Cover Crop Adoption and Climate Risks:

An Application of Causal Random Forests

David T. Quigley, University of Texas at Arlington, david.quigley2@uta.edu

Yuyuan Che, North Carolina State University, yche3@ncsu.edu

Mahmut Yasar, University of Texas at Arlington, myasar@uta.edu

Roderick M. Rejesus, North Carolina State University, rmrejesu@ncsu.edu

Selected Paper prepared for presentation at the 2023 Agricultural & Applied Economics Association Annual Meeting, Washington DC; July 23-25, 2023

Cover Crop Adoption and Climate Risks: An Application of Causal Random Forests

David T. Quigley, Yuyuan Che, Mahmut Yasar, and Roderick M. Rejesus

May 16, 2023

Abstract

This study applies the causal random forest approach to investigate the extent to which persistent warming temperatures might induce farmers to utilize cover crops. Causal random forests is a novel statistical approach that allows for consistent and flexible estimation of heterogeneous treatment effects. The causal random forest technique is then applied to a unique satellite-based panel data set on cover crop adoption at pixel-level resolution. On average, we find evidence that being exposed to higher-than-average maximum temperatures at the end of the harvest season for multiple consecutive years influence farmers' decisions to utilize cover crops. Moreover, our causal random forest analysis suggest that there is heterogeneity in the effect of persistent end-of-season heat events, and the heterogeneity is likely driven by precipitation and drought conditions at different points of the growing season. The evidence points to farmers increasingly utilizing cover crops voluntarily as weather patterns associated with global warming become increasingly common.

Keywords: Causal machine learning, extreme heat, generalized random forest, treatment effect heterogeneity

JEL Classifications: C8, Q12, Q18

1. Introduction

Climate change has become an important issue for agricultural production (Tack et al., 2015; Urban et al., 2015; D'Agostino and Schlenker, 2016). A large and growing literature has shown that extreme temperatures have negative effects on crop yields and will likely bring dramatic shifts in production agriculture (Burke and Emerick, 2016; Carleton and Hsiang, 2016; Chen et al., 2016; Zhang et al., 2017). For example, a warming increase of 2°C is expected to reduce agricultural output by almost 25% (IPCC, 2014). In addition, higher temperatures are predicted to lead to more intense rainfall events and greater occurrence of severe soil erosion and soil nutrient loss (Nearing et al., 2004; O'Neal et al., 2015). Therefore, reducing greenhouse gas (GHG) emissions and sequestering carbon in agriculture have been an important policy goal for a number of countries in order to mitigate the adverse effects of global warming. For instance, the United States (US) Department of Agriculture (USDA) recently provided up to $1 billion for pilot projects that would support adoption of climate-smart practices (including cover crops), with the intent of helping reduce the adverse impacts of climate-change-induced weather events (USDA, 2022).

Planting cover crops is regarded as a climate-smart soil health practice that can contribute to climate resilience in agriculture. Cover cropping as a practice has become increasingly relevant over the last decade as global warming changes the climate in which crops are grown (Arbuckle and Roesch-McNally, 2015; Kaye and Quemada, 2017; Olson et al, 2017; Shackelford et al, 2019). Cover crops are typically legumes, grasses, and brassicas that are planted to cover the soil between the growing seasons of regular crop production (Schnepf and Cox, 2006; Arbuckle and Roesch-McNally, 2015). Planting cover crops is expected to improve soil structure and soil health by increasing carbon sequestration potential and enhancing moisture infiltration and retention capacity (Lal, 2015). In addition, cover crops are often depicted as a win-win practice for both economic productivity and the environment. It has been shown that cover crops have the potential to increase long-term crop yields and reduce the input costs of tillage and fertilization (Montgomery, 2017; Shackelford et al., 2019). In addition, cover crops has the potential to provide large-scale environmental benefits by reducing soil erosion, assimilating excess nutrients, improving water quality, and providing habitat for beneficial insects and pollinators (Snapp et al., 2005; Schnepf and Cox, 2006; Tonitto et al., 2006; Laloy and Bielders, 2010; Castellano et al., 2012; Kladivko et al., 2014; Poeplau and Don, 2015; Hanrahan et al., 2018).

Despite the potential productivity and environmental benefits from planting cover crops, adoption of this soil health practice remains limited in the US. The 2017 US Census of Agriculture indicate that cover crop acres are only about 3.9% of all cropland acres nationwide (Zulauf and Brown, 2019). Previous literature has identified a range of factors that affect farmers' decision to adopt cover crops, as well as the challenges that may explain the relatively low adoption rate. Farmers' adoption of cover crops is mainly motivated by the private benefits of improving soil health and reducing soil erosion (Arbuckle and Roesch-McNally, 2015). In contrast, some studies suggest that uncertainty of net economic returns or cash crop yields in the short run is a barrier to the adoption of cover crops (Schipanski et al., 2014; Bergtold et al., 2017). Other possible factors associated with relatively low cover crop adoption in the US include biophysical factors (e.g., operating on steeper slopes, complementarity with other practices) (Lee and McCann, 2019), structural factors (e.g., farm size, lack of time or equipment) (Arbuckle and Roesch-McNally, 2015; Dunn et al., 2016; Roesch-McNally et al., 2018a, b), and policy support (Dunn et al., 2016; Park et al., 2022).

