SOCIETY OF ACTUARIES

Predictive Analysis: The Effects of Technology and Weather on Crop Yield

December 2019

SOCIETY OF ACTUARIES

Predictive Analysis: The Effects of Technology and Weather on Crop Yield

AUTHORS Lysa Porth, PhD, MBA
Mitchell Roznik, M.Sc.
Ken Seng Tan, PhD, ASA, CERA
Wenjun Zhu, PhD, FSA, CERA
C. Brock Porth, Ph.D., P.Eng.

SPONSOR SOA General Insurance Research Committee

CONTENTS

Predictive Analysis: The Effects of Technology and Weather on Crop Yield

Section 1: Introduction

Crop insurance is often delivered in partnership with governments and private industry. In Canada, 10 provincial government crop insurance companies develop and deliver crop insurance, with the federal government's oversight. The provincial and federal governments together subsidize the insurance premium by 60% and equally share in all administration costs. Each provincial government crop insurance company sets the premium rate, with the federal governments actuarial department's approval. In the U.S., as a comparison, approved private insurance companies deliver crop insurance, with the governments Risk Management Agency developing and pricing the policies. The Standard Reinsurance Agreement governs the relationship between the government and insurance companies, with the government subsidizing approximately 62% of the actuarily fair premium and covering all administration costs.

A necessary calculation in the crop insurance policy is establishing the Probable Yield (PY), which is typically based on the historical average yield at the farm or field level. The PY also sets the insured value (liability) by multiplying it by the selected coverage level and the average crop yield. Depending on the insurer, a time series of approximately 10 years is often used to calculate the PY, based on a simple average. In part, the decision on how much data to use depends on the availability and credibility of the data at the farm level, as well as program regulations. In agriculture, a particular challenge exists regarding data sparsity. In most regions, there is only one growing season each year and, hence, one data observation each year. Further, producers adapt crop rotations over time depending on commodity market price signals, soil conditions, etc. In some cases, crop types may be removed entirely from rotation, and this leaves gaps in the yield history to calculate the average crop yield. From a statistical inference point of view, there is a desire to utilize as much of the data as possible, but this is balanced against the need to ensure the data are representative of the current situation. This includes such considerations as changes due to programs, farming practices, biotechnology and weather experience, for example. As such, there is often a tradeoff between using a shorter time series and having a more responsive and representative PY or using a longer time series and having a more stable PY. In addition, actuaries also need to be very careful in using shorter time series, since if weather is treated as stochastic, shorter time series may not be representative.

To help overcome this challenge, some research has focused on restatement methods to adjust older observations so they are more comparable to current technology and growing conditions. It is well documented that yields for many crop types are increasing through time (Sherrick et al., 2004a; Lobell et al., 2011). There may be several contributing factors to increasing crop yield over time, and one important factor is technology. Under the hypothesis that technology positively contributes to yield over time, not accounting for the impact of technology may result in a PY that is too low. In this case, a farmer would not be able to receive adequate risk protection (i.e., coverage would be too low) (Skees and Reed, 1986; Coble et al., 2013). For example, consider Figure 1, which shows a hypothetical canola grower. Each year, if the yield improves due to technology (assumed to be an improvement of 1 bushel per acre, or bu/acre, per year), the PY will underestimate the producers' expected crop yield. In this example, the PY based on the simple average over 10 years using unadjusted yields is 37 bu/acre. However, if a linear trend is applied to the 10 years of historical data, then comparatively the trended PY is 42 bu/acre, resulting in a deficit of 5 bu/acre. This risk deficit is shown in Figure 1, and the yield deficit is shown in red. Similarly, in consideration of the belief that technology improvements are making crops more resilient, the risk profile of the restated losses may be reduced (Tack et al., 2012). If the yield trend is ignored in crop insurance, there may be negative welfare effects (Adhikari et al., 2012) and rating inefficiencies (Woodard et al., 2011).

Figure 1
Canola Yield 10-Year Technical Change

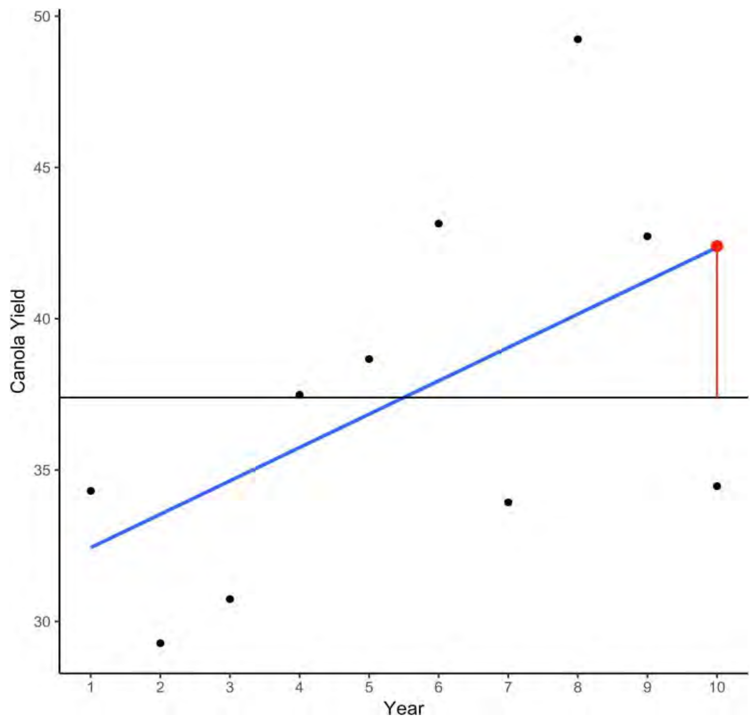

With a short yield history, it is difficult to determine the portion of yield gain attributable to technology and/or a positive weather trend over the period. Lobell and Asner (2003) investigated weather effects on yield trend for corn and soybeans in the U.S. over a 17-year period and reported that approximately 25% of corn yield trend and 32% of soybean yield trend can be explained by favorable temperature effects in the region of interest. Based on this finding, attributing 100 percent of yield gain due to technology over the period may result in an upward bias of insurable yields over the period. Some research has further studied the impact of specific weather variables on yield gains. For example, Tollenaar et al. (2017) showed that solar brightness over the period from 1980 to present has been increasing and that this increase, called a brightening effect, explains 27% of corn yield increase in the Midwestern U.S. This brightening effect may be linked to improved air-quality standards, resulting in less resistance in the atmosphere to incoming solar radiation over these growing regions. With a long history of yield data, much of these weather effects may balance out, but there may be some long-term prevailing changes that affect yield. These longer-term effects are likely fairly crop and region specific. With a shorter yield history, it is necessary to manage the effects of weather on crop yield to approximate the yield trend. This can be done by directly modeling weather effects on yield so the technology yield trend can be estimated, or, alternatively, a statistical-based method can be used to reduce the noise in the data introduced by weather effects. Therefore, a longer time series of yield is preferred because yield variation caused by growing season weather conditions becomes less influential, and approximating the yield trend becomes less sensitive to extreme low or high yields caused by growing season weather conditions.

In regard to premium ratemaking, researchers have attempted to evaluate whether the employed premium ratemaking methodology used for the federal crop insurance program in the U.S. is appropriate, given increasing crop yields (Woodard et al., 2011). They found that for the federal crop insurance program to provide unbiased rates, constant relative risk of crop yields must be assumed; however, some researchers believe as crop yields rise, crop risk decreases (Yu and Babcock, 2010; Woodard, 2014). Recent studies have assessed the impact of sample period length on premium rate setting to evaluate whether shorter historical yield sample periods are more effective for rate setting over longer histories by using more recent yield histories that better reflect current crop technologies at the expense of a less representative sample of weather events (Woodard, 2014). They found that the less representative sample, 30 years, of weather experience performed similarly to the more representative sample of weather events. Several methods have been proposed in literature for estimating yield trend; however, they are often designed based on approaches that use county level data or for several farms grouped together (Zhu et al., 2011; Ramirez et al., 2003; Just and Weninger, 1999; Finger, 2010; Goodwin and Mahul, 2004; Sherrick et al., 2004b; Ozaki et al., 2008; Ker and Goodwin, 2000). For crop insurance purposes, estimating the yield trend at an aggregate level may lead to inaccuracies, because high-technology adopting farms may receive a smaller yield trend adjustment compared to their actual yield trend, and low-technology adopting farms may receive a larger yield trend adjustment compared to their actual yield trend.

Less research has focused on yield trend estimation for individual farms, which is important for calculating the insured value (liability) for crop insurance. This could partially be due to the sparsity of farm-level data, which can make using farm-level data more challenging. As well, more generally, the ratemaking method for crop insurance uses a loss cost ratio (LCR) approach, which refers to the ratio of indemnities to liabilities. This normalized LCR approach is less sensitive to trending yields. A challenge with farm-level yield trend estimation is that the yield series are highly sensitive to annual variability in growing conditions. Therefore, when relatively shorter yield histories are used, there can be concerns over the accuracy of the trend adjustment. In the U.S., the trend adjusted yield factor is offered as an endorsement for the actual production history program, which is free to elect and is available to producers in counties where average county yield is trending. Trends are estimated for the county, and producers can elect for their yield history to be adjusted relative to the estimated county trend. And yields are conditioned for weather effects to account for favorable or unfavorable weather effects in the county. A possible negative side effect of the yield trend endorsement may be the fact that yield trends are estimated at the county level, rather than the farm level. This is because there may be a risk of adverse selection, because producers who are early adopters of technology may subsidize the slow technology adopting producers. However, the benefit of adopting a yield trend endorsement may far outweigh the added risk. In Canada, the actuarial ratemaking methodology is not publicly reported; however, discussions with the Agriculture and Agri-Food Canada revealed that each provincial government crop insurance agency is able to apply a trend factor to it PY method. However, a cap (maximum trend) is applied to each crop type in each region, which is approximately 2.5%.

The objective of this study is to propose a methodology to isolate crop yield improvement due to technological change, as well as crop yield improvement due to favorable and unfavorable weather. Two decomposition frameworks are presented here, including a fixed-effect approach and a relative effect approach. The findings of this study may be particularly useful for crop insurance and may be used to improve crop insurance by providing a procedure to adjust producers' PY for positive trends. The proposed methodology may help to isolate the influence of technological change from weather effects on crop yield over the sample period and may be useful for crop insurance in Canada, the U.S. and other countries.

Section 2: Motivation: Why Does the Effect of Technology on Crop Yield Matter?

To illustrate the empirical exercise, the authors present some motivation for the importance of technology and weather decomposition on crop yield using a simplified framework. Consider an agricultural loss random variable in year *t*, *Xt*, which has the following decomposition:

$$X_t = W_t - \eta_t + \varepsilon_t, \tag{1}$$

where W_t is the impact from weather variables, which has a distribution with mean $\overline{W}$ and variance σ_w^2 ; η_t is the technological impact/shock in year t, with mean $\overline{\eta_t}$ and variance σ^2; and ε_t is the idiosyncratic risk in year t, with mean 0 and variance σ^2, that cannot be explained by either the weather effect or technology shock. The impact from a technology shock usually includes an adoption of some new technology, improvement in farming practice, introduction of new seed varieties, etc. As a result, it is reasonable to assume that the effect of technology will always reduce agricultural losses, and hence there is a negative sign in the decomposition in Equation (1). The authors also assume that W_t, η_t, and ε_t are orthogonal with each other.

Now consider the wealth of an individual farmer, who is faced with agricultural loss. In the uninsured condition, the terminal wealth is $\omega_t = \omega_0 - X_t$, where ω_0 is the initial wealth level. Now suppose the farmer purchases θ units of insurance with price π. Then the insured wealth becomes $\omega_t = \omega_0 - X_t + \theta(X_t - \pi) = \omega_0 - (1-\theta)X_t - \theta\pi$.

The information asymmetry enters when the insurer is determining the contract price π. Usually the farmer has better information than the insurer. In the case of technological adoption, the farmer knows the level of the technological impact η_t, while the insurer usually does not have detailed information with respect to this. Now assume that the insurer prices the insurance contract following the most conservative assumption, that is, all farmers in the insurer's risk pool do not adopt any new technology. Hence, in such a conservative assumption, the insurer has a risk decomposition model $X_t = W_t + \varepsilon_t$, and the insurance policy has price π = E (X_t) = $\overline{W}$.

With this insurance contract available, the farmer needs to make the decision about how much insurance to purchase, i.e., to determine θ. Suppose the utility function of the representative farmer at wealth level ω is U (ω), then the farmer needs to maximize the following expected utility function:

$$\arg\max_{0\le\theta\le1} E[U(\omega_t)] = \arg\max_{0\le\theta\le1} E[U(\omega_0 - (1-\theta)(W_t - \eta_t + \varepsilon_t) - \theta\overline{W})]. \tag{2}$$

With a few steps of simple derivation, the calculation becomes:

$$U'(\overline{\omega})\overline{\eta} = -(1-\theta^*)\cdot(\sigma_w^2 + \sigma_\eta^2 + \sigma_\varepsilon^2)\cdot U''(\overline{\omega}). \tag{3}$$

For a strict concave utility function, that is, $U''(\omega) < 0$ for all ω, the coefficient of risk tolerance $\tau(\omega) = -\frac{U'(\omega)}{U''(\omega)}$ under the optimal condition can be expressed as

$$\tau(\omega) = \frac{(1-\theta^*)(\sigma_w^2 + \sigma_\eta^2 + \sigma_\varepsilon^2)}{\overline{\eta}}. \tag{4}$$

There are a few interesting implications from Equation (4):

- θ^*decreases as $\tau(\omega)$ increases. More specifically, when information asymmetry exists, the more risk tolerant the farmer, the less insurance the farmer will purchase.

- For each fixed level of risk tolerance, θ^*decreases with $\bar{\eta}$. This implies that farmers who adopt more advanced technology tend to purchase less insurance contracts under the information asymmetry. Due to this effect, in the long run, quick technology adopters will eventually leave the contract, and the insurer's risk pool will become unbalanced.

- For each fixed level of risk tolerance, θ^*increases with Var(X_t) = $\sigma_w^2 + \sigma_\eta^2 + \sigma_\varepsilon^2$. This is intuitive, because the number of insurance contracts purchased increases with the total underlying risk.

- Finally, since the $U'(\omega) < 0$ and $U''(\omega) < 0, \theta^* = 1$ if and only if $\bar{\eta}$ = 0. This result indicates that in the existence of information asymmetry, only those farmers who do not adopt any new technology will choose full insurance coverage. This result is important in the sense that the insurer can make use of the purchasing behavior of farmers to differentiate them with respect to their technology adoption.

The analysis of the above decomposition model illustrates that it is important to isolate crop yield improvement due to technical change from crop yield improvement due to weather.

Section 3: Proposed Methodology

Empirical models use underlying variables to estimate crop yield, and—if specified adequately—they can be extended to model the underlying technological trend. Crop yields likely depend on several factors, such as precipitation, temperature, solar radiation, soil effects, management, and technology. The influence of technology on crop yield may be geographically dependent, because some areas are early adopters of new technology and other areas adopt more slowly. Technology can be defined as any form of potential yield improvement and can be influenced by factors such as improved seeds, crop genetics, management practices, and precision agriculture improvements. This study attempts to separate the effects of weather and technology from the overall yield gain. This section discusses the proposed methodology. First, two decomposition frameworks are presented to separate technological gain from weather gain. This includes a fixed-effect approach and a relative effect approach. Second, the authors describe the data and processing.

3.1 DECOMPOSITION FRAMEWORK I: FIXED-EFFECT APPROACH

The first decomposition framework controls for heterogeneity between different geographical levels (such as farms, EcoDistricts, EcoRegions or counties) by estimating the fixed effects or separate intercept terms for each farm. Motivated by the method in Lobell et al. (2011), the weather effects are separated from the technology-related yield gain, by first controlling for the heterogeneity between groups and then modeling the weather and technology components. The model is specified below.

$$y = X\beta + D\alpha + \varepsilon, \qquad (5)$$

where *X* is a matrix of weather proxy variables; X_0 is the intercept, X_1 is the linear technology time trend with a polynomial time trend; and *D* is a dummy variable matrix composed of indicator variables identifying the geographical locations. This model is estimated by least squares methods. The least squares dummy variable model (LSDV) is more specifically described as follows,

$$y_{it} = \beta_1 X_{1it} + \beta_2 X_{2it} + \cdots + \beta_k X_{kit} + \emptyset(t) + \mathrm{D_i}\alpha_i + u_{\mathrm{it}} \qquad (6)$$

y_{it} represents the crop yield at farm/EcoDistrict/county *i* in year *t*,

- X_{it} represents weather proxy variables.
- ϕ(t), in the form ϕ(1) + ϕ(t) + ϕ(t^2) + ... + ϕ(t^p) where p is the degree of the polynomial, and may be specified as the polynomial technology time trend. In this study, the trend variable is specified as p = 1 i.e., a linear technology trend.

$\alpha_i + u_{\mathrm{it}}$ where α_i is the farm, EcoDistrict, EcoRegion or county specific unobserved heterogeneity, and u_{it} is the idiosyncratic error.

At the aggregate level, each of the weather proxy variables are estimated in a LSDV model. This model controls for heterogeneity between different geographic levels, indicated by the dummy variables, and adds a technology time trend to estimate the technological trend over the time period. Also, robust MM estimation is calculated. Finger (2013) suggested this method as an alternative to ordinary least squares (OLS) regression for yield trend estimation, because it may help to reduce the variance of the yield trend estimates. In this study, the authors estimate at the aggregate level a total of six models. The weather variables are expected to capture the effect of weather, and the technology trend captures the technological yield gains, such as improved genetics or improvements in management practices. The technology time trend is specified as a linear first-degree polynomial.

