Geostatistics, Basis Risk, and Weather Index Insurance

Michael Norton[1], Stephen Boucher[2], and Leslie Verteramo Chiu[3]

[1]Geography Graduate Group, University of California, Davis
[2]Agricultural and Resource Economics, University of California, Davis
[3]Dyson School of Applied Economics and Management, Cornell University

May 27, 2015

Abstract

This paper describes the application of geostatistics to weather index insurance in order to systematically analyze spatial basis risk inherent in index insurance contracts. The notion of spatial autocorrelation is in general overlooked by index insurance practitioners, but has profound implications for the effectiveness of the insurance offered. The analysis shows that it is possible to offer contracts from multiple weather stations to a single farmer, and that doing so will likely reduce the basis risk from a single contract. The two major implications of the paper are 1) that index insurance should be offered in more flexible contracts that allow farmers to hedge their production according to their perceptions of basis risk and their appetites for risk, and 2) the tradeoff between local (yield) correlation and spatial correlation needs to be more carefully considered, as it may even be better to offer contracts with poor yield correlation if they can include more spatial coverage.

Selected Paper prepared for presentation for the 2015 Agricultural & Applied Economics Association and Western Agricultural Economics Association Annual Meeting, San Francisco, CA, July 26-28.

1 Introduction

Weather index insurance removes some of the well-known problems with insurance products, such as adverse selection and moral hazard, and replaces them instead with the problem of basis risk. The basis risk in index insurance contracts arises from several sources, but is derived from the fact that a hedge on weather observations is not a perfect hedge for agricultural production at any given location. In addition, establishing new weather stations can be problematic since a fairly long time series of observations is needed to accurately price risk.

In short, basis risk is still a fundamental problem for index insurance practitioners. However, advances in other fields, especially spatial statistics, have still by and large not penetrated the index insurance literature. Highly visible papers in the past few years have either deliberately not confronted basis risk (Cole et al. 2013) or only sold insurance within a short distance of a weather station (Mobarak and Rosenzweig 2013 used 4 km.) Also, many practitioners attempt to create an index that is the best fit for the weather/yield relationship, but do not attempt to measure the spatial correlation. Although there have been some innovative solutions proposed (e.g. Elabed et al. 2013), a more careful study of the spatial factors of risk is needed for index insurance to achieve scale.

This paper outlines a framework for modeling spatial autocorrelation in index design. We demonstrate that basis risk is an endemic and fundamental problem of index insurance, and that practitioners should offer flexibility in their products so as to allow customers to manage basis risk according to their own preferences. As well, the design of weather indexes should more explicitly take spatial correlation and variability into account, and it may even be the case that an index that is a poor predictor of yield may be preferable if it has greater spatial coverage.

There are several advantages for a producer to hedge using multiple stations. It is straightforward and transparent to the insurance buyer to construct an estimation of the risk at a particular location by using existing weather observations. However, by writing a contract that utilizes a complicated mathematical function we are making a so-called "black box" which may not be accessible to the layperson. In the event of a dispute over the insurance payout, customers may not understand the exact reasons for the resolution of the insurance payment. By offering a hedge using existing weather stations, it is possible to offer more transparency to the contract using the same exact information that would be fed into a more complicated algorithm. Using the information to predict spatial autocorrelation will enable us to make a recommendation as to the exact proportions of insurance that should be bought at each weather station, but each hedger would ultimately be free to buy insurance in the proportions that they wish.

We demonstrate these ideas by designing several contracts with data from in and around the state of Iowa. The indexes are constructed with historical weather station data and are correlated with yield information from the Cornell University AgDB warehouse (Woodard 2014).

Offering suggestions for optimal hedges through space may not be applicable in all circumstances, but offers an additional method for hedging and insuring production in dispersed locations. The methods presented in this paper in an effort to introduce some of the empirics of geostatistics into the index insurance literature, and although the basics are presented here, it is possible to build on this work with additional advances from the geostatistics literature.

2 Literature

2.1 Basis Risk

Despite the fundamental and troubling presence of basis risk in index insurance contracts, relatively few studies have attempted to analyze basis risk with more specificity. In an early effort, Woodard and Garcia (2008) split basis risk into two distinct components: 1) "local" basis risk, or the imprecision in the weather/yield relationship, and 2) "spatial" basis risk, or the imprecision that results from distances away from weather stations. Elabed et al. (2013) propose the possibility of multi-scale insurance to avoid the basis risk inherent