The objective of this study is to investigate whether farmers' decision to adopt cover crops is influenced by occurrence of persistent warming events (i.e., multiple years of high temperatures at harvest). We utilize a large-scale satellite data set that has cover crop adoption information at the pixel level for soybean fields in Indiana for the 2008-2015 period (Seifert et al., 2019). The cover crop data is then merged with pixel-level weather variables (i.e., from the Daymet data set (see Thornton et al., 2014)), yield information, and soil quality measures (i.e., the National Commodity Crop Productivity Index (NCCPI)). A causal random forest estimation procedure is then used to estimate the impact of warming events on cover crop usage by farmers. Causal random forests allow for estimation and investigation of heterogeneous causal treatment effects.

There are very few studies that have investigated the relationship between farmers' adoption of cover crops and climate change (Gardezi and Arbuckle 2019; Lee and McCann 2019; Yoder et al., 2021; Dong, 2022; Kathage et al., 2022). Farmers' cover crop adoption decisions in response to weather events associated with climate change are found to be mixed (Yoder et al., 2021). Some research suggests that farmers are motivated to adopt cover crops after experiencing particularly acute climate risks, such as greater precipitation (Lee and McCann, 2019). However, other research finds that climate change considerations do not play a significant role in farmers' decisions regarding cover crop adoption (Kathage et al., 2022). Dong (2022) also find that previous occurrence of drought events do not affect farmers' cover crop adoption decisions. In addition, some studies suggest that cover crops are inadequate to address climate change impacts (Arbuckle and Roesch-McNally, 2015; Houser and Stuart., 2020). Others also suggest that planting cover crop introduce new risks and management difficulties, which can exacerbate climate change's negative impact on agricultural production (Roesch-McNally et al., 2018a). Overall, this literature indicates that farmers can be motivated to adopt cover crops in order to address climate risks, but at the same time, they can also be discouraged from adopting cover crops due to uncertainty in its benefits and possible introduction of new management risks (Yoder et al., 2021).

Despite the existing literature on cover crop adoption and climate risks, to the best of our knowledge, there has been no study that quantitatively examines how farmers' adoption of cover crops is affected by persistent occurrence of extreme heat events. We contribute to the literature in this regard. Previous literature typically evaluates the effects of general climate change considerations or specific weather events (such as excessive precipitation and drought, rather than persistent multi-year extreme heat events) on cover crop adoption. We focus on persistent extreme heat events since higher temperatures are directly correlated with global warming and they are always accompanied by frequent occurrences of other adverse events, such as soil erosion. This study also complements existing agronomic studies arguing that cover crops improve soil physical properties over time, which in turn enhance resilience to extreme heat events and mitigate soil erosion. For example, agronomists have argued that cover crops can modulate extremes in soil temperature in hotter regions and thus adapt to climate change (Kaye and Quemada, 2017). To be specific, cover crop residues can reduce soil temperatures (Scholberg et al., 2010) and standing cover crops reduce the amplitude of temperature variation because of the increased boundary layer effect of the canopy compared to bare soil (Dabney et al., 2001). In addition, erosion control is one of the core services that cover crops are typically seen to provide (Prado Wildner et al., 2004).

Second, we also contribute to the literature through the application of the causal random forest approach on a unique remote-sensing based data (i.e., from satellites) at a disaggregate pixel level. The causal random forest approach is a relatively new machine learning approach that allows for assessing heterogeneous causal relationships (i.e., as compared to traditional machine learning methods that are more concerned with prediction). To the best of our knowledge, only one recent study has used this method in the context of cover crop use (Deines et al., 2022). Traditional statistical approaches are limited to estimating average treatment effects, which can hide significant variation in how the treatment impacts different observations. Moreover, the use of a dense pixel-level data set based on satellite images is also a contribution since only a few studies have used satellite-based cover crop information for economic analysis (Seifert et al., 2019; Chen et al., 2021; Connor et al., 2021). Most studies that have examined cover crop adoption typically use either survey data sets (Lee and McCann 2019) or more aggregate (e.g., county-level) data sets (Chen et al., 2021; Connor et al., 2021; Chen et al. 2023).

Lastly, since our study examines whether experiencing persistent higher-than-average maximum temperature events for multiple consecutive years would encourage (or discourage) farmers to adopt cover crops, we also contribute to the literature that have investigated how experiencing recent extreme events can influence economic decision-making. There is a large literature showing how people's recent extreme experience has a role to play in the evaluation of subsequent choices, especially in different types of insurance markets (Cai and Song, 2017; Kousky, 2017; Bjerge and Trifkovic, 2018; Che et al., 2020). The general finding of these studies is that farmers are more likely to buy insurance (or use a risk mitigating instrument) after experiencing extreme weather events. Therefore, our study provides another context for which to evaluate whether recent extreme weather events influence adoption of cover crops (i.e., a practice that the agronomic literature consider a tool that can help manage or mitigate climate-related risks, much like crop insurance).