In addition to the aggregate level estimation, the models are fit at different levels of aggregation and at the farm level. For example, at the farm-level, separate regressions are estimated for each farm based on its yield history. Similarly, for the other levels of aggregation, including EcoDistrict, and EcoRegion (which the authors discuss in more detail in Section 4), the trends are separately regressed. Similarly, MM estimation is tested for each series at the various aggregation levels. These results are available in the Appendix and were conducted as robustness checks.

3.2 DECOMPOSITION FRAMEWORK II: RELATIVE EFFECT APPROACH

The second decomposition framework focuses on decomposing the weather and technological effect into two orthogonal parts as follows:

$$y_{i,t} = f(W_{i,t}) + \eta_{i,t} + \varepsilon_i, \tag{7}$$

where $y_{i,t}$ denotes the yield production in each year *t* at each geographic location *i*; $W_{i,t}$ is a vector of weather variables; ηi,t is a random variable that represents the technology improvement at each location *i*; and ε_i is an i.i.d.random variable representing the idiosyncratic residual. In this decomposition framework, there are two important assumptions: (1) the three components, f($W_{i,t}$), $\eta_{i,t}$, $\in_i$, which are orthogonal with each other; and (2) $\in_i$ is an i.i.d. random variable that does not depend on time.

Given the decomposition model in equation (7), estimating the corresponding technological effect $\eta_{i,t}$ for each location is still a challenging task. This is because it is difficult to find a good proxy of the technology effect for crop yield with available data. This is the reason why—in the first decomposition framework in this research, as well as in other current literature—polynomial functions of time are used as a technology effect proxy. Due to this difficulty, the framework proposed in this section considers a relative effect decomposition. For a certain estimated weather model $\hat{f}(W_{i,t})$, it is easy to estimate the residuals $\hat{\varepsilon}_{i,t} = y_{i,t} - \hat{f}(W_{i,t}) = \eta_{i,t} + \in_i$, but it is challenging to further identify $\eta_{i,t}$ and $\in_i$ individually. Note that $\eta_{i,t}$ = $\hat{\varepsilon}_{i,t}$ − $\in_i$, and taking the average of this equation over the sample period t = 1, . . . , T gives

$$\frac{1}{T}\sum_{t=1}^{T}\eta_{i,t} \equiv \overline{\eta_i} = \frac{1}{T}\sum_{t=1}^{T}\hat{\varepsilon}_{i,t} - \in_i \equiv \overline{\varepsilon_i} - \in_i. \tag{8}$$

Therefore, removing equation (8) from $\eta_{i,t} = \hat{\varepsilon}_{i,t} - \in_i$ gives

$$\eta_{i,t} - \overline{\eta_i} = \hat{\varepsilon}_{i,t} - \overline{\eta_i}, t = 1,2,\dots,T. \tag{9}$$

It is nice to have equation (9), because the idiosyncratic residuals disappears. Equation (9) leads to modeling the technology effect on a relative scale, by comparing to the historical average level of each location. Therefore, by demeaning equation (7) on both sides, the authors propose the following decomposition:

$$y_{i,t} - \overline{y_i} = [\hat{f}(W_{i,t}) - \overline{\hat{f}(W_{i,t})}] + [\eta_{i,t} - \overline{\eta_i}], \tag{10}$$

where $\overline{y_i} = \frac{1}{T}\sum_{t=1}^{T} y_{i,t}$ is the historical normal level of crop yields; $\overline{\hat{f}(W_{i,t})} = \frac{1}{T}\sum_{t=1}^{T}\hat{f}(W_{i,t})$ is the historical normal level of weather effects; and $\overline{\eta_i} = \frac{1}{T}\sum_{t=1}^{T}\eta_{i,t}$. The decomposition framework II proposed in this section includes the following three parts:

a) Total gain ($y_{i,t} - \overline{y_i}$): total production gain relative to the historical normal production level of location *i*.

b) Technology gain ($\eta_{i,t} - \overline{\eta_i}$): production gain from the technology effect relative to the historical normal technological level of location *i*.

c) Weather gain ($\hat{f}(W_{i,t}) - \overline{\hat{f}(W_{i,t})}$): production gain from weather effects relative to the historical normal weather level of location *i*.

A decomposition example following from equation (10) is displayed in Figure 2.

Figure 2

RELATIVE EFFECT APPROACH: A DECOMPOSITION EXAMPLE

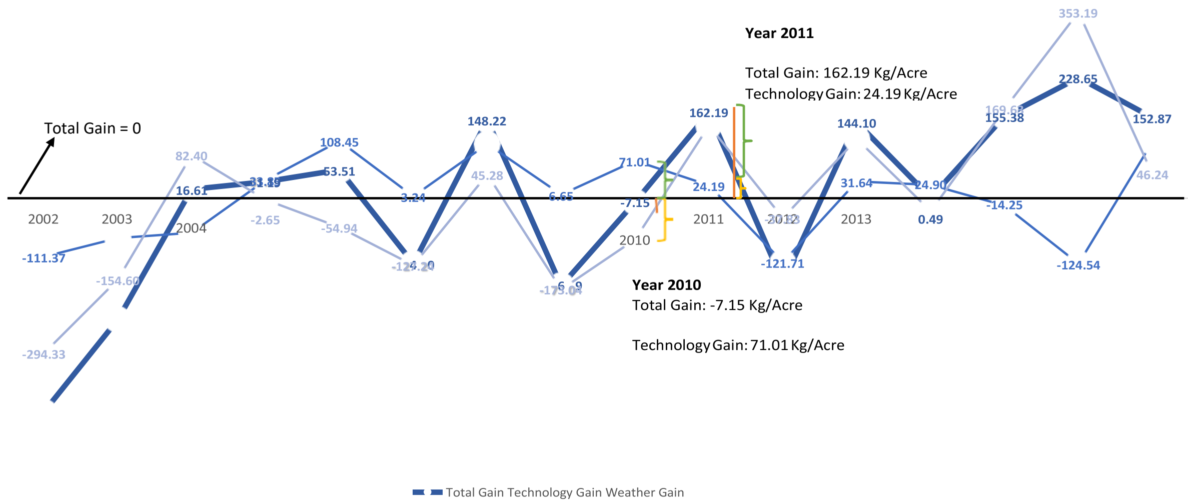

The horizontal black line is the zero total gain line, representing a production level equal to the historical average. Years above the line indicate a positive total gain compared to the historical level and vice versa. For example, for the year 2011, the total gain is positive and equal to 162.19 kg/acre, which means 2011 is performing better than the historical average, among which 24.19 kg/acre is due to the technology effect (technology gain), and 138.00 kg/acre is due to the weather effect (weather gain). In contrast, for the year 2010, the total gain is negative and equal to −7.15 kg/acre, which means 2010 is performing worse than the historical average level. Within the 7.15 kg/acre negative gain, 78.15 kg/acre is due to the negative weather gain (weather conditions worse than historical normal level), and 71.01 kg/acre is due to the positive technology gain (technology improvement better than historical average).

Figure 3

PUBLIC AGRICULTURAL R&D FUNDING INPUTS (IN BILLION USD) WITH TECHNOLOGY GAIN

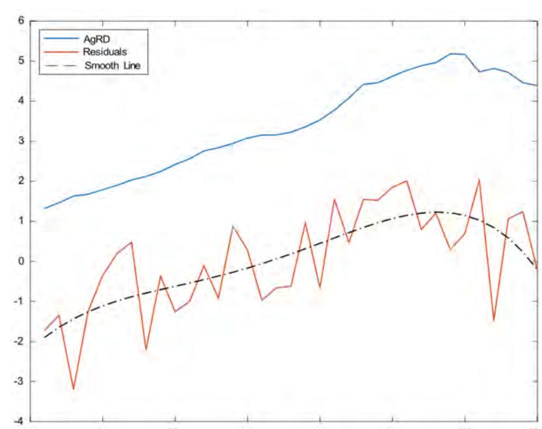

Dashed black lines are smoothing line with simple polynomials.

To help support the above decomposition framework, the authors compare the resulting technology gain to some technology improvement proxy. They consider the annual total public agricultural research and development funding inputs (AgRD) as a good proxy for the technological trend in the U.S., which they obtained from the Economic Research Service from the U.S. Department of Agriculture (USDA). Plotting the two series together, displayed in Figure 3, the technology gain shares much of the same pattern as the AgRD, and this provides some support for the authors' proposed decomposition framework. The next challenge to achieve an effective decomposition framework is to accurately and sufficiently model the weather component. In the empirical analysis of this research, the authors propose several machine-learning-based (likely nonlinear) predictive models to estimate $f\left(W_{i,t}\right)$.

Section 4: Data and Processing

In this research, the authors created two datasets for the empirical analysis. The first one is a county-level corn yield dataset from Iowa in the U.S., spanning 37 growing seasons from 1982 to 2018 for 99 counties, which gives a total of 3,663 yield observations. Weather proxy indices are then extracted for the corresponding counties and linked to yields. The second dataset is field-level canola yield data from the province of Alberta, Canada, spanning 16 years from 2002 to 2017. This amounts to 591,430 field-level yield observations corresponding to geo-located quarter sections for a total of 1,293 representative farms that have 16 years of canola history. Weather variables are extracted at the corresponding quarter section and then weighted to the farm-level. The dataset is an unbalanced panel; and due to farm crop rotation and changes in planting decisions, many producers do not produce canola each year.

For quality control of the extracted weather variables, the authors only considered for analysis field units greater than 150 acres. As a result, the number of producers who planted at least one canola field over 150 acres each year over the 16 growing seasons is reduced from 1,293 producers to 78 producers. The amount of producers with 10 years of canola production history with a field over 150 acres is 389, and at least five years of canola production history with a field over 150 acres is 883. A sensitivity analysis is conducted on the full 1,293 representative farms to compare the results of the quality controlled sample of 78 farms with the full dataset of 1,293 farms.

4.1 IOWA CORN DATASET

The first dataset includes county level yield data downloaded from the USDA's National Agricultural Statistics Service database for 99 counties in Iowa from 1982 to 2018 for a total of 37 years. These county yields, which are measured in bu/acre, are then joined with the weather proxy indices, weather variables, vegetation indices and biophysical indices. The weather variables, including minimum and maximum temperatures, precipitation and shortwave solar radiation are extracted from the University of Idaho Gridded Surface Meteorological Dataset weather grid, which is an interpolated weather grid that uses information from ground weather stations and ancillary information to provide estimates of weather variables over the continuous U.S. at a 4km-by-4km resolution grid. The weather variables of interest are extracted for each county for each day of the year and averaged over the county. This results in 365 weather observations for each variable for each respective county each year. These weather variables are aggregated to eight-day periods by averaging and are then compressed into a lower-dimensional dataset by using principal component analysis (PCA). This transforms the weather variables into orthogonal components that describe most of the information contained in the weather variables but with less dimensionality. These principal components (PCs) are saved for further analysis. The first 10 principal components described approximately 70% to 80% of the variation in weather observed in the multiple weather variables. Table 1 displays some descriptive statistics of the county corn yields.

Table 1
SUMMARY STATISTICS OF IOWA CORN YIELDS

Statistic	N Mean	St. Dev.	Min	Pctl (25)	Pctl (75)	Max
Yield*	3,645[a] 1 44.669	35.420	19.100	123.500	171.000	226.000

*Yield is the county level corn yield; [a] 3845 observation from 99 counties × 37 years

4.2 ALBERTA CANOLA DATASET

The next dataset is the field-level canola yield dataset from Alberta, Canada, spanning 16 years from 2002 to 2017. This amounts to 591,430 field-level yield observations corresponding to geolocated quarter sections for a total of 1,293 representative farms. The dataset is an unbalanced panel; and due to farm crop rotation and changes in planting decisions, many producers do not produce canola each year. For quality control of the extracted weather variables, the authors only considered for analysis field units greater than 150 acres. As a result, the number of producers who planted at least one canola field over 150 acres each year over the 16 growing seasons is reduced from 1,293 producers to 78 producers. The amount of producers with at least 10 years of canola production history with a field over 150 acres is 389 and those with at least five years of canola production history with a field over 150 acres is 883. The yield values are averaged at the farm-level and are not measured at the field unit. For example, a farmer in 2015 may have three canola fields, but there is only one yield value, which is downscaled to the field unit (i.e., the yields have been averaged over the three parcels and are not stored at the field unit). However, the weather and remote sensing indices are measured and extracted at the field level and then averaged based on the field unit acreage size proportionate to the total number of acres of canola grown on the farm for the year. As a result, the measurements for the weather variables are directly linked to the actual quarter section locations, which provides a higher degree of accuracy. Table 2 displays some descriptive statistics for canola yields.

Table 2

SUMMARY STATISTICS OF ALBERTA CANOLA YIELDS AND FARM SIZE

Statistic	N	Mean	St. Dev.	Min	Pctl (25)	Pctl (75)	Max
Total Acres*	1,248[a]	1,497.211	1,989.750	152	320	1,810.2	19,870
Yield**	1,248	720.836	317.835	0.000	483.765	942.557	1,624.059

Table 2 displays canola yield summary statistics of the 78 farms in Alberta, Canada. *Total Acres represents the total amount of canola acres insured. **Yield is canola yield averaged over the field units to the farm-level. [a] 1248 observations of 78 farms × 16 years.

Table 3

WEATHER VARIABLE DESCRIPTIONS

Variable	Units	Description
Min-temperature	C°	Daily minimum 2-meter air temperature.
Max-temperature	C°	Daily maximum 2-meter air temperature.
Shortwave radiation	W/m^2	Incident shortwave radiation flux density, taken as an average over the daylight period of the day.
Precipitation	mm	Daily total precipitation, sum of all forms converted to water-equivalent.

The authors collected weather variables from NASA's Daymet weather grid, which provides global daily weather variable measurement, derived from weather station and ancillary data, at a 1km-by-1km resolution grid. The weather variables, including minimum and maximum temperatures, precipitation and shortwave solar radiation are extracted for each field unit for each day of the year. Each of the 591,430 unique field units has 365 observations of each weather variable. These observations are then weighted based on the proportion of the farm's total acreage to the farm-level. The weather variables are then compressed using the same PCA procedure employed for the Iowa data. The first 10 principal components are used for analysis. Descriptions of the weather variables are available in Table 3.

The normalized difference vegetation index (NDVI) is derived from NASA's Moderate Resolution Spectral Radiometer (MODIS) using the MOD09A1 MODIS/Terra Surface Reflectance 8-Day L3 Global 500m product (Vermote, 2015). This product provides eight-day best pixel composite surface reflectance values that the authors then use to calculate NDVI. NDVI values are measured at a resolution of 500 meters and then extracted at the field unit level. These values are then smoothed and integrated to get the INDVI at each field unit location and then weighted proportionately to the farm level.

The leaf area index (LAI) and fraction of absorbed photosynthetic radiation (FAPAR) biophysical parameter indices are collected and aggregated in the same way as the NDVI vegetation index, however, they include additional spectral band information in their calculation. These indices are calculated and then weighted proportionately to the farm-level. Descriptions of the remote sensing variables are available in Table 4.

Table 4
REMOTE SENSING VARIABLE DESCRIPTIONS

Variable	Min	Max	Description
NDVI	−1	1	Normalized difference vegetation index
LAI	0	100	The one-sided green leaf area per unit of ground area
FAPAR	0	100	The FAPAR absorbed by the vegetation canopy

4.3 LEVELS OF AGGREGATION

Estimating farm level trends is sensitive to the sampled weather over the time period at those field units and may need to be aggregated further to consistently estimate the yield trend. As a result, the field unit canola yields from Alberta are aggregated to three different levels, from highest aggregation to the lowest: (1) EcoRegion, (2) EcoDistrict and (3) farm, respectively. The farm-level aggregation stage was discussed above, and the weather variables measured at the field units were aggregated to this farm level. The farm locations are joined with the Canadian Ecological Framework consisting of a nested hierarchy classifying Canada into EcoZones, EcoRegions and EcoDistricts characterized by distinctive landform, relief, surficial geologic material, soil, water bodies, vegetation and land uses.

EcoDistricts are the most detailed in the ecological hierarchy and are used to group the farms. In Alberta, there are more than 1,000 distinct EcoDistricts, and approximately 100 of these regions have canola grown in them. The EcoDistricts that have canola acres continuously over the period are kept for analysis for a total of 59 EcoDistricts. The mean canola yield across EcoDistricts are displayed in Figure 4. The figure shows that the average canola yield in these distinct EcoDistricts is likely heterogenous, which has been documented in the literature (See, for example, Barnett et al. (2005); Woodard et al. (2012); Zhu et al. (2015)). Producers management styles may vary between EcoDistricts to manage the distinct growing conditions. The heterogeneity of growing conditions across EcoDistricts is likely the cause of much of the yield differences across EcoDistricts.