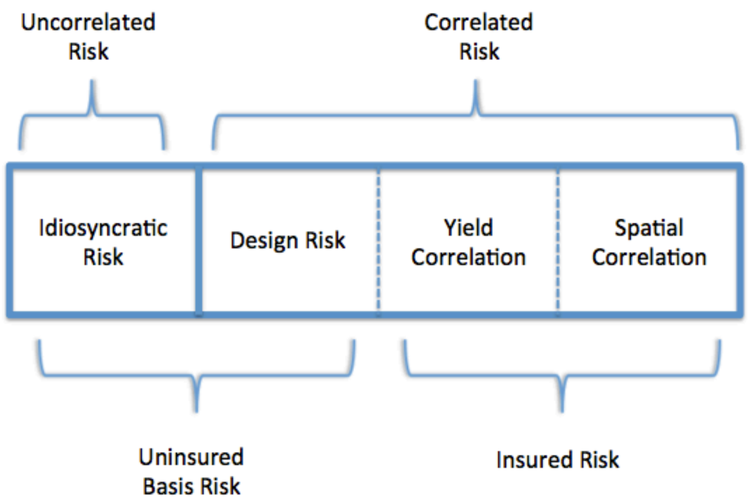

Figure 1: A decomposition of the index insurance contract.

in a single-trigger contract.

Figure 1, adapted from Elabed et al. (2013), illustrates the different types of basis risk present in the index insurance contract. The terms are defined as follows. Idiosyncratic risk is the risk that is outside the bounds of the index insurance contract. Some examples are pest damage or other, uncovered adverse weather events like hail. Correlated risk is the risk that is correlated with the index insurance payouts.

Where this figure departs from Elabed et al. (2013) is that it splits the insured risk into the two components: yield correlation and spatial correlation. Normally, the goal for designers of index insurance contracts is to correlate weather and yields in order to minimize the amount of Design Risk, which is correlated risk that is outside the insurance contract. Typically, index constructed to minimize design risk are very specific, detailed combination of one or more weather variables that describes the index/yield relationship at known points. The amount of spatial correlation will depend on the distance from a weather station, which is more explicitly rendered in the Empirical Model section below.

However, these highly specific indexes are likely to have worse spatial correlation that a simpler index, and there is a tradeoff between design risk, yield correlation, and spatial correlation. It may be that a simpler index with higher spatial correlation will provide better coverage than a product that is highly specific to a certain area.

Also of note is that the fundamental nature of basis risk also supports recent calls for risk layering in index insurance projects, with several recent papers investigating the role of index insurance alongside informal risk-sharing strategies (Mobarak and Rosenzweig 2013, Dercon et al. 2014).

2.2 Geostatistics and Spatial Econometrics

To model correlations through distance, we can use spatial statistics, using the philosophy of Tobler's First Law of Geography: "Everything is related to everything else, but near things are more related than distant things" (Tobler, 1970). The basic element of analysis for geostatistics is the semivariogram, which diagrams how spatial autocorrelation decays with distance. A sample semivariogram for June precipitation totals in and around Iowa is pictured in Figure 2.

In a semivariogram, the squared difference between the values of each pair of locations is plotted on the y-axis, and the distance between the points on the x-axis. The spatial autocorrelation will decrease with distance, which causes the difference in values between points (as pictured on the semivariogram) to rise. The observed points may be fitted with a concave function - the semivariogram model. The semivariogram

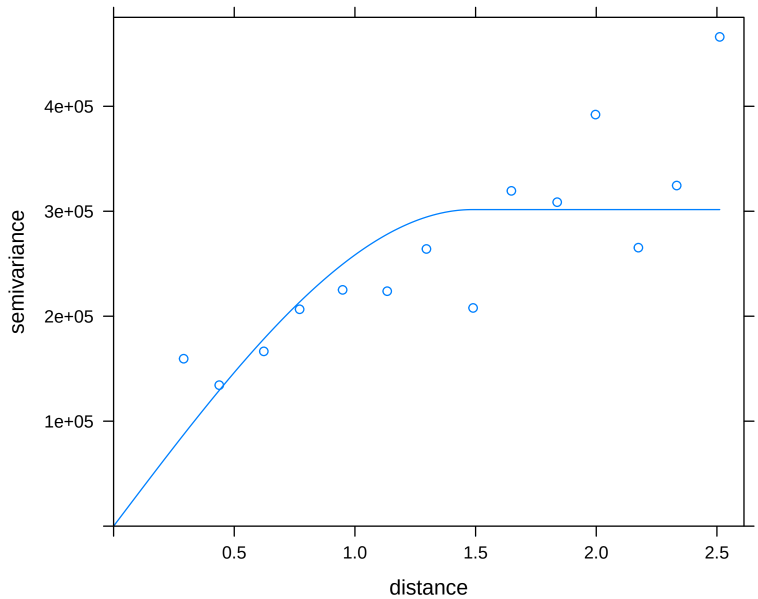

Figure 2: A sample semivariogram.

model will provide an estimate of spatial autocorrelation at each distance, from use in a linear regression to correct for that autocorrelation between data points.