Findings from this study suggest that being exposed to persistent higher-than-average maximum temperatures at the end of the harvest season likely induces farmers to utilize cover crops in their production system. We find evidence that this treatment effect is heterogenous and this heterogeneity appears to be driven by precipitation and drought conditions at various points in the growing season. The evidence points to farmers increasingly utilizing cover crops voluntarily as weather patterns associated with global warming (i.e., persistent high temperatures) become increasingly common. This insight can help provide a better understanding of farmer cover crop adoption behavior, and may help various institutions interested in encouraging more cover crop use (e.g., non-profit conservation groups, conservation-oriented government agencies, and university extension & outreach programs), especially with regards to targeting the timing and location of their educational programming.

2. Data and Empirical Approach

2.1 Data Description

The data for this study was created from satellite imagery as detailed in Seifert et al. (2019). Seifert et al. (2019) kindly provided their dataset for use in this study. For our analysis, we specifically focus on yearly data for soybeans in Indiana from 2008 to 2015. Individual observations are satellite imagery pixels corresponding to a growing season for a particular year. Individual pixels are 30m x 30m in size.

Cover crop use (*CC*) is the main dependent variable of interest in this study. It is a binary variable indicating whether cover crop usage is detected for a particular pixel before the next planting season. Detailed descriptions of how the cover crop data was generated can be found in Seifert et al. (2019). As reported in Seifert et al. (2019), the satellite-based data set developed had an out-of-sample accuracy of 91.2% and compares favorably with temporal and spatial trends observed in other third-party cover crop data sets (i.e., those from the Environmental Working Group (EWG) and USDA Natural Resource Conservation Service (NRCS)).

Aside from the main cover crop dependent variable, the Seifert et al. (2019) study also utilized the Daymet gridded daily weather information (see Thronton et al, 2014) to produce weather data that corresponds to the individual pixel-level observations for cover crops. In particular the weather variables in the data set includes the following: precipitation (*PRCP*) in meters, maximum temperature (*TMAX*) in °C, minimum temperature (*TMIN*) in °C, and vapor pressure (*VP*) and vapor pressure deficit (*VPD*) measures (both in kPa) for drought conditions. Each weather variable is measured before planting (*PRE-PLANT*), early in the season (*EARLY*), mid-season (*MID*), and at the end of the season (*END*). That is, there are four weather measurements at various points of the year. The summary statistics for the weather variables in the data are presented in Table 1.

Given our interest in the impact of persistent heat events on cover crop adoption decisions, it is critical to first define the "treatment" variables that correspond to occurrence of persistent heat events. Thus, based on the pixed-level weather data in Seifert et al. (2019), we create three binary treatment variables that correspond to whether or not a particular pixel-year observation experienced above average maximum temperatures at the end of the season for: (i) three consecutive years (*TMAX_END3*), (ii) four consecutive years (*TMAX_END4*);, and (iii) five consecutive years (*TMAX_END5*) (i.e., that is, *TMAX_END_J* where *J* = 3, 4, 5). These three treatment variables are our main independent variables of interest. While above-average temperatures may occur and have different impacts at different points in the farming season, this study examines above-average maximum temperatures at the end of the season. The reason is that this is the most recent weather event prior to the farmers' decision to use cover crop (or not).

Hence, it is reasonable to expect that these end-of-season heat events are the most likely to have an impact on cover crop utilization before the start of the next growing season. Note that the average maximum temperature is taken over the sample period to avoid increasing numbers of consecutive years with above-average maximum temperatures. Furthermore, these weather events are expected to be plausibly exogenous so that the effects estimated by the causal random forests will likely be consistent and unbiased.

In addition to the cover crop, weather, and treatment variables described above, the Seifert et al. (2019) data set utilized also includes two measures of soil quality – (i) a measure of water holding capacity (*WHC*) at the root zone, and (ii) a soil productivity measure called the National Commodity Crop Productivity Index (*NCCPI*) version 2 from the USDA NRCS' Soil Survey Geographic Database (SSURGO). We also use these two soil measures as observable covariates in the causal forest analysis. The summary statistics for the cover crop adoption variable, the three treatment variables, and the two soil quality measures are presented in Table 2.

2.2 Empirical Approach: Causal Random Forests

To explore the potential impacts of persistent end-of-season heat events on cover crop utilization (CC_{it}), we are conceptually interested in estimating an empirical specification of the form:

$$CC_{it} = f(TMAX\_END\_J_{it}, \mathbf{X}_{it}, \mu_i, \delta_t) + \varepsilon_{it} \qquad (1)$$

where CC_{it} is a binary cover crop use variable for pixel *i* in year *t*, $TMAX\_END\_J_{it}$ is the binary "treatment" variable representing persistent end-of-season heat events over *J* consecutive years (i.e, J = 3, 4, 5),[1] $\mathbf{X}_{it}$ is a vector of observable covariates (e.g., which includes the weather variables (see Table 1), the soybean yields (see Table 2), and the soil quality measures (see Table

[1] With J = 3, 4, 5, equation (1) is estimated three times – for three consecutive years of above average heat events (J = 3), for four consecutive years of above average heat events (J = 4), and for five consecutive years of above average heat events (J = 5).