Figure 4
HETEROGENEITY OF MEAN CANOLA YIELD ACROSS ECODISTRICTS

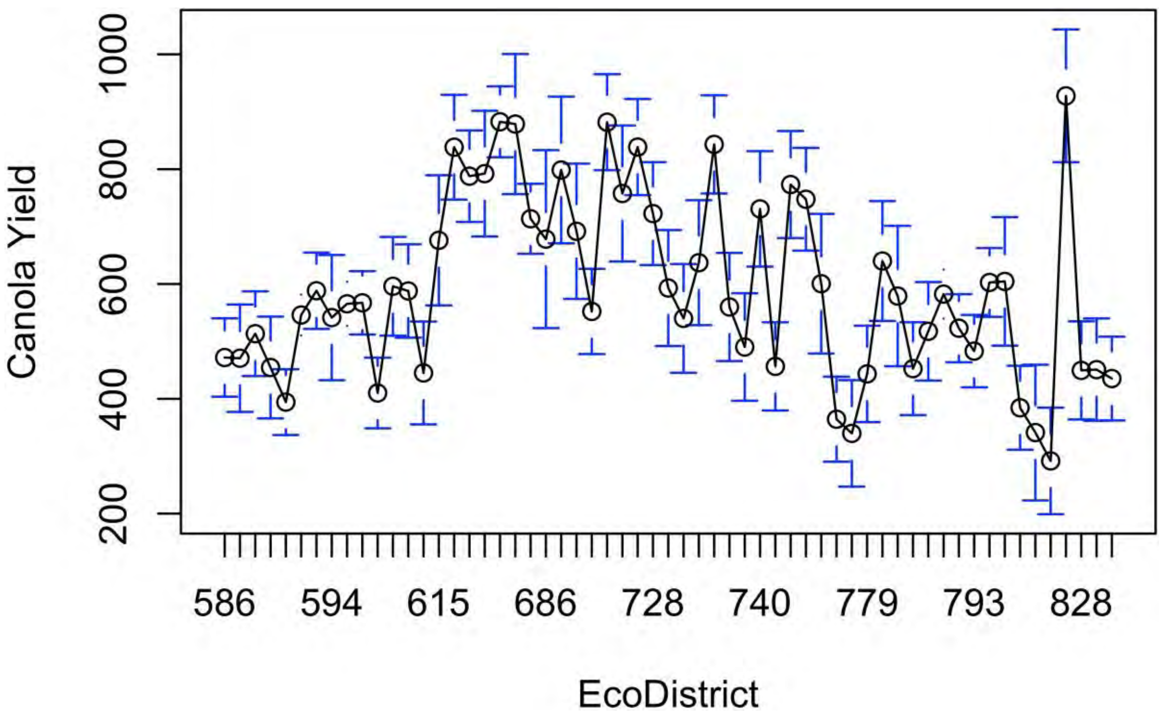

The x-axis displays the EcoDistrict ID number; the y-axis displays the EcoDistrict average canola yield.

EcoRegions are the next step in the hierarchy of land classification and are larger than EcoDistricts and generally are defined by a distinctive feature of the area. These regions aggregate the farms in a smaller set of regions than the EcoDistricts and may still represent the differences in production and yield described above. When grouped by EcoRegions, the farms are grouped into five subgroups that can be further analyzed. The number of farms grouped by EcoDistrict and EcoRegion are displayed in Figure 5 in Section 5.

Section 5: Empirical Analysis Results

In this section, the authors empirically evaluate the two proposed decomposition frameworks using the two datasets introduced in Section 4. Subsections 5.1 and 5.2 introduce the trend estimation results based on framework I and II, respectively.

5.1 EMPIRICAL RESULTS OF DECOMPOSITION FRAMEWORK I

Farm-level yield trend estimation can be sensitive to high and low yields and data sparsity. As a result, the estimated farm-level yield trends can be highly varied between farms, and this variability in the estimated trend coefficients can be assessed by the standard deviation of the estimated trends and the range of trend values. For crop insurance, trend values estimated at the farm level may be too variable to warrant a trend adjustment, and trend adjustment may be better conducted at an aggregated level.

Figure 5
ECODISTRICT AND ECOREGION AGGREGATION LEVEL OF THE ALBERTA CANOLA DATA

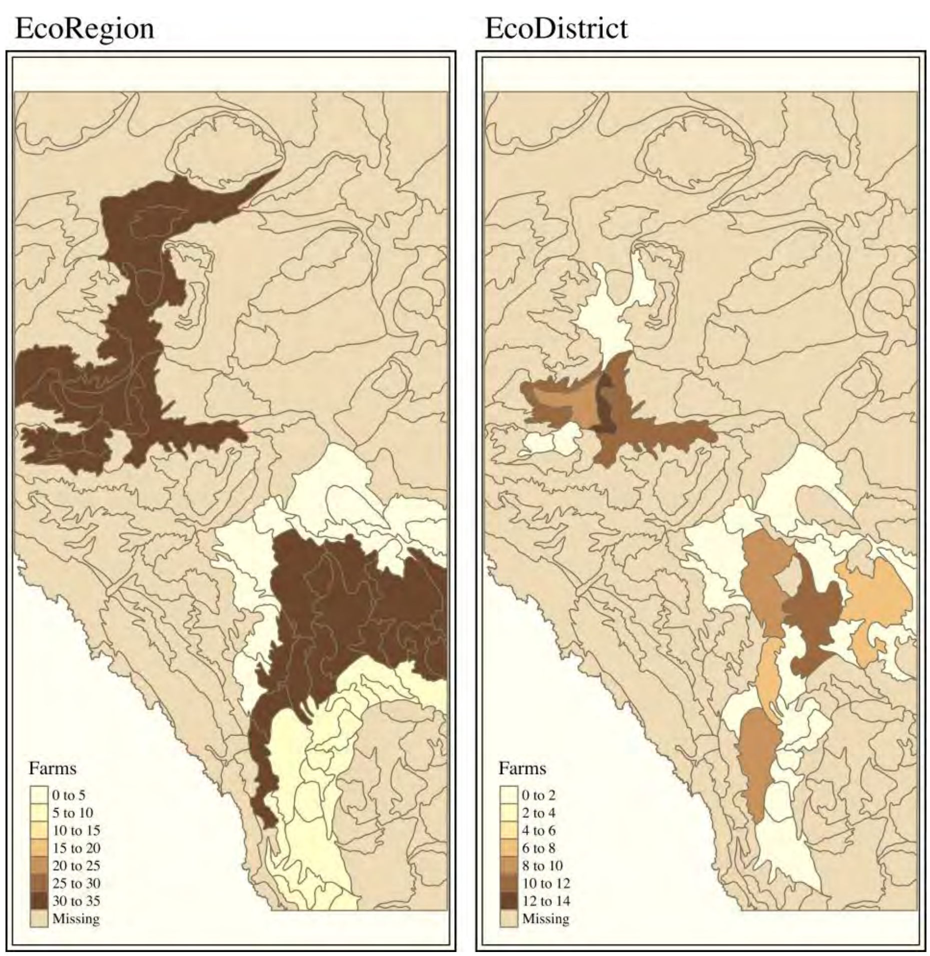

5.1.1 IOWA CORN DATASET

For the Iowa county-level corn yields dataset, the trend is first estimated at the state level using a LSDV regression. Next the state-level trend is estimated using the MM robust estimator using dummy variables and a linear time trend. More specifically, the following six LSDV regression models are estimated:

a) Model A: baseline model with a linear time trend estimated by OLS

b) Model B: model with a linear time trend estimated by MM robust estimation

c) Model C: integrated NDVI and a linear time trend

d) Model D: integrated LAI and a linear time trend

e) Model E: integrated FAPAR and a linear time trend

f) Model F: weather-variable principal components (minimum and maximum temperatures, precipitation, short-wave radiation) with a linear time trend

Once the models are estimated, the technology trend can be interpreted. The linear coefficient fitted to the technology trend variable can be interpreted to distinguish whether the effect is negative or positive, and the magnitude of technological change can be evaluated. The technology time trend captures the rate of technological change over the time period for corn yields.

The estimated yield trend for the six fixed effects regression models for the Iowa data set are displayed in Table 5. The baseline trend is used as a reference to approximate the weather effect over the period. The weather effect is defined as the baseline max technology minus the compared model technology.

In Table 5, Model A reports the baseline linear trend model estimated by OLS, and has an estimated technology trend of 2.420 bu/acre. That translates to a yield increase over the 37 years of $2.420 * 37 = 89.54$ bu/acre. Model B is a robust regression using MM estimation with a linear time trend, and the estimated yield trend is 2.382 bu/acre, which is slightly less than the baseline trend. This implies a slightly positive weather effect of $2.420 - 2.382 = 0.038$ bu/acre over the period. Model C, the linear time trend and NDVI model has a technology trend of 2.375 bu/acre, which is similar to the baseline trend and the robust regression model. The weather effect is 0.045 bu/acre over the period, translating to $0.045 * 37 = 1.665$ bu/acre gained due to favorable weather conditions over the time period.

Table 5

AGGREGATE REGRESSION RESULTS FOR THE IOWA CORN DATASET

Model	Description	Tech Trend	Weather Trend	%Tech Trend	% Weather Trend
A	$y \sim f(time)$	2.420	–	–	–
B	$y \sim f(time)$ MM	2.382	0.038	98.43%	1.57%
C	$y \sim f(time, NDVI)$	2.375	0.045	98.14%	1.86%
D	$y \sim f(time, LAI)$	0.934	1.486	38.60%	61.40%
E	$y \sim f(time, FAPAR)$	1.476	0.944	60.99%	39.01%
F	$y \sim f(time, PCs)$	1.950	0.47	80.58%	19.42%

Note: Yields and trends are measured in bu/acre.

Similarly, for the LAI model (Model D) and the FAPAR model (Model E), the technology trends are 0.934 and 1.476, respectively. Last, the weather principal component model (Model F) measures a technology trend of 1.950 bu/acre per year and a weather effect of 0.47 compared relative to Model A. Figure 6 shows the estimated trends for Model A to Model F based on using the Iowa dataset.

If the assumption is that weather has not significantly changed in the growing region over the time period, then the technology trend should converge to the long-term trend. Assuming that 37 years is a sufficiently long crop yield history to meet this long-term convergence and assuming weather has not drastically changed in the region, then the baseline model's estimated technology trend should be the effect of technology over the period. From this, the authors can establish the negative bias of technology that the model with NDVI has when estimating the technology trend due to the amount of technological change captured by NDVI over the period. This effect is measured by 2.420 − 2.375 = 0.45 bu/acre, or if using the robust regression, Model B, as the baseline true effect of technology over the period, then the amount of technological change captured by NDVI over the period is 2.382 − 2.375 = 0.007 bu/acre, which is a marginal amount over the period and likely of little economic significance.

Figure 6

ESTIMATED AGGREGATE TECHNOLOGY TRENDS FOR IOWA DATASET

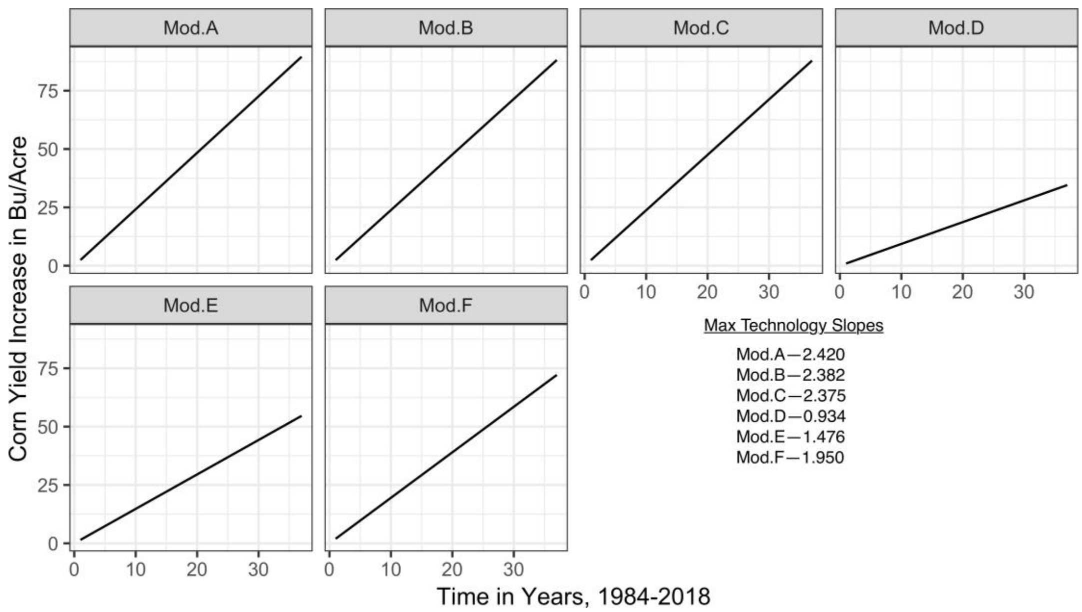

5.1.2 ALBERTA CANOLA DATASET

Next, the authors use the Alberta farm-level canola yield dataset to estimate the LSDV model and then the provincial level trend using the MM robust estimator using dummy variables and a linear time trend. More specifically, the following four LSDV regression models are estimated[1]:

a) Model A: baseline model with a linear time trend estimated by OLS

b) Model B: baseline model with a linear time trend estimated by MM robust estimation

c) Model C: integrated NDVI and a linear time trend

[1] Given the results from the Iowa Data, the LAI and FAPAR models were not estimated because they were found to capture a large amount of technological change.

d) Model D: weather-variable principal components derived from minimum and maximum temperatures, precipitation and short-wave radiation, with a linear time trend

Average canola yield over the period from 2002 to 2017 for the 591,430 canola fields displayed large yield increases and heterogeneity over the period[2]. Figure 7 shows the heterogeneity of cross-section average canola yield for each year over the time period. This yield trend is sometimes ascribed to improvements in crop technology, including enhanced seeds, improved farm management practices, etc. However, some of this yield increase may be attributed to favorable weather over the period. Since crop yields are observed only once per year, a biased pattern of subsequent highly favorable weather growing seasons may occur, and this pattern may increase average yield for a time before converging to a true mean yield.

Figure 7

HETEROGENEITY ACROSS YEARS FOR THE ALBERTA DATA

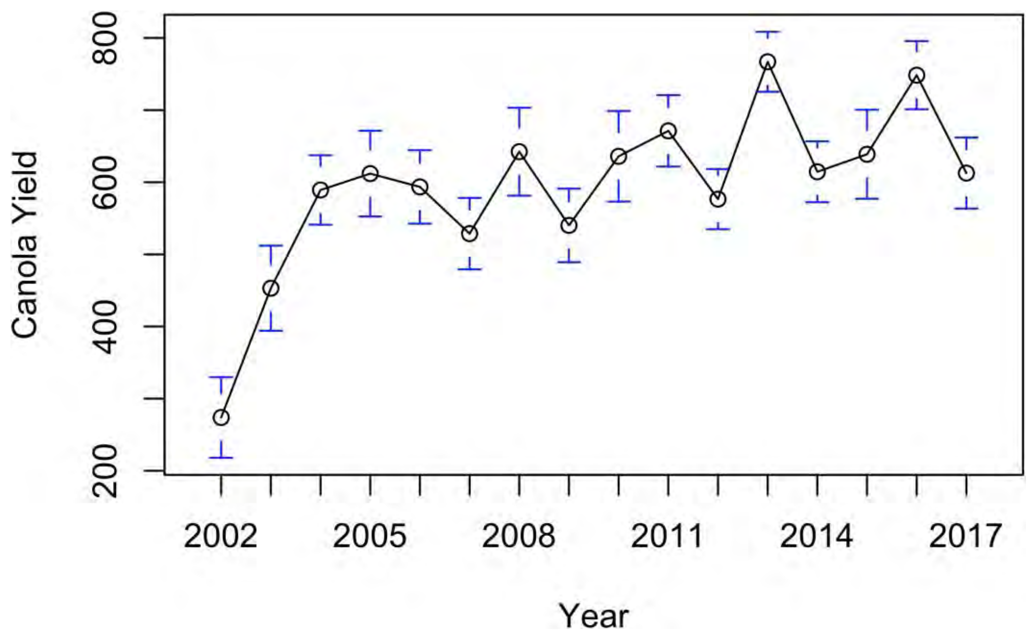

The x-axis displays the year, and the y-axis displays the provincial average canola yield. Table 6 displays the estimated technology and weather trends for the Alberta canola data. Model A has an estimated technology trend of 17.563 kg/acre, which is a large increase per year. This estimated technology trend translates to a 17.563 × 17 = 298.571 kg/acre increase over the period. Converted to bu/acre using a factor of 1/44 (approximate kg/acre to bu/acre approximation factor) results in a bu/acre increase over the period of 6.786 bu/acre.

Table 6

DATASET 2, CANOLA YIELDS, ALBERTA, CANADA

Model	Description	Tech Trend	Weather Trend	% Tech Trend	% Weather Trend
A	y~f(time)	17.563	–	–	–
B	y~f(time) MM	11.540	6.023	65.71%	34.29%
C	y~f(time, NDVI)	10.410	7.153	59.27%	40.73%

[2] Note that 2002–2003 was a drought year. This may impact the results.

D	y~f(time, PC)	6.118	11.445	34.83%	65.17%

Note: All yields and trends are measured in kilograms per acre (kg/acre).

Model B has a technology trend of 11.540 kg/acre, significantly less than the technology rate estimated by Model A. Model B is less influenced by yields, which are outside of the prevailing trend. Model A may be influenced by the beginning year 2001, which was a large crop-loss year. The technology gain over the 16 years estimated by Model B is 11.540 × 17 = 196.18 or 4.459 bu/acre. Model C has a technology trend of 10.410 kg/acre, which translates to a technology yield gain over the period of 10.410 × 17 = 176.970 or 4.022 bu/acre. Given that the amount of technology captured by NDVI in the Iowa dataset, Model C was economically not significant; and assuming that the same result holds for canola, then the NDVI model shows a yield increase due to technology of 4.022 bu/acre versus the baseline of 6.322 bu/acre for the period. This indicates that favorable weather may have improved canola yields by 6.786 − 4.022 = 2.764 bu/acre over the 16 years. Model D has a technology trend of 6.118 kg/acre, which is much lower than the other model results. Figure 8 shows the estimated trends for Model A to Model D.