A semivariogram typically contains several elements:

1. **Nugget** - The nugget is the covariance at distance zero. Normally, we would expect the covariance at distance zero to be zero, as comparing each point to itself should not create any different in values. The nugget effect accounts for phenomenon such as measurement errors, or differences in correlation at distances below which that are measured. Rainfall measurements in particular have a large nugget effect, as there are often large measurement errors (Villarini et al 2008).

2. **Sill** - The sill is the level at which the semivariogram flattens out. The sill is the variance of the sample, and the modeled covariance in the semi-variogram will eventually reach the sill, which is essentially the default value for uncorrelated pairs of points.

3. **Range** - The covariance model extends in a upward fashion to a point where it reaches the sill, where the spatial autocorrelation is no longer considered to be present. The range of spatial autocorrelation is a function of distance, and datasets with more spatial correlation will extend the range out further. For index insurance, the choice of index will impact the range of the semivariogram. It is important to understand where this range exists for each index, and how far from a station a project should be comfortable selling insurance.

The spatial correlation model may be used in a linear regression as off-diagonal elements of the covariance matrix as in econometric studies like Newey and West (1987). The formulation should be familiar as one possible manifestation of White's heteroskedastic-robust standard errors (White 1980).

Interpolation methods like kriging use the spatial autocorrelation model fitted to the semivariogram to derive a vector of linear weights for each point in space. The vector of linear weights allows us to recreate a value for unknown locations using a weighted combination of information from nearby stations.

In this way, the usual application of the geostatistics to provide an estimate at spatially diverse locations. However, in the insurance market, The variance can be calculated and pictured in a spatial dimension, such as in Figure 3. The variance at each location in space has a characteristic bulls-eye pattern, with the variance very low around the sampled points, and decaying with distance according to the spatial autocorrelation modeled in the semivariogram.

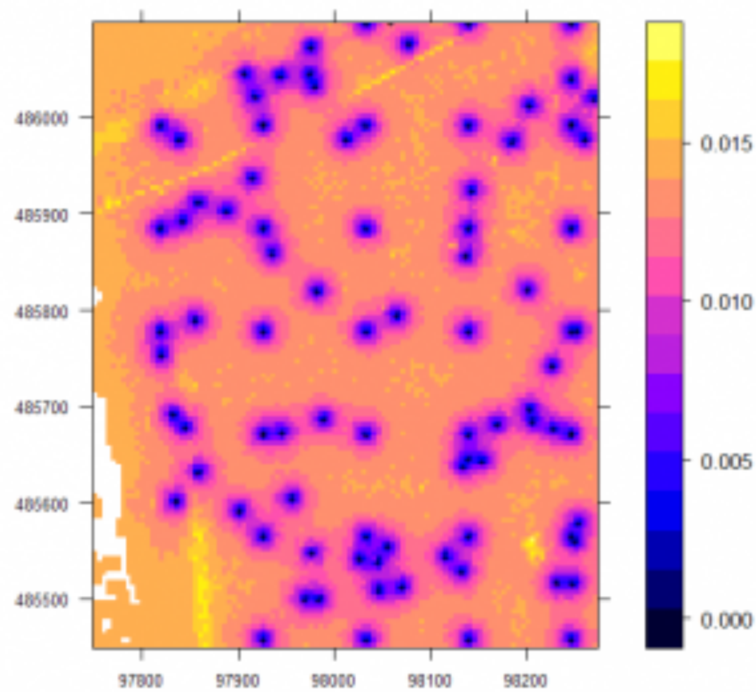

Figure 3: The bulls-eye pattern of spatial autocorrelation around sampled data points.

The spatial autocorrelation pictured in Figure 3 [1] does not have a very wide range, which is to say that the autocorrelation decays rapidly from each sampled point, and there are large gaps in between for which the sampled locations provide no additional information. In practice, there may be more overlap between sampled data points, so that the estimate of the value at a given location will depend on two or more points from the sample.

In index insurance terms, if each of those sampled points was a weather station, the range of the semivariogram would indicate the distance to which insurance could possibly be sold, with better correlations closer to sampled points (weather stations). Additionally, if a farm is within the range of spatial autocorrelation for more than one weather station, it may be prudent to offer insurance which is a combination of contracts from each station within range.

[1] Source: spatial-analyst.net

2.3 Relationship between geostatistics and Miranda's β

There are several concepts in the insurance and finance literatures that are related to the concepts presented above. Miranda (1991), on the subject of area yield insurance, introduced a β_i term intended to represent the individual difference of each term from the mean production in the area. If the localized farm was above or below a critical value β_c, the farm should buy the area yield insurance.