2)), μ_i is an unobservable time-invariant pixel-level effect, δ_t is an unobservable year effect, and ε_{it} is an idiosyncratic error term. The "treatment effect" of persistent above average heat events (i.e., $TMAX\_END\_J_{it}$ in equation 1) is then estimated using a recently developed machine learning algorithm called the causal random forest approach (Athey et al., 2019).[2]

The causal random forest approach can estimate non-linear and non-parametric heterogeneous treatment effects (as in the general form in equation 1) when the treatment propensity is dependent on observable variables (like $\mathbf{X}_{it}$), but treatment is exogenous with respect to unobserved variables. The non-linear and non-parametric nature of the approach contributes to its flexibility in estimation and reduces the risk of misspecification since one does not need to assume a specific parametric fuctional forms (that may or may not allow for non-linearity). The causal random forest procedure also allows for calculating the variance of the non-linear, non-parametric heterogeneous treatment effect, which then permits statistical inference on the estimated treatment effects (i.e., which is an improvement over traditional nonparametric methods that usually do not allow for statistical inference). Moreover, the causal random forest method makes it possible to identify the most important observable variables (in the vector $\mathbf{X}_{it}$) that most likely drives the heterogeneity in the estimated treatment effects.

Causal random forests are a recent adaptation of the classical random forest machine learning algorithm, which generates consensus predictions from many individual classification or

[2] The conceptual specification in (1) can also be expressed in the potential outcomes framework (see Imbens and Rubin, 2015), where the treatment effect (more precisely called the conditional average treatment effect (CATE)) is defined as $\tau(x_i) = \mathrm{E}\left[Y_i^1 - Y_i^0 \,|\, X_i = x\right] = \mu_1(x) - \mu_0(x)$, where $\tau(x_i)$ is the difference in potential outcomes between the treatment (1) and control (0) (i.e., the treatment effect) conditional on the observable covariates *x*. This is then equivalent to the difference in the conditional means ($\mu_1(x) - \mu_0(x)$), assuming that the four underlying assumptions of this framework holds (e.g., conditional independence, stable unit value assumption, overlap assumption, and exogeneity of covariates).

regression trees (Breiman, 2001). Similar to the classical random forest approach in predictive machine learning (Breiman, 2001), causal forests attempt to find "neighbourhoods" in the covariate space, also known as recursive partitioning (i.e., partitioning the data for treatment estimation). While a random forest is built from decision trees, a causal forest is built from causal trees, where the causal trees learn a low-dimensional representation of treatment effect heterogeneity. Importantly, the splitting criterion used in causal forests is optimized for finding splits associated with treatment effect heterogeneity. Assuming that the treatment effect is constant over a neighbourhood $N(x)$, then it is possible to estimate the average treatment effect in each neighborhood partition. The goal is to find tree "leaves" (i.e., the neighborhood defined by the partitions), where the treatment effect is constant but is different from other "leaves."

Therefore, a causal random "forest" is simply the average of a large number of causal trees, where the trees differ due to subsampling (Athey and Imbens, 2019). To create a causal forest from causal trees, it is necessary to estimate a weighting function and use the resulting weights to solve a local generalized method of moments (GMM) model to estimate the average treatment effect (conditional on the observable variables). To deal with overfitting, causal forests use an "honesty" condition. A tree is considered "honest", if for each training sample (i), it only uses the response (Y_i) to estimate the within-leaf treatment effect or to decide where to place the split, but not both (Jacob, 2021). Sample-splitting is used to create honest trees, where half the data is used to estimate the tree structure (i.e., a splitting subsample), and then the other half is used to estimate the treatment effect in each leaf (i.e., an estimating subsample). The prediction of treatment effects is the difference in the average outcomes between the treated and the control observations of the estimating subsample in terminal leaves. Intuitively, causal forests are a kind of matching approach where one aims to split the sample such that each leaf can be interpreted as a random experiment (conditional on controls) from which treatment effects can be calculated per leaf.

Using honest trees also allows for asymptotic normality in the estimator used to estimate the variance of the estimates, which then allows for reliable confidence intervals of the parameters estimated (Wager and Athey, 2018). This is important, because to obtain an accurate estimate, the bias should asymptotically disappear, such that the confidence intervals are minimized. Since the bias vanishes asymptotically, the causal forest estimates are consistent and asymptotically Gaussian, which means that together with the estimator for the asymptotic variance (honest trees), valid confidence intervals are ensured.