Figure 8

ESTIMATED AGGREGATE TECHNOLOGY TRENDS FOR ALBTERTA DATASET

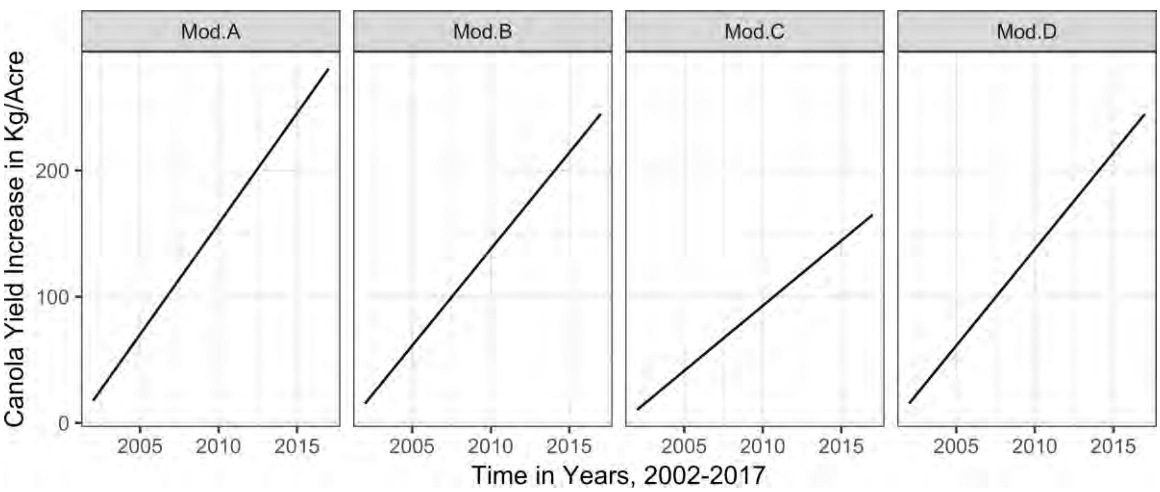

5.1.3 AGGREGATION RESULTS

For the Alberta farm-level canola yields, the effect of aggregation was tested on the crop yield trends. Technology may be adopted at different rates among farms, and by aggregating farms together for yield-trend estimation, some of this information may be lost. To test this hypothesis and to examine further the issues of aggregation, the authors use a Chow test procedure. This is motivated by Skees and Reed (1986), where a Chow test procedure is used to test different levels of aggregation (pooling of farms) to determine whether the yield trends observed on the farms are structurally different than the trends estimated at the aggregate level. Three aggregation levels are tested relative to the farm level trends, which includes the province level (highest), the EcoRegion level (medium) and the EcoDistrict level (lowest). With higher levels of aggregation, the technology and weather trends can be estimated with less uncertainty. A summary of the results are provided in Table 7. Note: Despite being a much higher level of aggregation, the EcoRegion (medium) level results are similar to the much lower EcoDistricts (lowest) aggregation level.

Table 7

CHOW TEST RESULTS, CANOLA YIELDS, ALBERTA, CANADA

Province Level[1]	
% Structural Difference	44.87% or 33/78 farms
EcoRegion Level[2]	
% Structurally Different	37.18% or 29/78 farms
EcoDistrict Level[3]	
% Structurally Different	32.05% or 25/78 farms

[1] Farm yields are aggregated in one group. [2] Farm yields are aggregated into five groups. [3] Farm yields are aggregated into 59 groups.

Table 7 shows the Chow test results for each tested level of aggregation. There is a balance between capturing the farm-specific rates of technology adoption while having enough data to accurately decompose the technology trend estimate from the overall trend. The results suggest that compared to the other levels of aggregation, the EcoRegion level may provide a good balance. For example, comparing the amount of farms that are structurally different from the aggregate trend of the EcoRegion level and the EcoDistrict level, the EcoRegion level has 29/78 farms, and the EcoDistrict level has 25/78 farms. The increase of four farms relative to the large increase in aggregation, 59 groups to five groups, may be an acceptable trade-off. In Appendix A, a full description of aggregation results is shown, including trend estimates at each level of aggregation for each respective model.

5.2 EMPIRICAL ANALYSIS RESULTS OF DECOMPOSITION FRAMEWORK II

This subsection summarizes the results of the second decomposition framework. Recall that the objective of this decomposition framework in equation (10) is to estimate the relative technology gain $(\eta_{i,t} - \overline{\eta_i})$ through reducing the weather gain $(\hat{f}(W_{i,t}) - \overline{\hat{f}(W_{i,t})})$ from the total gain $(y_{i,t} - \overline{y_i})$. As a result, the key task becomes how to model the weather gain as accurately as possible. Given the availability of weather proxy indices and other weather variables, in this research, the authors propose using some machine-learning-based (linear or nonlinear) models to estimate and predict the weather effect and, hence, estimate the technology effect. Those models are Principal Component Regression (PCR), Partial Least Square (PLS), Ridge Regression, Lasso Regression and Neural Network (NN). The remainder of this subsection briefly describes the five models the authors utilized in this research. For a more detailed introduction of the models, refer to, for example, Friedman et al. (2001) and James et al. (2013).

PRINCIPAL COMPONENT REGRESSION MODEL

The PCR approach involves constructing the first *m* principal components, where *m* can be much smaller than the original number of regressors and then using these components as the predictors in a linear regression model. More specifically, PCR uses Z = $(Z_1, \dots, Z_m)$, a set of linear combination of the original p-dimensional feature space, to construct the linear regression model:

$$Z_k = \sum_{j=1}^{p} \emptyset_{jk} X_j \qquad (11)$$

Then the linear regression model is fitted with *n* realizations of *Z* using OLS:

$$y_i = \theta_0 + \sum_{k=1}^{m} \theta_k z_{ik} + \epsilon_i, \qquad i = 1, \dots, n. \qquad (12)$$

The key advantages of the PCR are twofold:

1. Dimension reduction—often a small number of principal components are sufficient to explain most of the variability in the data, as well as the relationship with the response.

2. New feature space is orthogonal. Due to the way the principal components are constructed, they are uncorrelated with each other, which is very helpful to address the problem of collinearity.

PARTIAL LEAST SQUARE MODEL

PCR identifies linear combinations that best represent the original predictors $X_1, \dots, X_p$ in an *unsupervised* way, where the response variable is not used to help determine the principal components while performing the regression. Comparatively, PLS identifies new feature variables $Z_1, \dots, Z_m$ through a *supervised* way, which makes use of the response variable *Y* to identify new features that not only approximate the old features well but also the most related to the response variable.

More specifically, PLS obtains the first PLS component for predicting the response variable *Y* following two steps:

1. Compute $\emptyset_{j1}$ by fitting a simple linear model Y = $\beta_{j0} + \emptyset_{j1} X_j + \epsilon_j$, j = 1, . . . , p.

2. Construct the derived first PLS component $Z_1 = \sum_{j=1}^{p} \emptyset_{j1} X_j$.

Subsequent PLS components are found by taking residuals and then repeating the above steps. Since $\emptyset_{j1}$ are proportional to the correlation between *Y* and X_j, PLS places the highest weight on the variables that are most strongly related to the response variables. As a result, for PCR, there is no guarantee that the directions that best explain the predictors will also be the best directions to use in predicting the response variables. This potential drawback may be addressed by PLS through using the response variable to supervise the identification of the linear combinations.

RIDGE REGRESSION MODEL

The Ridge Regression model is a shrinkage method that constrains the coefficient estimates toward zero. The idea of shrinking the coefficients can significantly reduce the model variance and, hence, improve the model's predictive ability. The Ridge Regression estimates coefficients by solving the following problem:

$$\beta^R = \underset{\beta}{\operatorname{argmin}} RSS = \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_i x_{ij}\right)^2, \tag{13}$$

Subject to: $\sum_{j=1}^{p} \beta_j^2 \leq s,$

where *s* is the tuning parameter, which serves to control the relative impact of the constraint on the regression coefficient estimates. Usually a cross-validation method is used to search for the optimal tuning parameter.

LASSO REGRESSION MODEL

Another shrinkage method is the Lasso. The Lasso estimates the coefficients by solving the problem:

$$argmin_{\beta} \left\{ \sum_{i=1}^{n} (y_i - \beta_0 - \sum_{j=1}^{p} \beta_i x_{ij})^2 \right\} \tag{14}$$

Subject to: $\sum_{j=1}^{p} \left|\beta_j\right| \leq s,$

where s, similar to the Ridge, is the tuning parameter, which serves to control the relative impact of the constraint on the regression coefficient estimates. As in Ridge Regression, selecting a good tuning parameter is important, and cross-validation method can be used.

Lasso has a penalty function that forces some of the coefficient estimates to be exactly zero when the tuning parameter is sufficiently large. Hence, the Lasso yields sparse models that involve only a subset of the variables, which gives an idea of the most important variables in explaining and predicting each factor.

NEURAL NETWORK MODEL

NN encompasses a wide range of models and learning methods. NN are popular in the computer science literature and have been investigated for application in many areas. In particular, NNs have played an important role in pattern recognition and time series modeling (Ripley, 1993; Cheng and Titterington, 1994). In time-series forecasting with additional explanatory variables (regressors or features), future outcomes are predicted with some function of past observations, known as nonlinear (auto)regressive with external input. One key advantage of the NN method is that the predicting function does not need to be linear, which can be very important when compared to current most commonly used predictive methods in the literature. The architecture of a typical NN is displayed in Figure 9. This is a NN with one hidden layer of two neurons. The output layer (the forecast) is predicted through past observations in the input layer of $y_1, \ldots, y_T$.

Figure 9
AN ILLUSTRATION OF THE ARCHITECTURE OF A TYPICAL NN WITH ONE HIDDEN LAYER TWO NEURONS

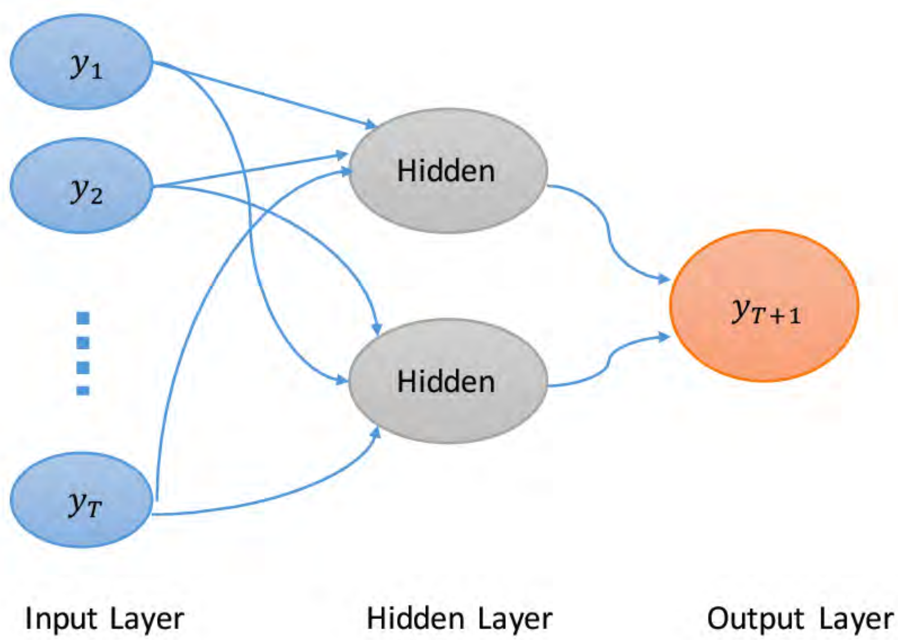

5.2.1 IOWA DATASET

The estimated technology gains and weather gains (in both absolute magnitudes and percentage) for the five machine-learning-based models for the Iowa dataset are displayed in Table 8. The percentage technology (weather) gains are calculated relative to the baseline model with a linear time trend estimated by OLS, similar to the decomposition framework in Subsection 5.1.

Table 8

ESTIMATED AGGREGATE TECHNOLOGY TRENDS FOR IOWA DATASET

Model	Description	Tech Gains	Weather Gains	% Tech Gains	% Weather Gains
A	PCR	53.01	23.34	69.43%	30.57%
B	PLS	64.98	28.53	69.49%	30.51%
C	Ridge	72.80	31.31	69.93%	30.07%
D	Lasso	63.21	36.81	63.20%	36.80%
E	NN	71.91	43.96	62.06%	37.94%

Note: Yields and trends are measured in bu/acre.

The technology gains on average estimated from the five models are quite stable, around 50 to 70 bu/acre. Take the results from Model A, the PCR model, as an example. The estimated average technology gains is 53.01 bu/acre compared to the historical yield average, or equivalently 69.43% relative to the baseline linear trend model. The decomposition results for the first models are very close.

Figure 10 shows the histogram plots of the decomposition with the Iowa dataset based on five weather models for all the counties of the Iowa dataset. Model A to Model D—i.e., PCR, PLS, Ridge and Lasso—display heavy tail properties for the estimated technology gains, while the NN produces technology gains with a more symmetric distribution. For the decomposed weather gains, it is interesting to note that all of the five models show a two-peak distribution pattern (i.e., a bimodal distribution). This finding indicates that from the 37-year period, the decomposition models identify two (good and bad) weather regimes. Figure 11 shows plots of the time series of the decomposition (in percentage) with the Iowa dataset based on the weather five models.