Compare the equation including β_i form Miranda (1991):

$$\tilde{y_i} = \mu_i + \beta_i \cdot (\tilde{y} - \mu) + \tilde{\epsilon_i} \tag{1}$$

The textbook definition of the simple kriging estimator (Goovaerts 1997, pg. 127) is identical to the above equation (except without the $\tilde{\epsilon_i}$ term.) The simple kriging estimator is just one of many interpolation algorithms in the kriging family. Simple kriging is, as its name suggests, only the most simple of the kriging estimators, which may be extended to multiple values in co-kriging (similar to multiple regression), with a surface trend model, or used non-parametrically in indicator kriging. Similarly, the search area may be modified to only consider the relationship of data points in one direction, or in an ellipse instead of a circle.

One difference between kriging and Miranda's β is in the purpose of the esimation. Miranda's β was intended to recover the yield at a given location. Kriging estimators are a linear regression on the covariance terms, which, as stated previously, provides a vector of linear weights. The weights will indicate the proportions of each of the known data points should be used to reconstitute value at the unknown location.

2.4 Relationship between geostatistics and the Optimal Hedge Ratio

Johnson (1960) and Stein (1961) introduced the risk-miniizing hedge ratio for futures markets based on theories of efficient portfolio selection of Markowitz (1952). Gagnon et al. (1992) extend the theory of the optimal hedge ratio into n different multiple markets, such as currency markets. The formula Gagnon et al. (1992) derive for the optimal hedge ratio is identical to both linear regression with heteroskedastic errors (Newey and West 1987) and the weights for the spatial smoothing process known as kriging. The similarity is not a coincidence, as each method seeks to minimize the variance in the prediction through the use of linear regression with a spatially consistent covariance matrix.

Although the Optimal Hedge Ratio addresses price risk and weather index insurance hedges production risk, the concept of hedging via multiple imperfect instruments is the same. To draw an analogy, if a sorghum farmer is not able to hedge his price risk through a futures contract, she may be able to hedge some of the risk through a linear combination of the corn and soybean markets. Likewise, a farmer in a given location could use a linear combination of weather stations to construct an insurance product tailored for that location.

More formally, in commodities futures markets, there is a well-known phenomenon known as "basis," defined as the discrepancy between price movements in the futures market and the local cash market. The optimal hedge ratio is derived from the variances and covariances of the local cash market and the futures market. The simplest method for determining the optimal hedge ratio is in the case of just one market for hedging production, when there is a local cash market for the good as well as a futures market for hedging. The expected return on a hedged position (R_h) is given by the return in the cash market R_c as well as the hedged return in the futures market (R_f).

$$E(R_h) = E(R_c) + hE(R_f) \tag{2}$$

To minimize the variance of the hedge, σ_h^2, by looking at its constituent parts, the variance of the cash market σ_c^2 and the variance of the futures market σ_f^2. The h variable represents the optimal percentage of the cash position to hedge in the futures market. h is often a number between 0 and 1, but can be greater than one.

$$min_h \sigma_h^2 = \sigma_c^2 + h\sigma_f^2 \tag{3}$$

The first order conditions:

$$\sigma_c^2 + h^2\sigma_f^2 - 2h\sigma_{cf} = 0 \tag{4}$$

For which we solve for optimal hedge ratio, h^*:

$$h^* = \frac{Cov(\sigma_c, \sigma_f)}{\sigma_f^2} = \rho\frac{\sigma_c}{\sigma_f} \tag{5}$$

Where ρ is the correlation between cash and futures prices. Equation 7 is intimately related to least squares regression, and in fact Equation 7 is identical to the equation for the slope coefficient of a least squares regression with one independent variable. The calculation is identical to regressing the time series of cash market prices with futures market prices from the same time period.

Gagnon et al. (1998) extend the theory to hedging a single transaction in multiple futures markets when no futures market exists for the commodity in question. One example of such a transaction would be for a currency for which there exists no futures market, in which case you might hedge in several closely related futures markets. In the case of a multi-market hedge, the optimal hedge ratio is a function of the correlation between the cash market and each futures market, as well as the correlation between each futures market.

$$h^* = \Sigma_{ii}^{-1}\Sigma_{li} \tag{6}$$

Where Σ_{li} is a vector of the correlations between the cash market and each futures market, and Σ_{ii} is an *n x n* matrix of the correlation of the insurance payouts at various weather stations to one another. The result fo Equation 6 provides a vector of linear weights for each hedging instrument, which will indicate the proportion of each instrument to buy. Instruments which are highly correlated to the hedged quantity will be weighted heavily, and the weights will be diminished by both poor correlations and high cross-correlations with other hedging instruments.