More broadly (and for a more general intuition of the approach), causal random forests estimate treatment effects by comparing outcomes for each treatment sample against available control samples which are weighted based on their similarity to the the treatment sample. Thus, in our study, the causal random forest procedure allow us to use each pixel's closest neighbor in covariate space to generate a counterfactual cover crop outcome estimate under the "control" scenario without recent persistent warming. Furthermore, causal forests guard against confoundedness, including by unobserved variables, by using a "doubly robust" ("honest") treatment estimation method that combines treatment propensity weighting (i.e., in our case, how likely is the pixel to experience persistent warming) and a procedure based on a model of the expected outcome (i.e., in our case, whether or not to adopt cover crops). This approach minimizes sensitivity to misspecification (Athey et al., 2019; Scharfstein et al., 1999). Causal random forests also generate valid confidence intervals (as alluded to above), and are robust to large number of covariates, nonlinear interactions, and overfitting without requiring an explicit parametric model specification (Athey et al., 2019; Athey and Imbens, 2016; Belgiu and Dragu, 2016; Wager and Athey, 2018). Recent studies have shown that causal random forests are better able to detect and quantify heterogeneous treatment effects as compared to traditional econometric methods (Baiardi and Naghi, 2020; Farbmacher et al., 2019; Strittmatter, 2019).

3. Results and Discussion

The estimated average treatment effects (ATEs) from the causal random forests are presented in Table 3. Specifically, we show the results for experiencing three, four, and five consecutive years of above-average maximum temperatures at the end of the season. Note that each causal random forest in our runs contains 400 individual trees. In addition, note that the methodology used to estimate the average treatment effects is overlap-weighted. As indicated in Li et al. (2018), when treatment propensity is characterized by poor overlap (i.e., propensities that are either very close to zero or very close to one), the estimated average treatment effects tend to be unstable. Hence, the overlap-weighted estimation procedure is necessary to avoid instability caused by treatment propensities close to the endpoints.

In general, the estimated average treatment effects in Table 3 suggest that farmers with fields (or pixels) experiencing persistent above-average maximum temperatures at harvest time would have higher likelihood of adopting cover crops. The estimated treatment effects (conditional on observable covariates) are statistically significant at the 1% level of significance. Three consecutive years of above-average maximum temperatures at harvest increases the probability of cover crop use by 0.8 percentage points. In addition, more years of persistent warming experiences at harvest time consequently increases the likelihood of adopting cover crops. For example, four consecutive years of above-average maximum temperatures at the end of the season increases the probability of cover crop use by 1.3 percentage points, and five consecutive years of persistently high temperatures at harvest increases the probability of cover crop use even more (i.e., by 4 percentage points). This pattern of results indicate that soybean growers in Indiana likely respond to how frequent they experience heat events, and the more frequent they experience above-average temperatures increases the probability that farmers will use cover crops in their production system. We posit that this behavioral response may be due to farmers' knowing that cover crops' can help modulate soil temperatures in the presence of warming (see Scholberg et al., 2010 and Dabney et al., 2001). It is also possible that since warming temperatures is often associated with soil erosion, the behavioral response to use cover crops may also be related to anticipated soil erosion problems. That is, the persistent above average temperatures may be a signal for impending soil erosion problems, and so farmers react by using cover crops to help address (or preempt) potential future soil erosion issues (Prado Wildner et al., 2004; Chen et al., 2023).

To better contextualize the magnitude of the estimated treatment effects, we conduct a simple back-of-the-envelope calculation using cover crop adoption figures based on data from the 2017 Census of Agriculture. Zulauf and Brown (2019) estimated that the cover crop adoption rate in Indiana (for all crops) is around 7% based on the cover crop acreage reported in the 2017 Census of Agriculture. Based on our estimated treatment effects, a three-year heat event would increase cover crop adoption in Indiana from 7% to 7.8%. This is about an 11.42% increase in cover crop adoption rate due to three consecutive years of above average maximum temperatures at harvest (i.e., 0.8 divided by 7). Moreover, a four-year heat event at harvest would result in a 18.57% increase in cover crop adoption, and a five-year heat event at harvest would result in a 57% increase in cover crop adoption. These estimate magnitudes indicate that the impact of persistent above average temperatures on cover crop usage rates is economically meaningful.

Our causal random forest analysis also indicates that there is substantial heterogeneity in the estimated treatment effect of persistent above average temperatures at harvest. In Figures 1, 2, and 3, we present the unweighted histograms of the estimated treatment effects based on the causal random forests for the three, four, and five year persistent heat treatments. It is apparent from these figures that the effect of experiencing multi-year above average temperatures at harvest varies across observations in the sample. Moreover, the figures also suggest that the variation tend to be larger (i.e., less narrow bell shaped curves) for the four and five year treatment effect as compared to the three year treatment effect. This supports the notion that a persistent warming event over five-years gives a “strong” signal for use of cover crops (e.g., more farmers wanting to adopt to address soil problems) such that the upper tail of the distribution “spreads” to the right and increases heterogeneity of impact.