Figure 10

HISTOGRAM PLOTS OF THE DECOMPOSITION WITH THE IOWA DATASET BASED ON FIVE WEATHER MODELS

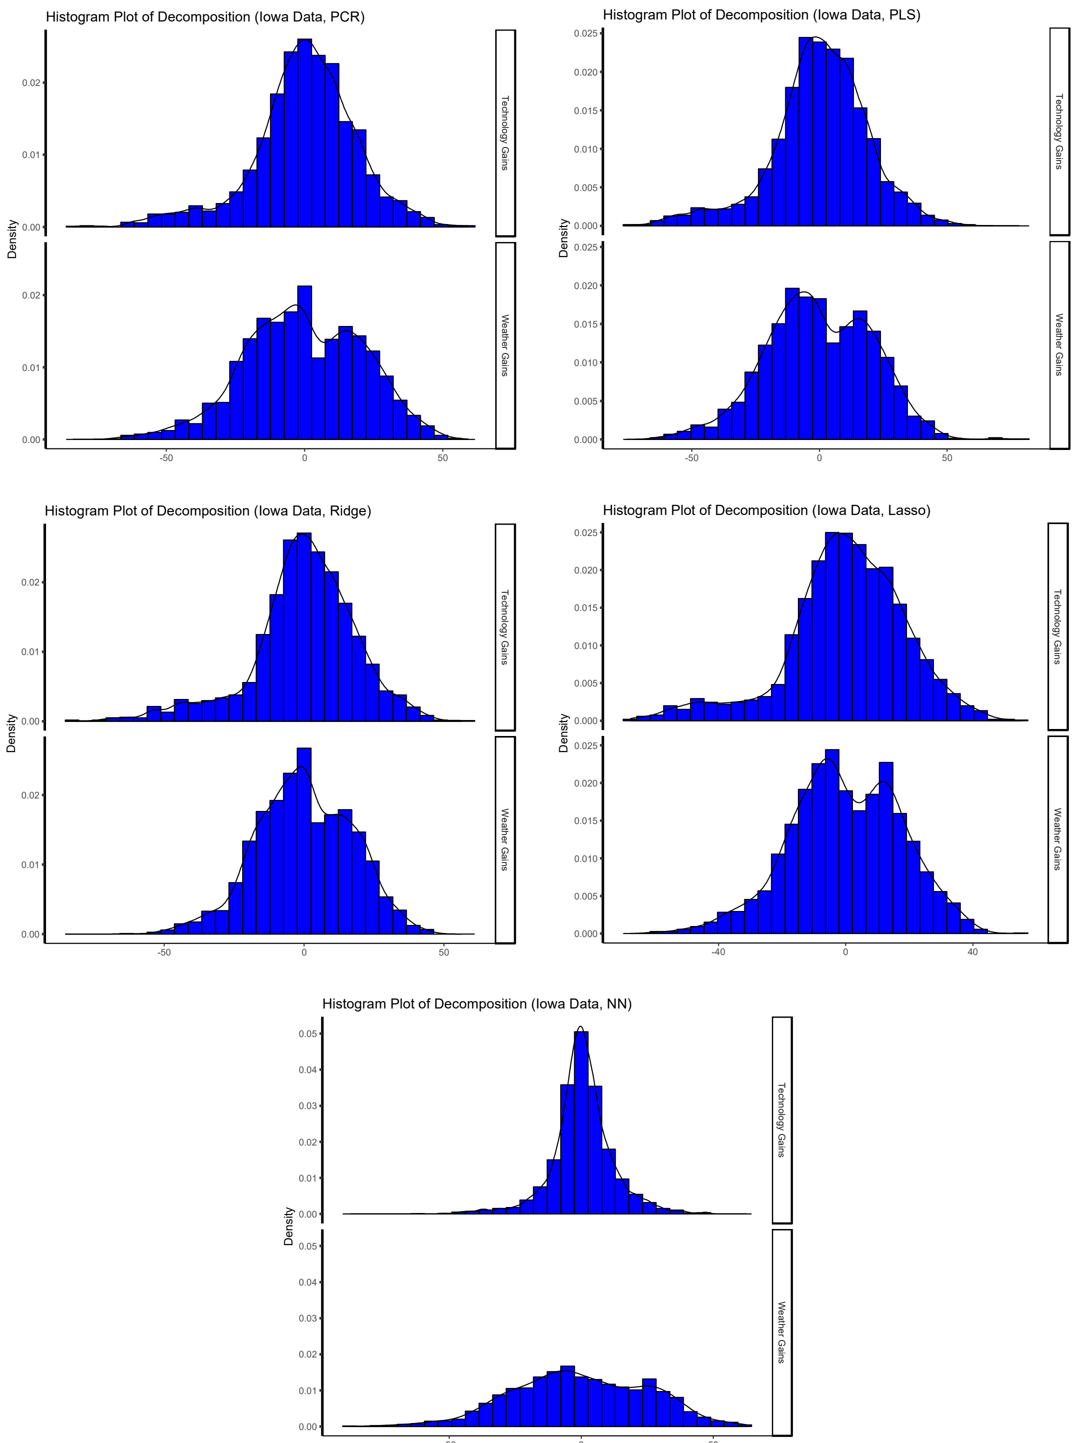

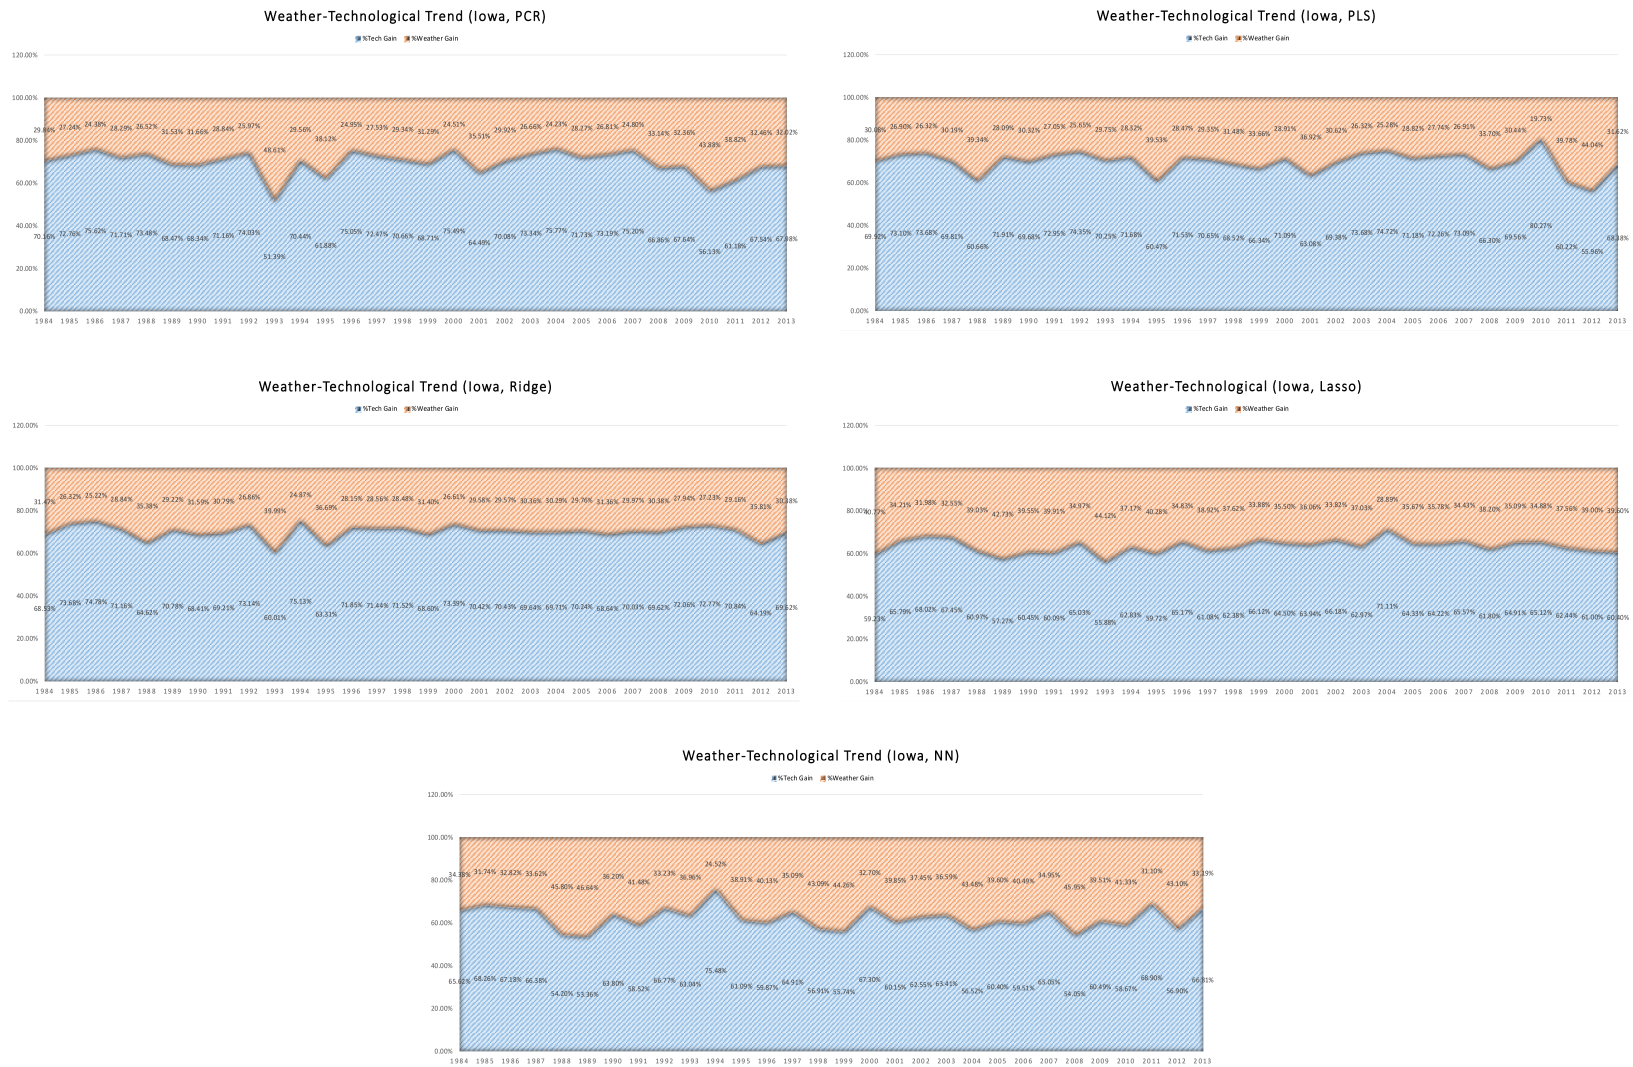

Figure 11

DECOMPOSITION FOR FIVE WEATHER MODELS OVER THE HISTORY. THE FIGURES SHOW PLOTS OF THE TIME SERIES OF THE DECOMPOSITION WITH THE IOWA DATASET BASED ON FIVE WEATHER MODELS

5.2.2 ALBERTA DATASET

The estimated technology gains and weather gains (in both absolute magnitudes and percentage) for the five machine-learning-based models for the Alberta dataset are displayed in Table 9. The technology gains on average estimated from the five models are quite stable, around 160 to 240 kg/acre. For example, according to the PCR model, the estimated average technology gain is 212.87kg/acre (or equivalently, 4.8 bu/acre) compared to the historical yield average, or equivalently 75.03% relative to the baseline linear trend model. Note that for the Alberta dataset, the NN model has a trend decomposition that is quite different from the other weather models.

Figure 12 shows the histogram plots of the decomposition with the Alberta dataset. The distribution of technology gains for the Alberta data are less asymmetric compared to the Iowa results. However, some heavy tail evidence from all of the models is still observed. For the decomposed weather gains, there is now bimodal pattern as with the Iowa data. One possible explanation is that the 16-year time series for the Alberta dataset is not sufficiently long to identify historical weather patterns. It is also possible, however, that the Alberta dataset does not contain the similar two-weather regimes observed with the Iowa dataset.

Table 9

ESTIMATED AGGREGATE TECHNOLOGY TRENDS FOR ALBTERTA DATASET

Model	Description	Tech Gains	Weather Gains	% Tech Gains	%Weather Gains
A	PCR	212.87	70.86	75.03%	24.97%
B	PLS	231.07	77.54	74.87%	25.13%
C	Ridge	237.76	71.00	77.00%	23.00%
D	Lasso	239.80	73.87	76.45%	23.55%
E	NN	163.98	82.05	66.65%	33.35%

Note: Yields and trends are measured in kg/acre.

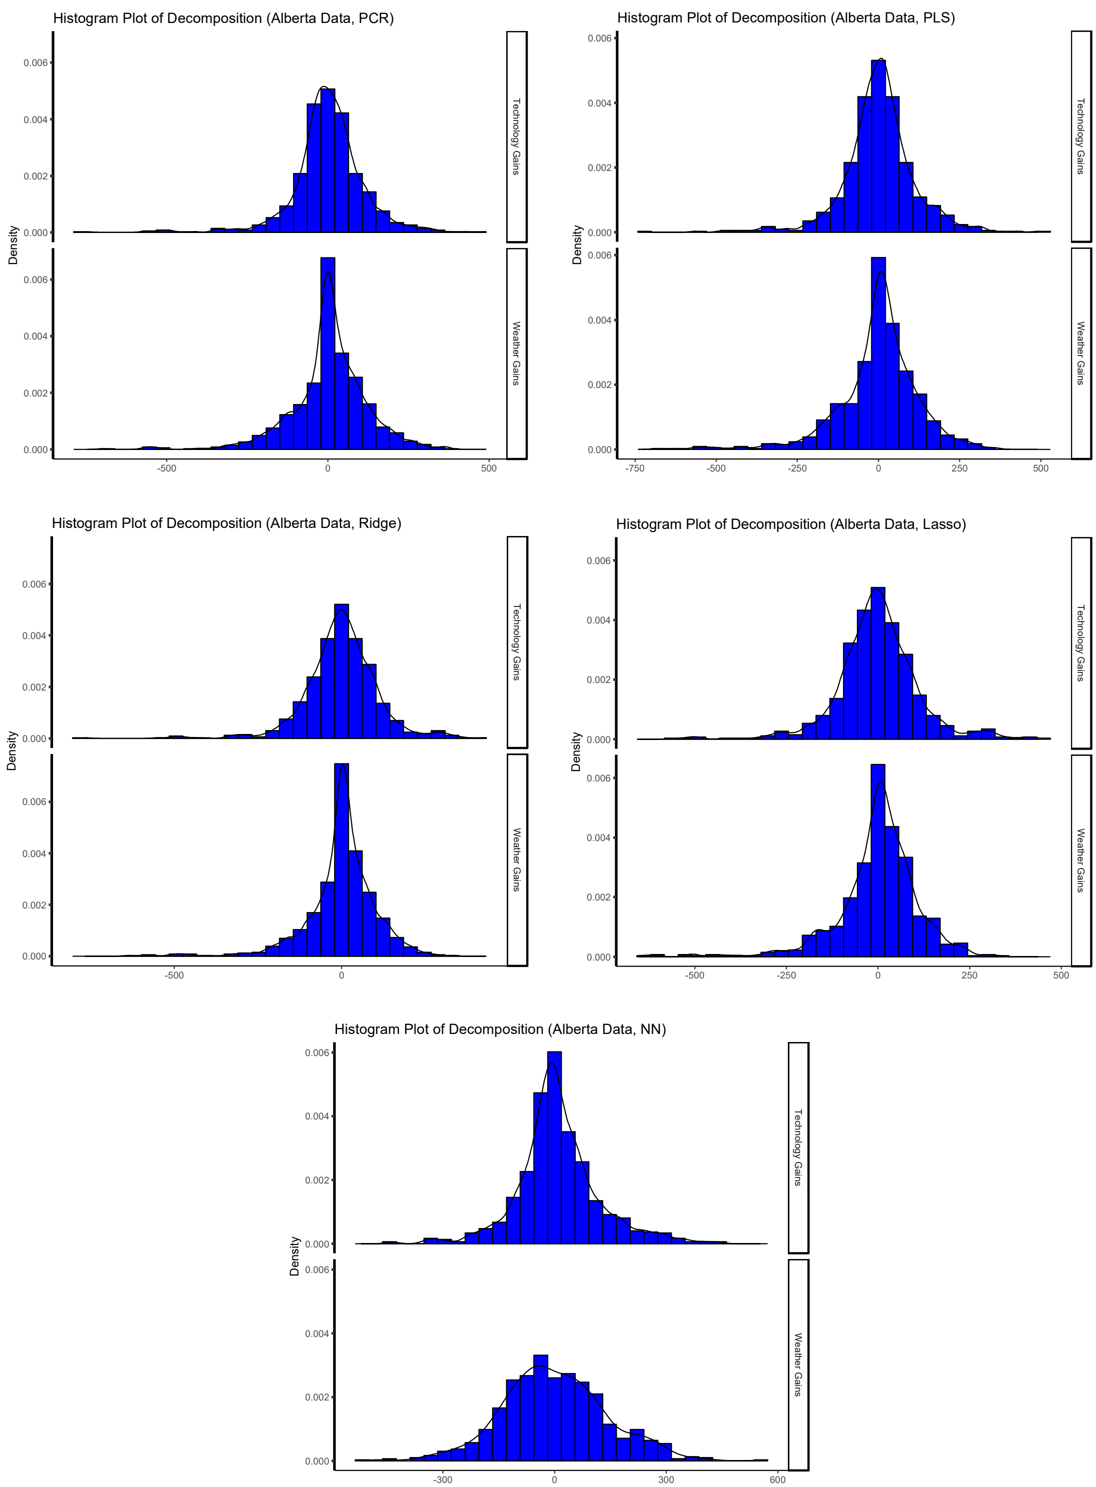

Figure 12

HISTOGRAM PLOTS OF THE DECOMPOSITION WITH THE ALBERTA DATASET BASED ON FIVE WEATHER MODELS

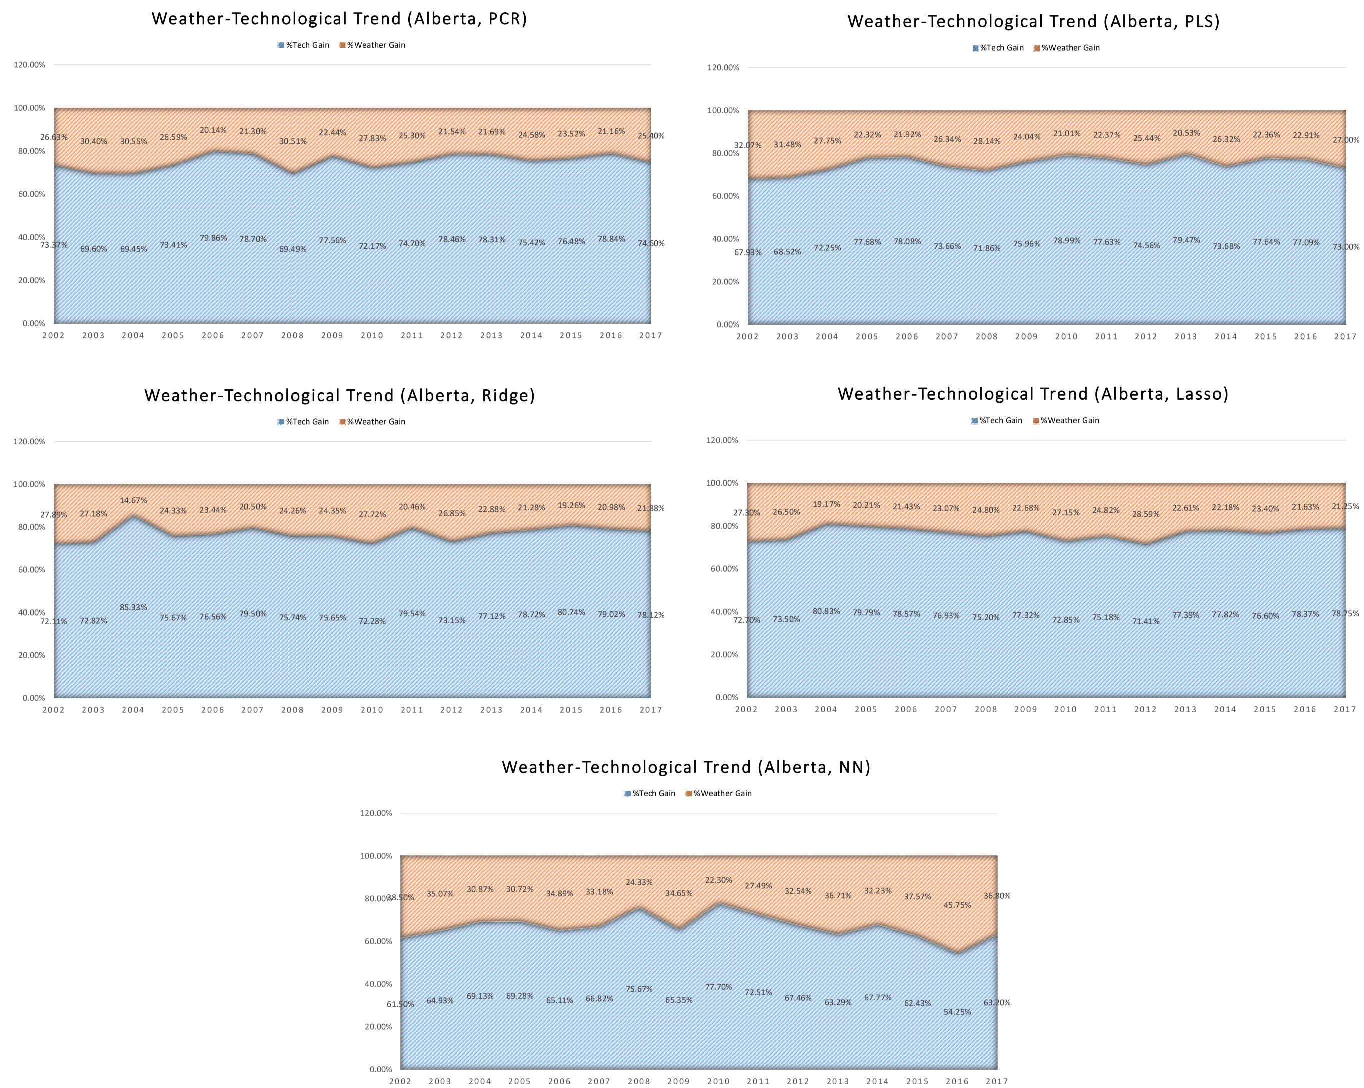

Figure 13
DECOMPOSITION FOR FIVE WEATHER MODELS OVER THE HISTORY. THE FIGURES SHOW PLOTS OF THE TIME SERIES OF THE DECOMPOSITION WITH THE IOWA DATASET BASED ON FIVE WEATHER MODELS.