Gagnon et al. (1998) demonstrate how Equation 6 could, with a few assumptions, be used in a utility maximization framework. Those assumptions include a martingale, so that future performance is not predictable from past results, and that the agent is a mean-variance utility maximizer. Perhaps unsurprisingly for a utility function that maximizes utility by minimizing variance, Equation 6 is very similar in formulation and derivation to the ordinary least squares estimate $b = (X'X)^{-1}X'y$, defined as the best linear unbiased estimator (BLUE). Equation 6 is also identical in formulation and derivation to kriging estimators.

3 Empirical Model

We propose a model for estimating the amount of insured risk at any location via a linear model with spatial covariance matrix.

$$\beta^* = (X'\Sigma X)^{-1}X'\Sigma y \tag{7}$$

Where $(X'\Sigma X)^{-1}$ is the variance/covariance matrix for the weather stations, and the $X'\Sigma y$ term captures the covariance between each weather station and the point of interest. The end result of the regression, the β terms, are a vector of linear weights for each station according to how much weight should be placed on each station. For locations that are near to a weather station, that station will likely dominate the weights and will perhaps even be the only station that is considered in the model. or locations that are between weather stations, this model will help determine the weights to place on each weather station according to a spatial autocorrelation model as represented by the Σ term.

Conceptually, starting with Figure 1, we can extend the conception of basis risk to include distances from a weather station. Specifically, the amount of spatial correlation will depend on the distance of the farm from the weather station. At the weather station (distance=0), if we assume perfect information the spatial correlation will be perfect, so that the uninsured basis risk will include only idiosyncratic risk and design risk. (If we don't want to assume perfect information we can incorporate the "nugget" effect into our spatial model, described in the Geostatistics and Spatial Econometrics section above.)

However, as distances increase from the weather station, a new term is added to the uninsured basis risk - spatial basis risk. The inclusion of spatial basis risk is conceptualized in Figure 4. In Figure 4, the insured risk (green area) decreases as the distance from the weather station increases, with the difference made up by spatial basis risk (the pink area). The spatial basis risk will increase according to the spatial autocorrelation model used, as analyzed using the semivariogram of the index in question. The green/pink areas are drawn as triangles in Figure 4, but this conception would only be appropriate for a linear model of spatial autocorrelation. However, spatial autocorrelation models can have many shapes, and are usually curved to account for decreasing marginal correlation.

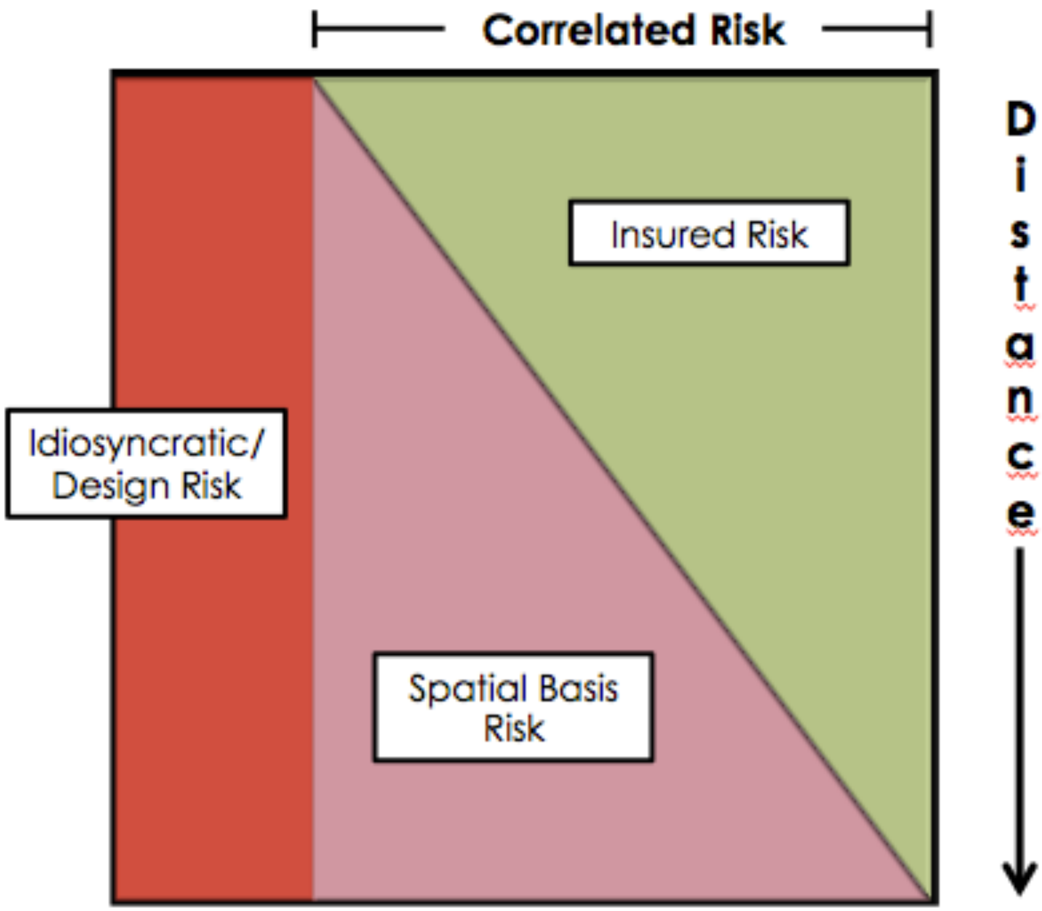

Figure 4: Conceptualizing the index insurance contract with spatial basis risk.