One advantage of the causal random forest approach is that it allows for calculation of a simple measure of “variable importance” to see which observable variables are most closely associated with the heterogeneity of the treatment effects. The variable importance measure we use is a simple weighted sum of how many times a particular variable was used as the splitting variable in each tree weighted by the depth at which the split occurred. Splits towards the top of the tree are presumed to be more important for the heterogeneity of the treatment effect compared to splits further down the tree. Splits are only considered down to a depth of 10 since the individual trees are rather large, so calculating this measure for each full tree would be rather time-consuming, and splits further down are unlikely to be very important. Figure 4 shows the results of this measure. Higher numbers for variable importance indicate that the variable is deemed as contributing more toward the heterogeneity in the treatment effect. The variable importance across each weather event from the causal random forests are broadly similar with a few exceptions. Surprisingly, usage of cover crops in the previous season and crop yield in the preceding season do not appear to be very important for the heterogeneity in the treatment effects. Variables associated with drought conditions such as precipitation before planting, vapor pressure early in the season, and the vapor pressure deficit at the end of the season are more strongly associated with the heterogeneity in the treatment effects. The minimum temperate early in the season also appears to be an important contributor, perhaps because it suggests to farmers that there is a more permanent change in climate conditions rather than a temporary warm weather event.

4. Conclusions

Although there have been numerous previous studies that explored factors that may affect cover crop adoption, overall understanding of the nature and scope of the relationship between climate risks and cover crop adoption remains limited. In this paper, we investigate the impact of persistent adverse warming events on subsequent cover crop utilization. In particular, this study examines how farmers' adoption of cover crops responds after experiencing consecutive years of above-average maximum temperatures at the end of the farming season. A causal random forest approach applied to a unique pixel-level satellite-based data is used to determine whether persistent warming events influence farmer adoption of cover cropping. Results from the causal random forest analysis suggest that farmers do respond to experiencing persistent warming events, and that this response is likely heterogeneous across observations in the sample. In addition, we find that more years of experiencing warming events also lead to a larger effect on cover crop adoption probabilities. Hence, it is likely that farmers will respond to warming climates by adapting climate-smart practices like cover crops. We posit that cover crops' ability to manage soil temperatures and to address soil erosion in the presence of warming contributes to this observed cover crop response of farmers to persistent warming temperatures.

Results from our study point to a couple of policy implications. Since farmers will likely increase cover crop use in response to persistent warming temperatures, this insight implies that it may be important to target educational programming and conservation payment programs to locations that have experienced warming events over the last several years. Based on our results, farmers in these “warming” areas are already more likely to utilize cover crops. Thus, educational programs about the resilience benefits of cover crops may help encourage more farmers in the area to use cover crops. Conservation payment programs in these areas may also push farmers “on the fence” about cover crops to ultimately use this practice in the presence of warming events. Extension programming efforts would also likely be more fruitful in areas already experiencing warming temperatures that can lead to soil health problems that cover crops can help address.

While our research represents a step forward in understanding cover crop adoption and climate risks, it is important to acknowledge study limitations and issues that deserve future attention. First, even though our satellite-based data set covers a large scale pixel-level adoption rate of cover crops in Indiana, we still do not consider other major agricultural production regions in the United States. Second, we only focus on farmers’ responses after experiencing extreme heat events. The analysis will be more comprehensive if one can conduct further studies to show how farmers’ decisions would change after experiencing various climate events. Third, our study has shown that consecutive heat events would increase the adoption of cover crops. We have not explored whether farmers’ use of cover crops would respond to sudden weather shocks. We leave all these suggested research directions for future work.

References

Arbuckle, J. G., and Roesch-McNally, G. (2015). Cover crop adoption in Iowa: The role of perceived practice characteristics. *Journal of Soil and Water Conservation*, 70(6):418–429.

Athey, S. and G.W. Imbens. (2019). Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*. 11:685-725

Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences of the United States of America*. 113(27): 7353–7360.

Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized random forests. *Annals of Statistics*, 47(2):1148–1178.

Baiardi, A., & Naghi, A. (2020). The value added of machine learning to causal inference: Evidence from revisited studies. Tinbergen Institute Discussion Paper 2021-001/V.

Belgiu, M., & Drăgu, L. (2016). Random forest in remote sensing: A review of applications and future directions. *ISPRS Journal of Photogrammetry and Remote Sensing*. 114: 24–31.

Bergtold, J. S., Ramsey, S., Maddy, L., and Williams, J. R. (2019). A review of economic considerations for cover crops as a conservation practice. *Renewable Agriculture and Food Systems* 34(1):62–76.

Bjerge, B., & Trifkovic, N. (2018). Extreme weather and demand for index insurance in rural India. *European Review of Agricultural Economics*, 45(3):397–431.

Breiman, L. (2001). Random forests. *Machine Learning*. 45(1): 5–32.

Burke, M., and Emerick, K. (2016). Adaptation to climate change: Evidence from U.S. agriculture. *American Economic Journal: Economic Policy*, 8(3):106–40.

Cai, J., & Song, C. (2017). Do disaster experience and knowledge affect insurance take-up decisions?. *Journal of Development Economics*, 124:83–94.

Carleton, T. A., and Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, 353(6304): aad9837.