Section 6: Sensitivity Analysis

A sensitivity analysis is conducted to compare results derived from the quality-controlled Alberta farm-level dataset containing 78 farms and the larger unfiltered dataset with 1,293 farms. The 1,293 farms were filtered down to 78 farms on the condition that they grew at least one 150-acre field of canola over the past 16 years. This was done to ensure the geolocated weather variables were properly extracted from the field locations. The authors use the full Alberta farm-level canola yields dataset here to estimate the LSDV model and then estimate the provincial level trend using the MM robust estimator using dummy variables and a linear time trend. The estimated models are as follows:

a) Model A: baseline model with a linear time trend estimated by OLS

b) Model B: baseline model with a linear time trend estimated by MM robust estimation

c) Model C: integrated NDVI and a linear time trend

d) Model D: weather-variable principal components derived from minimum and maximum temperatures, precipitation, short-wave radiation, with a linear time trend

Table 10 shows the yield trends estimated using the full dataset. The magnitudes of the coefficients in general are larger, and this may be due to sampling bias. For example, the 78 farms in the analysis above had a lower average canola yield compared to the full sample. The results are similar to the 78 farm sample, with the exception of the MM estimate being more similar to the estimate of the baseline OLS model. This could be because, as the sample increases, the outlier yields become less influential and the MM estimation method becomes more similar to OLS. For the NDVI model, there are similar results as the previous analysis, and the weather PC model gives similar results as well. The authors also analyzed log transformed yields as a robustness check.

Table 10
DATASET 2, CANOLA YIELDS, ALBERTA, CANADA

Model	Description	Tech Trend	Weather Trend	% Tech Trend	% Weather Trend
A	y~f(time)	28.386	–	–	–
B	y~f(time) MM	26.555	1.831	93.55%	6.45%
C	y~f(time, NDVI)	19.370	9.016	68.24%	31.76%
D	y~f(time, PC)	11.204	17.182	39.47%	60.53%

Note: All yields and trends are measured in kg/acre.

Table 11 shows the estimated trend coefficients for the log transformed yield estimation models using the full sample of farms. The noticeable difference between the log transformed models and the level yield models shown in Table 10 is that the MM model gives a lower technology trend and also the PC model gives a lower technology trend. This could be related to the transformation not being appropriate because the sample increases and the influence of the zero values and extreme yield values begins to decrease.

Table 11

DATASET 2, LOG CANOLA YIELDS, ALBERTA, CANADA

Model	Description	Tech Trend	Weather Trend	% Tech Trend	% Weather Trend
A	y~f(time)	0.084	–	–	–
B	y~f(time) MM	0.024	0.06	28.57%	71.43%
C	y~f(time, NDVI)	0.041	0.043	48.81%	51.19%
D	y~f(time, PC)	0.002	0.082	2.38%	97.62%

Note: All yields and trends are measured in kg/acre.

Section 7: Conclusion and Additional Remarks

7.1 SUMMARY

A necessary component of the crop insurance policy is determining producers' average crop yield, which is central to calculating the probable yield (PY). The PY multiplied by the coverage level sets the insured value (liability) of the insurance policy. The PY is typically calculated using an average of approximately 10 years of crop yield data for each farm or field. Some producers may desire to use a shorter yield history to compute their PYs, because the result may be more current and representative of their expected production for the current season from a technology gain perspective. However, when a shorter yield time series is used, the PY can be less stable, and rates and coverage levels may vary substantially from year to year. Another concern with using a shorter yield history is that concurrent years of favorable (or unfavorable) weather can occur and may lead to inflated crop yields over the period. When producers observe crop yield increases, they may put pressure on insurers to make crop yield trend adjustments to increase their PY and ultimately their coverage level. However, the observed yield trend may need to be first decomposed into the weather and technology effects to control for favorable or unfavorable weather conditions over the period.

This study's objective was to propose a methodology to isolate crop yield improvement due to technological change from crop yield improvement due to favorable and unfavorable weather at the farm level. Two decomposition frameworks were proposed, including a fixed-effect model and relative effect model. Several weather proxy variables were proposed that could be used to approximate the weather effects over the period and used to isolate the crop yield gain relative to the technological improvement. The proposed decomposition frameworks are empirically tested using two datasets, including (1) 37 years of county-level corn yield data for 99 counties in Iowa and (2) 16 years of farm-level canola yield in Alberta. These datasets were merged with large geospatial datasets, which were geolocated at the county and field level. Collected variables included daily values of minimum and maximum temperatures, precipitation, short wave radiation, and satellite remote sensing information.

For the first decomposition framework, six models were proposed: a baseline model estimated by OLS with a linear time trend, a robust regression using MM estimation with a linear time trend, a linear time trend and NDVI model, a linear time trend and LAI model, a linear time trend and FAPAR model, and a linear time trend and weather principal components model. For the second decomposition framework, five machine-learning-based models were used to model the weather effect: PCR, PLS, Ridge Regression, Lasso Regression, and NN. Further, the effects of aggregation were tested, and Chow tests were used as a robustness check to determine whether farm-level yield technology trends were structurally different than the aggregate trend. These robustness tests showed the proportion of farms that exhibit yield trends that are different from the aggregate-level trend. This may be useful for determining the level of aggregation technology trends should be estimated at for use in crop insurance. The main findings from the two decomposition frameworks are described next.

In general, the results of this study showed that the larger number of years of data used in the trend analysis, the more stable the trend calculation. Based on the empirical analysis presented here, it seems that with the greater number of years utilized, more of the trend is due to technology, while a smaller number of years often means more of the trend is due to weather. Another observation is that when the trend for a larger geographic regions is calculated, the trend is more stable. As well, while the trend could vary from farm to farm, in practice it would be difficult to calculate a trend for each farm due to missing data and the length of data that is required to establish a trend for each farm. Finally, another issue with trying to calculate a trend at the farm level is that at times you get a negative trend. This would present several challenges in practice, because it would be difficult to explain to producers.

7.1.1 FIXED EFFECTS APPROACH

More specifically, the empirical results using the first decomposition method (fixed-effects approach) for the first dataset—the county corn yields in Iowa—showed that the OLS method with a linear trend had a corn yield trend of 2.420 bu/acre per year. The robust MM estimation method, which is less sensitive to yields that fall outside the main trend, and the NDVI method, which uses the NDVI vegetation index to control for weather conditions, resulted in similar yield trends of 2.382 and 2.375, respectively. The weather trend over the 37 years that these models indicated was marginal. According to the NDVI model, technology increased corn yields in Iowa by 87.875 bu/acre, and weather was marginally favorable, contributing to an increase in yields of 1.665 bu/acre. A possible explanation for this is that the 37 years of crop history is a large enough sample of yields, and over time the affect of weather on corn yields begins to stabilize. However, these models may have failed to capture the affect of weather on corn yields. The weather principal component model had a technology trend of 1.950, which was lower than the other models, including the NDVI model. This implied that the weather effect was larger compared to the NDVI model over the period and accounted for approximately 20% of yield increase over the period. Results from the second dataset showed that the baseline OLS linear trend model had a technology trend of 17.563 kgs/acre per year, and the other models— robust MM regression, NDVI model and weather principal components—had a technology trend of 11.540, 10.410 and 6.118, respectively. According to these models, over the 16-year time period, the weather was favorable for canola yields, and approximately 40% of the yield gain can be attributed to favorable weather. The baseline OLS model showed a yield increase due to technology of 281 kg/acre, or 6.39 bu/acre. From the NDVI model, the increase in canola yields due to technology was 166.56 kg/acre, or approximately 3.79 bu/acre, and weather contributed to 114.45 kg/acre, or 2.601 bu/acre.

7.1.2 RELATIVE EFFECTS APPROACH

The empirical results using the second decomposition method (relative effects approach) for the first dataset, the county corn yields in Iowa, using several different machine-learning approaches to model weather found that weather contributed more to crop yield gain than the first decomposition indicated. According to the model results, positive weather over the 37-year time period contributed between 30% to 38% of relative corn yield gain over the period, and technology contributed between 62% to 69% of corn yield gain. These models used minimum and maximum temperatures, precipitation and short wave radiation to model the weather component and its effect on crop yields in a linear and nonlinear way. Research has attributed approximately 20% of corn yield gain in Iowa to a solar brightening effect. Due to less aerosols in the atmosphere over the past several decades, short wave solar radiation is able to reach corn canopies at greater intensity, and this increases crop yield. These models used in decomposition two, due to the flexibility they provide in modeling weathers effects on crop yield, may be capturing this brightening effect. However, caution should be applied to these results, because there are two components to the brightening effect. First, there is the actual increase in the availability of short wave radiation during critical periods of corn growth, which is exogenous to technology. Second, there is the adaptation of the crop to increase the day length of the stage of crop development in which additional short wave radiation is most beneficial. This is a component of plant technology. These models may not distinguish between these two separate effects and, as a result, may be capturing the changes in crop technology and ascribing their effects on yield to weather instead of technology. Results from the second dataset, Alberta canola yield data, showed that weather contributed approximately 23% to 33%, and technology contributed between 70% to 82% of canola yield gain over the 16 years. This result indicated that weather was favorable for the 16 years over Alberta for canola production.

The findings of this study may be particularly useful for crop insurance and may be used to improve crop insurance by providing more accurate PY calculations, which serve as the foundation for setting coverage levels and insurance liabilities. The proposed methodology may help isolate the influence of technological change from weather effects on crop yield over the sample period and may be useful for crop insurance in Canada, the U.S. and other countries.

7.2 RECOMMENDATIONS

In this subsection, the authors make practical recommendations for use in crop insurance based on the results obtained in this study. They first provide a detailed summary of the advantages and limitations of each proposed model, followed by the recommendations of applying each method in practice.

Decomposition 1: Advantages and Limitations

For the first decomposition framework, four models were evaluated: the full yield trend estimated by OLS; the full yield trend estimated by MM estimation; a robust estimation technique, a yield trend and NDVI model that used the normalized difference vegetation index (NDVI) as a weather proxy variable intended to capture the effects of weather on crop yield; and a yield trend model with weather PCs that are derived through principal component analysis on daily minimum and maximum temperature, precipitation and short wave radiation. Further models were testing including models that used other remote-sensing-derived indices as weather proxies; these included the LAI and FAPAR models. After evaluation on the first dataset, the authors determined those were unfit for use as weather proxies, because they seem to capture too much crop technology; as a result, the authors did not describe or test them further. The advantages and limitations of each model are summarized in Table 12.

Table 12

ADVANTAGES AND DISADVANTAGES OF DECOMPOSITION 1 YIELD TREND ESTIMATION METHODS

Model A: y~f(time) OLS	
Advantages	Provides good estimates of yield trend with a long yield history. May be an appropriate estimation method when yield history is long or when aggregation is high.
Disadvantages	Does not control for weather. Give poor estimates with short histories and may be sensitive to outliers. Trend estimates are highly variable at the farm-level. Estimates at the aggregate level may still be affected by favorable/unfavorable weather over the time period.
Model B: y~f(time) MM	
Advantages	Estimation procedure is more robust to influential yield values than OLS, but still requires a long yield history to establish yield trend.
Disadvantages	Does not model the weather effect and simply reduces the influence of weather for estimating the yield trend. Trend estimates are still highly variable at the farm-level.

Model C: y~f(time, NDVI)	
Advantages	NDVI provides a proxy for weathers affect on crop yield, high NDVI over the growing season indicates favorable growing conditions, and low NDVI indicates unfavorable growing conditions. This method may control for the weather effects so that the maximum technology yield trend can be approximated.
Disadvantages	NDVI may capture some of the technology related yield gain over the period. However, as shown using the Iowa data, corn yields in Iowa, NDVI captures only a small amount of the technological change. This may not hold for other crops, and more research is needed to determine the level of technology that is implicit in NDVI. This approach worked well and gave similar results to Model C, the robust MM estimation. Similar to the other models
Model D: y~f(time, PC)	
Advantages	Weather derived principal components are a useful way to model the effect of weather on yield, and unlike NDVI the weather derived PCs are independent from crop technology.
Disadvantages	Principal components analysis (PCA) reduced the high dimensional weather variable dataset, but after this reduction the number of PCs used to explain the variation in weather remained high. At the farm-level, with less observations per farm, the results were quite different than the other models and is likely a failure of estimating the model. An out of sample cross validation set may be necessary, however due to the scarcity of yield history, simpler approaches may be preferred.

Decomposition 2: Advantages and Limitations

For the second decomposition framework, the authors evaluated five models: PCR, PLS, Ridge Regression, Lasso regression, and NN model. The advantages and limitations of each model are summarized in Table 13.

Table 13
ADVANTAGES AND DISADVANTAGES OF DECOMPOSITION 2 YIELD TREND ESTIMATION METHODS

Model A: PCR	
Advantages	Provides good estimates of yield trend with a long yield history. May be an appropriate estimation method when yield history is long or when aggregation is high. It is an effective dimension reduction method, which can also solve the problem of colinearity.
Disadvantages	Principal components regression (PCR) identifies linear PC's in an unsupervised way, in the sense that the response variable (i.e., the crop yield) is not used to help determine the principal components while performing the regression. There is no guarantee that the selected PC's are also the best to use in predicting the yields.

Aggregation

Farm-level yield trend estimation can be sensitive to high and low yields and data sparsity. Aggregation can be used to reduce the variability of the crop yields. Pooling the farm yields reduces much of the weather variability because each farm, although spatially related, observes different weather. The farms were pooling into three different levels based on their location in reference to the Canadian ecological hierarchy. The farms were classified into three aggregation levels: provincial (highest), EcoRegion (medium) and EcoDistrict (lowest). Next, Chow tests were used to compare the aggregate yield trend relative to the farm-level yield trend to determine whether there was a structural difference. Results suggested that, compared to the other aggregation levels, the EcoRegion-level aggregation provided detailed geographic information of where yield trends differed while maintaining much of the unique farm-specific information. Other aggregation levels could be explored, and Alberta's crop risk zones should be investigated and compared to the EcoRegion aggregation level.

Recommendation

For crop insurance yield trend adjustment, the context must be considered before determining which yield trend adjustment method is appropriate. Two applications should be considered separately: actual production history trend adjustment for liability calculation and trend adjustment for premium calculation.

For liability calculation, which uses the actual production history (LTAY) with or without a trend adjustment times the producer's elected coverage level, LTAY +*trend adjustment× coverage level*, adjusting the farmers LTAY based solely on their yield history is challenging due to the high variability in crop yields. Aggregating the farms to the EcoRegion reduces the effect of weather on crop yields while maintaining much of the individual farm specific trends. For decomposing the crop yield trend, several methods produced similar results. The methods that seemed most effective were Model B; the robust MM estimator, which helps reduce the effect of weather in the yield series; Model C that used the NDVI as a weather proxy variable to capture the effect of weather on crop yield; and Model D, a yield trend model with weather PCs that are derived through principal component analysis on daily minimum and maximum temperature, precipitation, and short wave radiation. It is favorable to examine each of these methods, but for simplicity, the NDVI model seems to provide stable results.

For premium pricing, the second decomposition model is preferred. In general, the premium of insurance loss X can be in general expressed as $\Pi(\mathrm{X}) = \mathrm{E}(\mathrm{X})(1 + \Theta_{\Pi}(\mathrm{X}))$, where $\Theta_{\Pi}(X)$ is the risk loading (Zhu, et al, 2019). The relative decomposition framework provides a very convenient way to adjust for technological effect in Θ_{Π}. In this study, several methods were evaluated, including PCR, PLS, Ridge Regression, Lasso regression, and NN models. Each seem to provide relatively similar results with the exception of the NN method. In the situation when data are very limited (the Alberta case), the NN model provides very volatile results that are a far way from the other models. This indicates that due to the scarcity of yield history, alternative models may be preferred than NNs.

Section 8: Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Project Oversight Group and others for their diligent work overseeing questionnaire development, analyzing and discussing respondent answers, and reviewing and editing this report for accuracy and relevance.