Different indexes will have different levels of local correlation as well as different levels of spatial autocorrelation. In this diagram, if an index is a better predictor of yields, the horizontal width of the insured risk will be wider. Likewise, if an index has a higher spatial correlation the green polygon will extend further, but if an index does not have a high spatial correlation then the green polygon will disappear at a short distance from a weather station.

The model that we use in our analysis is not a linear model, but rather the spherical model (Cressie 1985), chosen because it resolves to a horizontal sill at a given range.

$$\gamma(d) = \begin{cases} 1.5\frac{d}{r} - 0.5(\frac{d}{r})^3 & \text{if d < r} \\ \sigma^2 & \text{o.w.} \end{cases} \quad (8)$$

Where d is the distance from the weather station to the farm, and r is the range of the variogram. In practical terms, the modeled semivariance increases through a concave function to a given point r, at which it resolves to the background variance σ^2. The fit for the variogram function was computed by using the R package 'gstat' (Pebesma 2004).

For this iteration of the research, we consider each year of weather observation is a random draw from a distribution, although this approach neglects the possibility of long term trends like decadal processes or climate change. These will be issues to revisit in future research, however, by treating each year as a random draw, we may start to understand the differences in spatial correlation between the different indexes.

4 Results

To test our theory, we analyzed weather data from the Global Historical Climatology Network (GHCN) for the years 1901-2014 in and around the state of Iowa. The weather stations in question are mapped in Figure 5. Iowa was chosen for analysis for its heavily agricultural landscape as well as relatively homogenous topography. GHCN COOP data is available from the National Centers for Environmental Information of the National Oceanic and Atmospheric Administration (NOAA)[2].

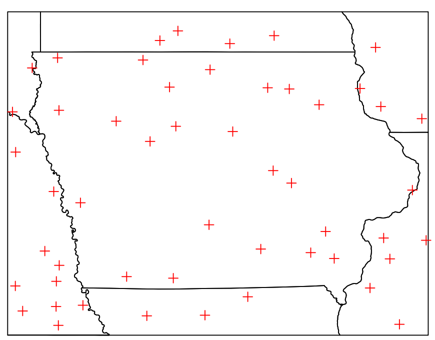

Figure 5: Location of GHCN weather stations in and around Iowa.

To illustrate the differing spatial characteristics of different types of indexes, we calculated two indexes from the weather observations. First, Heating Degree Days (HDD) for each station for the months June-August (strike value 29 degrees C) to provide a measure of heat stress for the crops. The value of 29 degrees C was chosen based on the evidence of Schlenker and Roberts (2006), who note the non-linear effects of excess temperature on corn yields. Second, we calculated the total precipitation for each weather station for the month of June.

The June precipitation was shown to have some spatial correlation, with the ranges for each year pictured in Figure (histogram). The mean range for the variogram was calculated at 123 km, and the median range as 117 km. For the HDD index, there is very little spatial correlation, with a mean value of 110 km and a

[2]ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily/

median value of 25 km. The median and the mean are so different because there are a large number of years in which the variogram did not converge to a solution, and the was given a value of 0.

The lesser spatial correlation of the temperature index may be surprising, given the higher spatial correlation of temperature values (Norton et al. 2012), however, when arranged into an HDD index the spatial correlation is much worse than that of the precipitation totals.

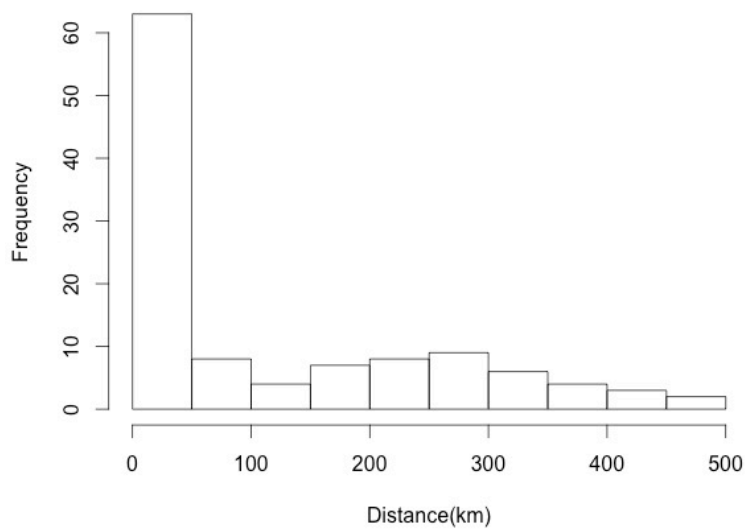

Figure 6: Ranges for HDD totals.

