin et al. (2003), "The Collaborative Historical African Rainfall Model: Description and Evaluation." *International Journal of Climatology*, 23: 47-66

- Galvin, K.A., Boone R., Smith N., and S. Lynn (2001), “Impacts of climate variability on East African pastoralists: Linking social science and remote sensing.” *Climate Research* 19: 161-172
- Hansen, G.(2001), “A Bias-Corrected Least Squares Estimator of Dynamic Panel Models,” *Allgemeines Statistisches Archiv* 85: 127-140
- Howe, P. and S. Devereux (2004), “Famine Intensity and Magnitude Scales: A Proposal for an Instrumental Definition of Famine”, *Disasters* 28(4):353-372
- Kaitho, R., J. Stuth, J. Angerer et al. (2003), “A New System to Forecast Near-Term Forage Conditions for Early Warning Systems in Pastoral Regions of East Africa”, Research Brief 03-02-LEWS
- Lybbert, T., C. Barrett, S. Desta and D. Coppock (2004), “Stochastic Wealth Dynamics and Risk Management Among A Poor Population”, *The Economic Journal* 114 (498): 750-777.
- McKenzie, D. (2004), “Asymptotic theory for heterogeneous dynamic pseudo-panels”, *Journal of Econometrics* 120: 235-262
- McPeak, J. (2004), “Contrasting Income Shocks with Asset Shocks: Livestock Sales in Northern Kenya.” *Oxford Economic Papers* 56: 263-284.
- Nicks, A.D. (1974), “Stochastic generation of the occurrence, pattern, and location of maximum amount of daily rainfall”. p. 154–171. *In Proc. Symp. on Statistical Hydrol.*, Tucson, AZ. Misc. Publ. 1275. USDA, Washington, DC.
- Rubin, D.B., and N. Schenker (1986). “Multiple Imputation for Interval Estimation From Simple Random Samples With Ignorable Nonresponse”, *Journal of the American Statistical Association*, 81 (394): 366-374
- Stuth J., R. Kaitho, J. Angerer et al. (2003), “Integrating Information and Communication Technology for the Livestock Early Warning System (LEWS) in East Africa” Research Brief 03-01-LEWS
- UN Sub-Committee on Nutrition (UNSCN) (1999), “Refugee Nutrition Information System (RNIS), No. 26, Report on the Nutrition Situation of Refugee and Displaced Populations; see <http://www.unsystem.org/scn/archives/rnis26>

Vella V., Tomkins A., Ndiku J., Marshal T., Cortinovic I. (1994), “Anthropometry as a predictor for mortality among Ugandan children, allowing for socio-economic variables”, *European Journal of Clinical Nutrition* 48:189-197

World Bank Group (2003), “KENYA: Arid Lands Resource Management Project”, Project Information Document No. AB39

World Food Programme (2000), *Food and Nutrition Handbook*. WFP, Rome