Castellano, M. J., Helmers, M.J., Sawyer, J.E., Barker, D.W., and Christianson L. (2012). Nitrogen, Carbon, and Phosphorus Balances in Iowa Cropping Systems: Sustaining the Soil Resource. *Proceedings of the 24th Integrated Crop Management Conference*, 145–56. Ames: Iowa State University Digital Repository.

Che, Y., Feng, H., & Hennessy, D. A. (2020). Recency effects and participation at the extensive and intensive margins in the U.S. Federal Crop Insurance Program. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 45(1):52–85.

Chen, S., Chen, X., and Xu, J. (2016). Impacts of climate change on agriculture: Evidence from China. *Journal of Environmental Economics and Manage*ment, 76:105–124.

Chen, B., Gramig, B. M., and Yun, S. D. (2021). Conservation tillage mitigates drought induced soybean yield losses in the us corn belt. *Q Open*, 1(1):qoab007.

Chen, L., R.M. Rejesus, S. Aglasan, S. Hagen, and W. Salas. (2022). "The Impact of Cover Crops on Soil Erosion in the US Midwest." *Journal of Environmental Management*. 324: 1-15.

Connor, L., Rejesus, R. M., and Yasar, M. (2021). Crop insurance participation and cover crop use: Evidence from Indiana county-level data. *Applied Economic Perspectives and Policy*, 44(4), 2181–2208.

Dabney, S. M., Delgado, J. A., & Reeves, D. W. (2001). Using winter cover crops to improve soil and water quality. *Communications in Soil Science and Plant Analysis*, 32(7-8):1221–1250.

D'Agostino, A. L. and Schlenker, W. (2016). Recent weather fluctuations and agricultural yields: implications for climate change. *Agricultural economics*, 47(S1):159–171.

Deines, J. M., Guan, K., Lopez, B., Zhou, Q., White, C. S., Wang, S., & Lobell, D. B. (2022). Recent cover crop adoption is associated with small maize and soybean yield losses in the United States. *Global Change Biology*, 00:1–14.

Dong, F. (2022). Cover crops, drought, yield, and risk: an analysis of U.S. soybean production. NBER Chapters.

Dunn, M., Ulrich-Schad, J. D., Prokopy, L. S., Myers, R. L., Watts, C. R., and Scanlon, K. (2016). Perceptions and use of cover crops among early adopters: Findings from a national survey. *Journal of Soil and Water Conservation*, 71(1):29–40.

Farbmacher, H., Kogel, H., & Spindler, M. (2019). Heterogeneous effects of poverty on cognition. http://www.farbm acher.de/worki ng_papers/Farbmacher_etal_2019.pdf

Gardezi, M., and Arbuckle, J. G. (2019). The influence of objective and perceived adaptive capacity on Midwestern farmers' use of cover crops. *Weather, Climate, and Society*, 11(3):665–679.

Hanrahan, B. R., Tank, J. L., Christopher, S. F., Mahl, U. H., Trentman, M. T., and Royer, T. V. (2018). Winter cover crops reduce nitrate loss in an agricultural watershed in the central U.S. *Agriculture, Ecosystems & Environment*, 265:513–523.

Houser, M., & Stuart, D. (2020). An accelerating treadmill and an overlooked contradiction in industrial agriculture: Climate change and nitrogen fertilizer. *Journal of Agrarian Change*, 20(2):215–237.

Imbens, G.W. and D.B. Rubin. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, Cambridge MA: 1- 495.

IPCC, "Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change," Technical Report, IPCC, Geneva, Switzerland 2014. [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)].

Jacob, D. (2021). CATE meets ML - Conditional Average Treatment Effect and Machine Learning. (March 30, 2021). Available at SSRN: https://ssrn.com/abstract=3816558 or http://dx.doi.org/10.2139/ssrn.3816558

Kathage, J., Smit, B., Janssens, B., Haagsma, W., & Adrados, J. L. (2022). How much is policy driving the adoption of cover crops? Evidence from four E.U. regions. *Land Use Policy*, 116:106016.

Kaye, J. P., and Quemada, M. (2017). Using cover crops to mitigate and adapt to climate change. A review. *Agronomy for Sustainable Development*, 37(1):1–17.

Kladivko, E. J., Kaspar, T. C., Jaynes, D. B., Malone, R. W., Singer, J., Morin, X. K., & Searchinger, T. (2014). Cover crops in the upper midwestern United States: Potential adoption and reduction of nitrate leaching in the Mississippi River Basin. *Journal of Soil and Water Conservation*, 69(4):279–291.

Kousky, C. (2017). Disasters as learning experiences or disasters as policy opportunities? Examining flood insurance purchases after hurricanes. *Risk Analysis*, 37(3):517–530.

Lal, R. (2015). Soil carbon sequestration and aggregation by cover cropping. *Journal of Soil and Water Conservation*, 70(6): 329–339.

Laloy, E., and Bielders, C. L. (2010). Effect of intercropping period management on runoff and erosion in a maize cropping system. *Journal of Environmental Qua*lity, 39(3):1001–1008.