5 Next Steps

Geostatistics is a rich field in the literature, and there are several potential improvements to the application of geostatistics to index insurance. For one, there are many possible functions to use in the semivariogram model, each of which will produce different results. Future research should test the applicability of different variogram shapes.

Second, it should be possible to write a computer program to solve for the optimal correlation for the yield/index relationship. Given the nearly infinite number of combinations for a weather index, it will be necessary to apply an algorithmic approach to sift through the various possibilities. The optimaization program should be run both with and without a spatial component to show the effects of spatial optimization.

Third, there are types of spatial interpolation that may be of interest to index insurance researchers. In particular, the concept of indicator kriging (Solow 1986). Indicator kriging is a non-parametric approach that provides a probability at each given location that a measured value is over a defined threshold. The possibility of spatial analysis with a threshold is promising for index insurance because of the presence of the contract trigger at some definied threshold that is intended to represent yield loss due to weather patterns. Thus, indicator kriging could provide a probability of weather observations being over a given threshold at each distributed spatial location.

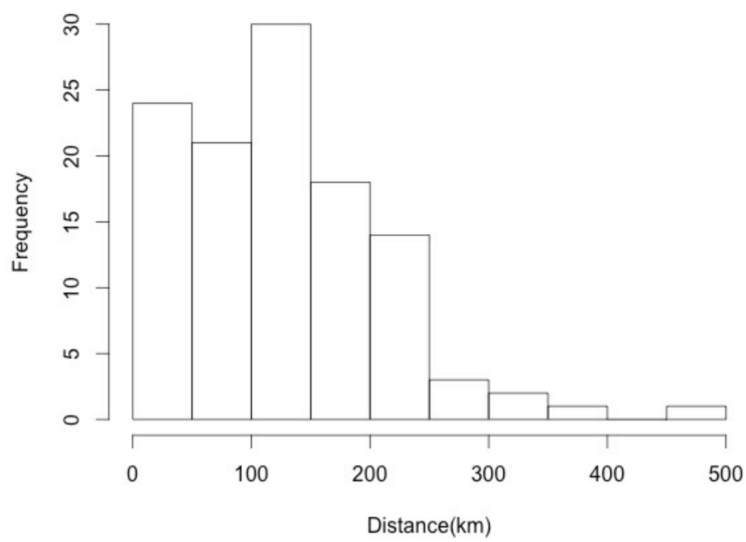

Figure 7: Ranges for precipitation totals.

6 Discussion and Conclusion

The above discussion has two implications. First, it will almost never be optimal to buy index insurance to hedge the full amount of production as there will always be some level of basis risk. Index designers should not assume that the optimal strategy for hedging is for the full value of production, because the hedge is by nature imperfect. Because the optimal amount of insurance will be less than the total production, insurers should not offer a single contract to insure the full value of production. Instead, insurers should offer farmers flexible contracts that insure different fractions of production so as to allow their customers to approach basis risk as they see fit. One such example is the HARITA project described in Norton et al. (2014), where the insurance was offered in five different increments of expected production (from 20

The second implication is that index designers face a tradeoff between index accuracy and spatial basis risk. Practitioners often attempt to devise indexes that are the best predictors of yields at known locations, but these very specific indexes might not be valid for more than a few kilometers from a weather station. For example, the occurrence of rainfall on any particular day has a higher spatial correlation than the amount of rainfall (Greatrex and Grimes 2012), but the amount of rainfall is a better predictor of the yield. In some cases, an index based on a weak local predictor like rainfall occurrence may be preferable to indexes based on stronger local predictors like rainfall amounts if the spatial correlation dominates the yield correlation.

However, it is important to note that this paper only addresses basis risk in the spatial dimension, and there may even be sources of uncertainty that are subjective in nature. For example, another potential source of covariance from the perspective of the buyer of index insurance is trust that the insurance company will make a payment when they say they will, which is a subjective assessment of the covariance in the contract. Trust is an issue that has been identified many times in the literature as a barrier to take-up of index insurance in developing countries (e.g. Cole et al. 2013), especially among people who have limited experience in financial transations.

Thus, while there may be other considerations the consumers should include when making the decision to purchase insurance, this analysis limits itself only to the case of spatial variability in order to take advantage of existing methodologies in the spatial statistics literature. The estimates therefore should be considered something of an upper bound of the amount of insurance that would minimize the variance of the hedged portfolio.