Project Oversight Group Members:

Rémi Villeneuve, FSA, FCIA (Chair)

Flavien S. Apedjinou, ASA

Janet E. Duncan, FSA, MAAA, FCAS

Teddy T. Wong

Yuanjin Liu, ASA

Yukki Yeung, FSA, MAAA

At the Society of Actuaries:

W. Scott Lennox, FSA FCAS FCIA, SOA Staff Fellow

Robert P. Montgomery, ASA, MAAA, SOA Research Project Manager

Erika Schulty, Research Associate

Appendix A: Robustness Checks for Levels of Aggregation

Farm-level yield trend estimation can be sensitive to high and low yields and data sparsity. Therefore, in this section, several models of different levels of aggregation are tested, and Chow tests are conducted to determine whether the farm-level yield is structurally different from the yield trend estimated at the aggregate level.

A.1. IOWA CORN DATASET

Following the state-level regressions in Subsection 5.1, the individual county series are regressed to establish the trend values for each county. These individual county trends are tested using a Chow test to determine whether they are structurally different from the overall state trend estimated using the LSDV model. These individual county trends are evaluated by examining the minimum and maximum estimated trends, the standard deviation, and the magnitude of the linear trend coefficient. Further, each weather proxy is separately estimated in this way and evaluated. There are 99 individual counties, and each is regressed individually using the same six models without the county indicator variable. The results are summarized in Table 14. Chow tests are conducted by testing whether the maximum technology trend of the baseline model estimated separately for each county is structurally different than the trend estimated at the state level using the LSDV regression baseline model.

Table 14
INDIVIDUAL COUNTY REGRESSIONS, IOWA DATASET.

Model	Description	Technology Trend	Weather Trend	Min	Max	Sd
A	y~f(time) OLS	2.4165	-	1.8559	3.0816	0.2986
B	y~f(time) MM	2.3792	-	1.8412	3.0030	0.2755
C	y~f(time,ndvi)	2.3504	0.0661	1.6093	3.0627	0.3556
D	y~f(time,PCs)	2.0307	0.3858	1.2546	3.0170	0.3867
Chow Tests, % Structurally Different			36.36% or 36/99 counties			

At the county level with 37 years of corn yield history, the yield trends estimated from models A through D are quite stable. For example, the baseline Model A has a mean technology trend of 2.4165, a minimum of 1.8559, and a maximum of 3.0816. Also, the standard deviation of trend estimates is 0.2986. This is quite a narrow range of trend estimates. The Chow tests showed that 36/99 counties or 36.36% of counties exhibit a trend that is different from the State trend. In this case, if a yield trend adjustment was made for Iowa corn based on the aggregate-level trend shown in Table 14, then 36 counties would have either too high or too low of a trend adjustment.

A.2 ALBERTA CANOLA DATASET

Using the Alberta farm-level canola yields, varying degrees of aggregation are estimated and tested relative to this provincial trend. Individual trends grouped by EcoRegion and estimated by OLS and MM estimation are evaluated by examining the standard deviation, minimum trend, and maximum trend values. Further, each weather proxy is separately estimated in this way and evaluated. The EcoRegion trend values are then tested using the Chow procedure to determine whether they differ from the aggregate provincial level trend. Next, the EcoDistrict grouped trends and the farm-level grouped trends are estimated and evaluated using this same procedure.

EcoRegion Aggregation

Table 15 shows the yield trends estimated for farms grouped by EcoRegion, and farms are segregated into five groups and estimated. Model A, the baseline linear trend model estimated by OLS, has an estimated technology trend of 19.7967 kg/acre. This estimated technology trend translates to a 19.7967 × 17 = 336.5439 kg/acre increase over the period. Converted to bu/acre by a factor of 1/44 (approximate kg/acre to bu/acre approximation factor) results in a bu/acre increase over the period of 7.6487 bu/acre.

Table 15
ALBERTA DATASET GROUPED BY ECOREGION

Model	Description	Technology Trend	Weather Trend	Min	Max	Sd
A	y~f(time) OLS	19.7967	-	3.8446	29.7867	12.0385
B	y~f(time) MM	19.1411	0.6556	2.6974	30.6081	13.0905
C	y~f(time,ndvi)	12.6267	7.1699	0.1506	19.7895	8.7420
D	y~f(time,PCs)	8.3275	11.4691	1.0207	17.0786	7.2088

Table 16 shows the trends when estimated using the log transformation for yield. Model A has a technology yield trend of approximately 5.683% per year; this is much higher than the other models. The technology yield trend of Model B is 2% per year; and for Model C, the NDVI model, it is 2.175% per year.

Table 16
ALBERTA DATASET GROUPED BY ECOREGION WITH LOG-TRANSFORMED CANOLA YIELD

Model	Description	Technology Trend	Weather Trend	Min	Max	Sd
A	y~f(time) OLS	0.05683	-	0.00787	0.10600	0.04448
B	y~f(time) MM	0.02000	0.03682	-0.00193	0.03815	0.01758
C	y~f(time,ndvi)	0.02175	0.03507	0.00087	0.03900	0.01643
D	y~f(time,PCs)	0.02860	0.02822	0.00385	0.05861	0.02335

Table 17 shows the technology yield trend estimated for each EcoRegion group for each model. The level yields and log-transformed yields are displayed. EcoRegions 156, 158 and 157, are exhibiting large maximum technology trends with 28.82, 26.61, and 29.79 for the baseline model, respectively. Robust regression trend estimates are similar to the baseline model when nontransformed (level) yields are used. The NDVI model seems to result in lower trend estimates, and the PC model is similar. When yields are log-transformed, the robust estimator model gives more similar results to the NDVI model.

Table 17
TECHNOLOGY TREND VALUES FOR EACH ECOREGION, ALBERTA DATA

	Level Yields				Log Transformed Yields			
EcoRegion	Baseline	Robust	NDVI	PC	Baseline	Robust	NDVI	PC
156	28.82	29.00	16.57	7.90	0.1060	0.0274	0.0376	0.0453
158	26.61	26.04	19.79	14.15	0.0556	0.0317	0.0390	0.0310
157	29.79	30.61	19.73	16.35	0.0967	0.0381	0.0187	0.0588
138	3.84	2.70	0.15	1.86	0.0079	0.0047	0.0009	0.0047
149	9.92	7.36	6.90	0.07	0.0180	-0.0019	0.0126	0.0060

EcoDistrict Aggregation

The farm-level yields are grouped by 29 EcoDistricts, and 25 of these EcoDistricts had more than one farm located within the EcoDistrict boundaries. Table 18 shows the yield trends estimated for each EcoDistrict group. Model A has a higher maximum technology yield trend of 20.34281 compared to the other models. Model B gives similar results. Model C and Model D, the NDVI and the PC models, result in lower estimated yield trends.

Table 18
ALBERTA DATASET GROUPED BY ECODISTRICT.

Model	Description	Technology Trend	Weather Trend	Min	Max	Sd
A	y~f(time) OLS	20.35281	-	-7.007636	56.11393	17.26609
B	y~f(time) MM	19.99060	0.36221	-7.856972	55.98914	18.05681
C	y~f(time,ndvi)	13.93837	6.41444	-7.128040	59.53139	16.47031
D	y~f(time,PCs)	8.76522	11.58759	-22.163287	96.88465	23.91647

Table 19 shows the technology yield trends estimated using the log-transformed canola yield data for the EcoDistrict groups. Similar to Table 18, the maximum technology yield trend is highest for Model A, the baseline model. Model B has a lower trend compared to when the nontransformed level yields are used.

Table 19
ALBERTA DATASET GROUPED BY ECODISTRICT WITH LOG-TRANSFORMED CANOLA YIELD

Model	Description	Technology Trend	Weather Trend	Min	Max	Sd
A	y~f(time) OLS	0.06346	-	-0.00730	0.25865	0.07984
B	y~f(time) MM	0.02101	0.04246	-0.00700	0.06458	0.01890
C	y~f(time,ndvi)	0.02940	0.03406	-0.01423	0.09103	0.03254
D	y~f(time,PCs)	0.03460	0.02886	-0.02147	0.27624	0.06737

Table 20 shows the estimated technology yield trends estimated at the EcoDistrict grouping level using each proposed method. The level yields and the log-transformed technology yield trends are displayed. The 25 EcoDistricts that had two or more farms located within them are shown. Despite the aggregation, the maximum technology yield trends widely vary, with some trends being large and positive and others being large and negative. A higher degree of aggregation may be necessary to estimate the technology yield trends.

Table 20
TECHNOLOGY TREND VALUES FOR EACH ECODISTRICT, ALBERTA DATA.

	Level Yields				Log Transformed Yields			
EcoDistrict	Baseline	Robust	NDVI	PC	Baseline	Robust	NDVI	PC
744	18.38	21.27	9.59	-0.50	0.0364	0.0243	0.0190	0.0049
798	26.61	26.04	19.79	14.15	0.0556	0.0317	0.0390	0.0310
781	42.02	42.44	27.96	13.90	0.1890	0.0307	0.0719	0.0621
731	34.44	34.61	18.48	5.89	0.1142	0.0213	0.0448	0.0387
727	24.21	23.46	15.10	2.03	0.0377	0.0166	0.0178	0.0021
593	14.85	14.38	8.95	10.95	0.0242	0.0232	0.0139	0.0174
597	-22.51	-18.72	-28.00	-20.98	-0.0546	-0.0516	-0.0652	-0.0555
595	-1.79	-2.21	-2.58	-2.37	-0.0022	-0.0035	-0.0038	-0.0046
599	18.10	18.13	11.84	22.44	0.0258	0.0252	0.0166	0.0341
596	2.26	-0.08	-3.04	1.36	0.0194	-0.0040	0.0076	0.0226
592	1.47	-2.61	-7.13	-8.71	0.0011	-0.0012	-0.0142	-0.0164
678	11.35	11.81	11.62	-13.39	0.0119	0.0112	0.0121	-0.0079
728	11.77	8.91	-2.03	5.45	0.0259	0.0097	0.0019	0.0130
730	33.51	33.63	18.81	12.98	0.2586	0.0287	0.0819	0.0902
729	46.01	48.69	34.18	23.14	0.2365	0.0518	0.0910	0.1001
738	31.37	30.97	27.53	13.08	0.0757	0.0391	0.0612	0.0384
746	20.33	19.59	17.54	2.77	0.0283	0.0286	0.0242	-0.0031
681	-7.01	-7.86	-6.14	-22.16	-0.0073	-0.0070	-0.0064	-0.0215
791	17.59	16.97	16.85	12.21	0.0253	0.0247	0.0246	0.0184
793	12.96	22.33	3.09	5.82	0.0235	0.0445	0.0037	0.0078
588	3.21	3.80	-0.20	-1.13	0.0063	0.0085	-0.0022	-0.0042
684	11.87	17.70	2.27	254.25	0.0266	0.0340	0.0115	0.4288
692	2.36	5.45	-5.19	-8.24	0.0023	0.0103	-0.0093	-0.0158
739	-7.89	-4.74	-43.89	60.86	-0.0032	0.0005	-0.0627	0.0725

For the EcoRegion and the EcoDistrict aggregation levels the yield trends are plotted geographically in Figure 14 so the spatial relationship between maximum technology yield trend can be visually inspected. Figure 14 shows the

EcoDistrict and EcoRegion level yield trends estimated for each grouping. The highest yield trends is shown to be in the southern growing region of Alberta and the Northern growing region seems to show lower levels of trend. The EcoRegion level grouping seems to capture the spatial relationship of trend well and the spatial differences between the EcoDistrict and EcoRegion trends visually appear to be marginal. The higher level of aggregation the EcoRegion grouping seems to be appropriate.

Figure 14
ECODISTRICT AND ECOREGION LEVEL CANOLA YIELD

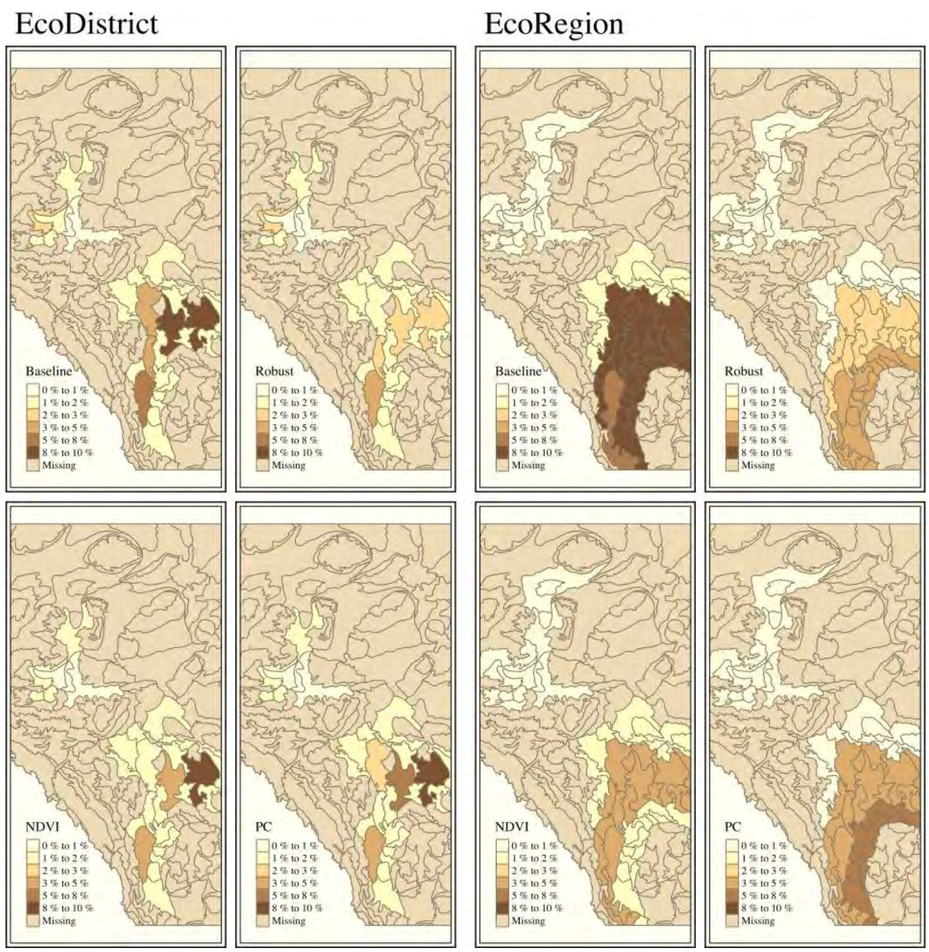

Farm-level Aggregation

There is a total of 78 farms that have continuous histories of canola production over the past 16 years. These farms are estimated separately and their trends examined. Table 21 shows the farm-level yield trends. A sensitivity analysis by estimating the different models with various levels of trend was conducted to show the difficulty of estimating farm-level yield trend when there is not a long yield history and results are displayed. As the number of years used for estimation decreases, the standard deviation of the yield trend increases and the trend estimates become unreliable. Aggregation helps reduce this effect.

Table 21
FARM-LEVEL TECHNOLOGY TREND FOR ALBERTA DATASET

Model	Description	Sample Size	Max Technology Trend	Weather Trend	Min	Max	Sd
A	y∼f(time)	16	17.547	-	-22.510	70.583	17.988
A	y∼f(time)	10	12.65	4.897	-72.076	120.106	27.285
A	y∼f(time)	5	-14.63	32.177	-243.446	213.828	67.085
B	y∼f(time) MM	16	15.296	2.251	-18.724	70.307	17.696
B	y∼f(time) MM	10	11.687	5.86	-74.138	120.724	28.500
B	y∼f(time) MM	5	-12.060	29.607	-240.033	283.895	69.002
C	y∼f(time, ndvi)	16	9.798	7.749	-27.859	64.790	15.969
C	y∼f(time, ndvi)	10	11.223	6.324	-125.613	92.539	30.675
C	y∼f(time, ndvi)	5	-28.746	46.293	-196.853	91.021	67.776
D	y∼f(time, PCs)	16	15.296	2.251	-18.724	70.307	17.696
D	y∼f(time, PCs)	10	11.687	5.860	-74.138	120.724	28.500
D	y∼f(time, PCs)	5	-12.06	29.607	-240.033	283.895	69.002

Table 22 shows the farm-level yield trends estimated with the log-transformed canola yields. A sensitivity analysis by estimating the different models with various levels of trend was conducted to show the difficulty of estimating farm-level yield trend when there is not a long yield history. As the number of years used for estimation decreases, the standard deviation of the yield trend increases, and the trend estimates become unreliable. Aggregation helps reduce this effect. Further aggregation or longer yield history may be needed to provide stable trend estimates.

Table 22

LOG-TRANSFORMED FARM-LEVEL TECHNOLOGY TREND FOR ALBERTA DATASET

Model	Description	Sample Size	Max Technology Trend	Weather Trend	Min	Max	Sd
A	y~f(time)	16	0.05734	-	-0.05464	0.26862	0.07716
A	y~f(time)	10	0.02209	0.03525	-0.34030	0.38914	0.05494
A	y~f(time)	5	-0.03011	0.08745	-1.83218	1.74753	0.14856
B	y~f(time) MM	16	0.01626	0.04109	-0.05157	0.09039	0.02301
B	y~f(time) MM	10	0.01673	0.04061	-0.11154	0.14380	0.04152
B	y~f(time) MM	5	-0.01910	0.07644	-0.43200	1.77837	0.11877
C	y~f(time, ndvi)	16	0.02122	0.03612	-0.06517	0.14093	0.03416
C	y~f(time, ndvi)	10	0.01856	0.03878	-0.03481	0.11684	0.03821
C	y~f(time, ndvi)	5	-0.03231	0.08965	-0.22285	0.12561	0.07290
D	y~f(time, PCs)	16	0.02177	0.03556	-0.05555	0.14294	0.03952
D	y~f(time, PCs)	10	0.17793	-0.12059	-0.26543	0.72766	0.19771
D	y~f(time, PCs)	5	0.05167	0.00567	-0.99038	0.76536	0.34611

Chow Test Results

Chow tests are conducted to determine whether the farm-level maximum technology trends are structurally different from the trends estimated at the more aggregated levels: 1) the province, 2) the EcoRegion and 3) the EcoDistrict. Each individual farm-level trend was tested against the trend for the province, the EcoRegion the farm is located and the EcoDistrict the farm is located. These tests may help determine the appropriate level of aggregation that should be used for crop insurance yield trend adjustment. The level of aggregation could be chosen relative to the stability of the maximum technology estimates, and also the level that the aggregate yield trend reflects producers' trends, determined by the Chow test.