What is unknown, however, is whether or not farmers perceive spatial basis risk and make their purchase decisions accordingly. Demand for index insurance a topic that is still a popular topic in the literature and actively under investigation by projects around the world, and farmer perceptions of the usefulness of the insurance could have profound implications on demand for index insurance. An empirical test of whether or not farmers that are closer to stations (i.e. encounter less basis risk) do in fact buy more insurance would likely be of interest for academics and index insurance projects.

Further, the application of index insurance is not only for its ability to cover downside losses, but also in its ability to encourage productive risks. Farmers' behavior is likely to be excessively conservative if they are near a point from which they cannot recover. The analysis in this paper is only for the case of hedging, but the model could be extended to include the possibility of additional productive activities.

References

Clarke, Daniel J. A theory of rational demand for index insurance. Department of Economics, University of Oxford, 2011.

Cole, Shawn, Gine, X., Tobacman, J., Topalova, P., Townsend, R., Vickery, J. "Barriers to household risk management: evidence from India." American Economic Journal. Applied Economics 5.1 (2013): 104.

Cressie, Noel. "Fitting variogram models by weighted least squares." Journal of the International Association for Mathematical Geology 17.5 (1985): 563-586.

Dercon, Stefan, et al. "Offering rainfall insurance to informal insurance groups: evidence from a field experiment in Ethiopia." Journal of Development Economics 106 (2014): 132-143.

Driscoll, John C., and Aart C. Kraay. "Consistent covariance matrix estimation with spatially dependent panel data." Review of Economics and Statistics 80.4 (1998): 549-560.

Elabed, Ghada, et al. "Managing basis risk with multiscale index insurance." Agricultural Economics 44.4-5 (2013): 419-431.

Gagnon, Louis, Gregory J. Lypny, and Thomas H. McCurdy. "Hedging foreign currency portfolios." Journal of Empirical Finance 5.3 (1998): 197-220.

Goovaerts, Pierre. Geostatistics for natural resources evaluation. Oxford university press, 1997.

Greatrex, Helen, and David I.F. Grimes Capturing the spatial variability of rainfall in weather-based index insurance. In The challenges of index-based insurance for food security in developing countries. Proceedings of a technical workshop organised by the EC Joint Research Centre (JRC) and the International Research Institute for Climate and Society (IRI, Earth Institute, Columbia University), JRC Ispra, Italy, 2 and 3 May 2012. R Gommes and F. Kayatikire eds. (2012)

Johnson, Leland L. "The theory of hedging and speculation in commodity futures." The Review of Economic Studies (1960): 139-151.

Markowitz, Harry. "Portfolio selection." The journal of finance 7.1 (1952): 77-91.

Miranda, Mario J. "Area-yield crop insurance reconsidered." American Journal of Agricultural Economics 73.2 (1991): 233-242.

Mobarak, Ahmed Mushfiq, and Mark R. Rosenzweig. "Informal risk sharing, index insurance, and risk taking in developing countries." The American Economic Review 103.3 (2013): 375-380.

Newey, Whitney K., and Kenneth D. West. "A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix." (1986).

Norton, Michael T., Calum Turvey, and Daniel Osgood. "Quantifying spatial basis risk for weather index insurance." Journal of Risk Finance, The 14.1 (2012): 20-34.

Norton, Michael, et al. "Evidence of demand for index insurance: experimental games and commercial transactions in Ethiopia." Journal of Development Studies ahead-of-print (2014): 1-19.

Pebesma, Edzer J. "Multivariable geostatistics in S: the gstat package." Computers and Geosciences 30.7 (2004): 683-691.

Schlenker, Wolfram, and Michael J. Roberts. "Nonlinear effects of weather on corn yields." Applied Economic Perspectives and Policy 28.3 (2006): 391-398.

Solow, Andrew R. "Mapping by simple indicator kriging." Mathematical Geology 18.3 (1986): 335-352.

Stein, Jerome L. "The simultaneous determination of spot and futures prices." The American Economic Review (1961): 1012-1025.

Tobler, Waldo R. "A computer movie simulating urban growth in the Detroit region." Economic geography (1970): 234-240.

Villarini, Gabriele, et al. "Rainfall and sampling uncertainties: A rain gauge perspective." Journal of Geophysical Research: Atmospheres (19842012) 113.D11 (2008).

White, Halbert. "A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity." Econometrica: Journal of the Econometric Society (1980): 817-838.

Woodard, J.D., and Philip Garcia. "Basis risk and weather hedging effectiveness." Agricultural Finance Review 68.1 (2008): 99-117.

Woodard, J.D., AgDB Data Warehouse, Cornell University, 2014, Available online at: http://128.84.12.99/agriskmanageme