Table 23 shows that as aggregation decreases the percentage of farms that have structurally different maximum technological yield trend decreases. However, results show that despite the provincial and EcoRegion level being more aggregated than the EcoDistrict level, the percentage of farms that structurally differ from the aggregate level maximum technology trend is similar. This may indicate that aggregating may be appropriate.

Table 23
DATASET 2, CHOW TEST RESULTS, CANOLA YIELDS, ALBERTA, CANADA

Province Level	
% Structurally Different	44.87% or 33/78 farms
EcoRegion Level	
% Structurally Different	37.18% or 29/78 farms
EcoDistrict Level	
% Structurally Different	32.05% or 25/78 farms

Appendix B: Individual Farm Technology Effect Estimates

The individual estimated farm technology yield trends are summarized in this section. Each of the 78 farms estimated technology canola yield trends estimated using the full yield history are displayed in Table 24. The results are highly variable between farms, with some farms having large magnitude yield trend values and others having negative values.

Table 24
DATA SET 2, GROUPED BY FARMS, MAX TECHNOLOGY TREND VALUES

	Level Yields				Log Transformed Yields			
Farm	Baseline	Robust	NDVI	PC	Baseline	Robust	NDVI	PC
1	11.60	24.44	−6.69	−14.74	0.0283	0.0214	0.0003	−0.0087
2	31.65	27.02	18.77	23.80	0.0577	0.0192	0.0331	0.0414
3	17.47	15.80	15.79	5.03	0.0310	0.0167	0.0270	0.0132
4	47.26	49.71	23.32	12.20	0.2386	0.0487	0.0409	0.1028
5	42.15	43.94	36.83	26.77	0.0965	0.0376	0.0770	0.0680
6	20.12	−1.20	−3.58	−16.50	0.0774	−0.0027	0.0093	0.0116
7	19.25	19.56	14.68	9.22	0.0888	0.0490	0.0550	0.0400
8	46.29	47.82	37.73	7.41	0.2478	0.0640	0.0737	0.0859
9	33.43	33.44	25.53	13.75	0.0349	0.0237	0.0232	0.0125
10	37.95	30.02	34.20	9.41	0.0546	0.0224	0.0451	0.0016
11	12.15	9.91	2.52	−4.12	0.0699	0.0008	0.0136	0.0153
12	13.19	15.60	3.08	12.36	0.0127	0.0210	−0.0046	0.0121
13	22.57	23.63	16.95	18.24	0.0527	0.0531	0.0403	0.0372
14	0.25	0.33	−8.22	−1.41	−0.0006	0.0001	−0.0121	−0.0023
15	−22.51	−18.72	−28.00	−20.98	−0.0546	−0.0516	−0.0652	−0.0555
16	25.38	24.72	25.65	29.54	0.0359	0.0362	0.0362	0.0423
17	18.10	18.13	11.84	22.44	0.0258	0.0252	0.0166	0.0341
18	30.41	30.68	31.87	10.91	0.0433	0.0431	0.0456	0.0108
19	−10.71	−10.91	−12.75	-7.35	−0.0210	-0.0202	-0.0245	−0.0124
20	−0.20	−2.00	−5.67	−8.26	−0.0036	−0.0044	−0.0123	−0.0154
21	−8.34	−12.30	−8.13	−7.54	−0.0179	−0.0245	−0.0174	−0.0172
22	−15.26	−15.47	−21.37	−26.01	−0.0253	−0.0247	−0.0355	−0.0394
23	7.06	5.88	7.19	−7.58	0.0071	0.0059	0.0072	−0.0068

Table 24
DATA SET 2, GROUPED BY FARMS, MAX TECHNOLOGY TREND VALUES (CON'D)

	Level Yields				Log Transformed Yields			
Farm	Baseline	Robust	NDVI	PC	Baseline	Robust	NDVI	PC
24	32.05	28.18	17.29	2.27	0.0420	0.0136	0.0199	0.0041
25	4.56	5.88	2.74	−0.18	0.0096	0.0109	0.0060	0.0014
26	9.63	10.77	3.57	6.28	0.0102	0.0111	0.0039	0.0069
27	−16.88	−15.50	−16.18	−20.78	−0.0253	−0.0245	−0.0243	−0.0291
28	−9.01	−8.33	−8.93	−1.21	−0.0152	−0.0159	−0.0153	−0.0045
29	−6.28	−7.66	−6.28	5.67	−0.0091	−0.0100	−0.0091	0.0056
30	5.64	2.13	−6.21	−3.99	0.0223	−0.0069	−0.0039	0.0026
31	50.47	52.35	38.60	23.50	0.2337	0.0494	0.1035	0.0975
32	29.20	26.16	20.28	20.93	0.2168	0.0346	0.0516	0.1135
33	35.50	14.01	17.70	17.65	0.0566	0.0118	0.0233	0.0270
34	70.59	70.31	64.56	56.31	0.2686	0.0904	0.1409	0.1429
35	14.75	13.76	15.20	7.92	0.0228	0.0216	0.0224	0.0140
36	25.78	9.30	8.25	−13.24	0.0726	0.0045	0.0229	−0.0047
37	24.22	19.23	11.41	−1.89	0.0800	0.0174	0.0288	0.0173
38	32.17	13.28	19.05	2.02	0.2134	0.0130	0.0999	0.0928
39	55.79	35.68	31.85	22.33	0.1322	0.0321	0.0487	0.0569
40	56.53	55.20	52.28	47.39	0.0591	0.0425	0.0517	0.0476
41	26.30	23.46	17.42	−1.22	0.0732	0.0168	0.0437	0.0126
42	32.07	30.34	22.28	8.22	0.2299	0.0412	0.1036	0.0547
43	21.43	17.32	3.19	−10.11	0.2044	0.0071	0.0397	0.0568
44	35.06	35.25	22.43	9.02	0.0749	0.0576	0.0422	0.0200
45	19.67	17.52	9.63	−2.12	0.0491	0.0179	0.0209	−0.0007
46	0.91	1.11	−1.92	−11.50	0.0002	0.0015	−0.0041	−0.0230

Table 24
DATA SET 2, GROUPED BY FARMS, MAX TECHNOLOGY TREND VALUES (CON'D)

	Level Yields				Log Transformed Yields			
Farm	Baseline	Robust	NDVI	PC	Baseline	Robust	NDVI	PC
47	25.86	22.22	18.08	12.85	0.0308	0.0052	0.0220	0.0174
48	11.04	11.30	10.26	−10.78	0.0080	0.0082	0.0072	−0.0117
49	2.37	3.60	−2.96	−8.89	0.0123	0.0051	0.0003	−0.0133
50	13.00	12.55	1.61	−1.98	0.0307	0.0285	0.0072	−0.0025
51	-4.13	−2.75	−9.69	−8.06	−0.0087	-0.0064	−0.0192	−0.0100
52	−12.05	−11.77	−10.68	−29.97	−0.0114	−0.0102	−0.0100	−0.0279
53	1.80	6.37	−6.54	−6.71	0.0062	0.0103	-0.0106	-0.0109
54	8.72	9.30	5.68	1.79	0.0169	0.0162	0.0100	-0.0035
55	36.68	35.78	19.60	−0.96	0.0991	0.0202	0.0265	0.0257
56	1.32	−13.33	−1.60	−0.33	-0.0021	-0.0163	−0.0056	−0.0044
57	4.47	−1.95	−0.33	17.89	0.0055	0.0002	−0.0004	0.0269
58	30.22	26.96	21.98	29.83	0.0486	0.0487	0.0352	0.0496
59	14.58	14.63	3.68	16.62	0.0227	0.0226	0.0065	0.0291
60	5.63	5.11	4.91	4.25	0.0101	0.0112	0.0086	0.0077
61	15.37	15.31	10.91	6.39	0.0197	0.0197	0.0135	0.0077
62	8.10	4.57	3.91	8.56	0.0130	0.0107	0.0044	0.0155
63	21.25	20.90	16.61	6.67	0.0281	0.0246	0.0205	0.0035
64	12.96	22.33	3.09	5.82	0.0235	0.0445	0.0037	0.0078
65	27.57	27.65	7.88	1.77	0.0397	0.0156	0.0040	-0.0075
66	37.36	35.38	23.41	7.14	0.2222	0.0283	0.0795	0.0906
67	18.44	5.73	10.99	1.40	0.0206	0.0045	0.0119	0.0022
68	34.39	33.26	26.66	27.37	0.0881	0.0296	0.0596	0.0684
69	15.50	18.98	7.75	5.13	0.0274	0.0310	0.0150	0.0132

Table 24
DATA SET 2, GROUPED BY FARMS, MAX TECHNOLOGY TREND VALUES (CON'D)

Level Yields					Log Transformed Yields			
Farm	Baseline	Robust	NDVI	PC	Baseline	Robust	NDVI	PC
70	−1.23	−0.59	−7.04	−7.48	−0.0021	−0.0012	−0.0152	−0.0151
71	1.90	3.51	−2.39	−3.29	0.0026	0.0103	−0.0082	−0.0103
72	20.46	12.55	12.22	6.19	0.2102	0.0282	0.0492	0.0707
73	16.02	12.42	3.61	9.94	0.1008	0.0129	0.0642	0.1249
74	9.18	15.31	10.43	10.29	0.0197	0.0308	0.0210	0.0174
75	30.52	26.90	5.12	19.70	0.2158	0.0223	0.0066	0.1106
76	4.17	−14.00	1.85	9.59	−0.0010	−0.0250	−0.0049	0.0072
77	18.71	18.53	16.38	10.67	0.0443	0.0304	0.0390	0.0263
78	10.37	9.69	9.10	7.58	0.0269	0.0136	0.0241	0.0200

References

Adhikari, Shyam, Thomas O., Knight, and Eric J., Belasco. 2012. Evaluation of Crop Insurance Yield Guarantees and Producer Welfare With Upward-Trending Yields. *Agricultural and Resource Economics Review* 41, no. 3:367–376.

Barnett, Barry J., J. Roy, Black, Yingyao, Hu, and Jerry R., Skees. 2005. Is Area Yield Insurance Competitive With Farm Yield Insurance? *Journal of Agricultural and Resource Economics* 30, no.2:285–301.

Cheng, Bing, and D. M. Titterington. 1994. Neural Networks: A Review From a Statistical Perspective. *Statistical Science* 9, no. 1:2–30.

Coble, Keith, Thomas O., Knight, Mary Frances, Miller, Barry J., Goodwin, Roderick M., Rejesus, et al. 2013. Estimating Structural Change in US Crop Insurance Experience. *Agricultural Finance Review* 73, no. 1:74–87.

Finger, Robert. 2010. Revisiting the Evaluation of Robust Regression Techniques for Crop Yield Data Detrending. *American Journal of Agricultural Economics* 92, no. 1:205–211.

———. 2013. Investigating the Performance of Different Estimation Techniques for Crop Yield Data Analysis in Crop Insurance Applications. *Agricultural Economics* 44, no. 2:217–230.

Friedman, Jerome, Trevor, Hastie, and Robert, Tibshirani. 2001. *The Elements of Statistical Learning*. Vol. 1. New York: Springer.

Goodwin, Barry K., and Olivier, Mahul. Risk Modeling Concepts Relating to the Design and Rating of Agricultural Insurance Contracts. *The World Bank*, August 2004, https://elibrary.worldbank.org/doi/abs/10.1596/1813-9450-3392 (accessed on November 18, 2019).

James, Garth, Daniela, Witten, Trevor, Hastie, and Robert, Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 112. New York: Springer.

Just, Richard E., and Quinn, Weninger. 1999. Are Crop Yields Normally Distributed? *American Journal of Agricultural Economics* 81, no. 2:287–304.

Ker, Alan P., and Barry K., Goodwin. 2000. Nonparametric Estimation of Crop Insurance Rates Revisited. *American Journal of Agricultural Economics* 82, no. 2:463–478.

Lobell, Daniel B., and Gregory P., Asner. 2003. Climate and Management Contributions to Recent Trends in US Agricultural Yields. *Science* 299, no. 5609:1032–1032.

Lobell, Daniel B., Wolfram, Schlenker, and Justin, Costa-Roberts. 2011. Climate Trends and Global Crop Production Since 1980. *Science* 333, no. 6042:616–620.

Ozaki, Victor A., Barry K., Goodwin, and Ricardo, Shirota. 2008. Parametric and Nonparametric Statistical Modelling of Crop Yield: Implications for Pricing Crop Insurance Contracts. *Applied Economics* 40, no. 9:1151–1164.

Ramirez, Octavio A., Sukant, Misra, and James, Field. 2003. Crop-Yield Distributions Revisited. *American Journal of Agricultural Economics* 85, no. 1:108–120.

Ripley, B. D. 1993. Statistical Aspects of Neural Networks. *Networks and Chaos—Statistical and Probabilistic Aspects* 50:40–123.

Sherrick, Bruce J., Fabio C., Zanini, Gary D., Schnitkey, and Scott H., Irwin. 2004a. Crop Insurance Valuation under Alternative Yield Distributions. *American Journal of Agricultural Economics* 86, no. 2:406–419.

———. 2004b. Crop Insurance Valuation Under Alternative Yield Distributions. *American Journal of Agricultural Economics* 86, no. 2:406–419.

Skees, Jerry R., and Michael R., Reed. 1986. Rate Making for Farm-Level Crop Insurance: Implications for Adverse Selection. *American Journal of Agricultural Economics* 68, no. 3:653–659.

Tack, Jesse, Ardian, Harri, and Keith, Coble. 2012. More than Mean Effects: Modeling the Effect of Climate on the Higher Order Moments of Crop Yields. *American Journal of Agricultural Economics* 94, no. 5:1037–1054.

Tollenaar, Matthijs, Jon, Fridgen, Priyanka, Tyagi, Paul W., Stackhouse Jr, and Saratha, Kumudini. 2017. The Contribution of Solar Brightening to the US Maize Yield Trend. *Nature Climate Change* 7, no. 4:275.

Vermote, Eric. 2015. MOD09A1 MODIS/Terra Surface Reflectance 8-Day L3 Global 500m SIN Grid V006. Distributed by NASA EOSDIS Land Processes DAAC.

Vermote, Eric, Chris, Justice, Ivan, Csiszar, Jeff, Eidenshink, Ranga, Myneni, et al. 2014. NOAA Climate Data Record (CDR) of Normalized Difference Vegetation Index (NDVI), Version 4. NOAA National Climate Data Center.

Viovy, Nicholas, O., Arino, and Alan, Belward. 1992. The Best Index Slope Extraction (bise): A Method for Reducing Noise in NDVI Time-Series. *International Journal of Remote Sensing* 13, no. 8:1585–1590.

Woodard, Joshua D. 2014. Impacts of Weather and Time Horizon Selection on Crop Insurance Ratemaking: A Conditional Distribution Approach. *North American Actuarial Journal* 18, no. 2:279–293.

Woodard, Joshua D., Gary D., Schnitkey, Bruce J., Sherrick, Nancy, Lozano-Gracia, and Luc, Anselin. 2012. A Spatial Econometric Analysis of Loss Experience in the US Crop Insurance Program. *Journal of Risk and Insurance* 79, no. 1:261–286.

Woodard, Joshua D., Bruce J., Sherrick, and Gary D., Schnitkey. 2011. Actuarial Impacts of Loss Cost Ratio Ratemaking in US Crop Insurance Programs. *Journal of Agricultural and Resource Economics* 36, no. 1:211–228.

Xiao, Zhiqiang, Shunlin, Liang, TongTong, Wang, and Bo, Jiang. 2016. Retrieval of Leaf Area Index (LAI) and Fraction of Absorbed Photosynthetically Active Radiation (FAPAR) from VIIRS Time-Series Data. *Remote Sensing* 8, no. 4:351.

Yu, Tian, and Bruce A., Babcock. 2010. Are US Corn and Soybeans Becoming More Drought Tolerant? *American Journal of Agricultural Economics* 92, no. 5:1310–1323.

Zhu, WenJun, Lysa, Porth, and Ken Seng, Tan. 2015. A Credibility-Based Yield Forecasting Model for Crop Reinsurance Pricing and Weather Risk Management. *Agricultural Finance* Review, forthcoming.

———. 2019. Agricultural Insurance Ratemaking: Development of a New Premium Principle. *North American Actuarial Journal*, forthcoming.

Zhu, Ying, Barry K., Goodwin, and Sujit K., Ghosh. 2011. Modeling Yield Risk Under Technological Change: Dynamic Yield Distributions and the US Crop Insurance Program. *Journal of Agricultural and Resource Economics* 36, no.1:192–210.

About The Society of Actuaries

The Society of Actuaries (SOA), formed in 1949, is one of the largest actuarial professional organizations in the world dedicated to serving more than 30,000 actuarial members and the public in the United States, Canada and worldwide. In line with the SOA Vision Statement, actuaries act as business leaders who develop and use mathematical models to measure and manage risk in support of financial security for individuals, organizations and the public.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and nonactuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.

Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org