Lee, S., and McCann, L. (2019). Adoption of cover crops by U.S. soybean producers. *Journal of Agricultural and Applied Economics*, 51(4):527–544.

Li, F., Morgan, K. L., and Zaslavsky, A. M. (2018). Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, 113(521):390–400.

Montgomery, D. R. (2017). Growing a revolution: bringing our soil back to life. WW Norton & Company, Inc, New York.

Nearing, M. A., Pruski, F. F., and O'neal, M. R. (2004). Expected climate change impacts on soil erosion rates: a review. *Journal of Soil and Water Conservation*, 59(1), 43–50.

Olson, K. R., Al-Kaisi, M., Lal, R., & Morton, L. W. (2017). Soil ecosystem services and intensified cropping systems. *Journal of Soil and Water Conservation*, 72(3):64A–69A.

O'Neal, M. R., Nearing, M. A., Vining, R. C., Southworth, J., & Pfeifer, R. A. (2005). Climate change impacts on soil erosion in Midwest United States with changes in crop management. *Catena*, 61(2–3), 165–184.

Park, B., Rejesus, R. M., Aglasan, S., Che, Y., Hagen, S. C., and Salas, W. (2022). Payments from agricultural conservation programs and cover crop adoption. *Applied Economic Perspectives and Policy*, 1-24.

Poeplau, C., and Don, A. (2015). Carbon sequestration in agricultural soils via cultivation of cover crops–A meta-analysis. *Agriculture, Ecosystems & Environment*, 200:33–41.

Prado Wildner, L. D., Freitas, V., & McGuire, M. (2004). Use of green manure/cover crops and conservation tillage in Santa Catarina, Brazil. In Green Manure/Cover Crop Systems of Smallholder Farmers (1-36). Springer, Dordrecht.

Roesch-McNally, G. E., Arbuckle, J. G., and Tyndall, J. C. (2018a). Barriers to implementing climate resilient agricultural strategies: The case of crop diversification in the US Corn Belt. *Global Environmental Change*, 48:206–215.

Roesch-McNally, G. E., Basche, A. D., Arbuckle, J. G., Tyndall, J. C., Miguez, F. E., Bowman, T., and Clay, R. (2018b). The trouble with cover crops: Farmers' experiences with overcoming barriers to adoption. *Renewable Agriculture and Food Systems*, 33(4):322–333.

Scharfstein, D. O., Rotnitzky, A., Robins, J. M. (1999). Adjusting for nonignorable drop-out using semiparametric nonresponse nodels. *Journal of the American Statistical Association*. 94(448): 1096–1120.

Schipanski, M. E., Barbercheck, M., Douglas, M. R., Finney, D. M., Haider, K., Kaye, J.P., Kemanian, A.R., Mortensen, D.A., Ryan, M.R., Tooker, J. and White, C. (2014). A framework for evaluating ecosystem services provided by cover crops in agroecosystems. *Agricultural Systems*, 125:12–22.

Schnepf, M., and Cox, C. (2006). Environmental benefits of conservation on cropland: the status of our knowledge. Soil and Water Conservation Society, Ankeny, IA.

Scholberg, J., Dogliotti, S., Leoni, C., Cherr, C. M., Zotarelli, L., and Rossing, W. A. (2010). Cover crops for sustainable agrosystems in the Americas. *Genetic Engineering, Biofertilisation, Soil Quality and Organic Farming*, 23–58.

Seifert, C. A., Azzari, G., and Lobell, D. B. (2019). Corrigendum: Satellite detection of cover crops and their effects on yield in the Midwestern United States. *Environ. Res. Lett.*, 14(3):1–14.

Shackelford, G. E., Kelsey, R., and Dicks, L. V. (2019). Effects of cover crops on multiple ecosystem services: Ten meta-analyses of data from arable farmland in California and the Mediterranean. *Land Use Policy* 88:104204.

Snapp, S.S., Swinton, S.M., Labarta, R., Mutch, D., Black, J.R., Leep, R., Nyiraneza, J. and O'neil, K. (2005). Evaluating cover crops for benefits, costs and performance within cropping system niches. *Agronomy Journal*, 97(1):322–332.

Strittmatter, A. (2019). What is the value added by using causal machine learning methods in a welfare experiment evaluation? *ArXiv*. http://arxiv.org/abs/1812.06533v2

Tack, J., Barkley, A., and Nalley, L. L. (2015). Effect of warming temperatures on U.S. wheat yields. *Proceedings of the National Academy of Sciences*, 112(22):6931–6936.

Thornton P.E. R. Shrestha, Y. Wei, P.E. Thornton, S. Kao, and B.E. Wilson. 2014 Daymet: Daily Surface Weather Data on a 1 km Grid for North America, Version 2 (Oak Ridge, TN: Oak Ridge National Laboratory)

Tonitto, C., David, M. B., and Drinkwater, L. E. (2006). Replacing bare fallows with cover crops in fertilizer-intensive cropping systems: A meta-analysis of crop yield and N dynamics. *Agriculture, Ecosystems & Environment* 112(1):58–72.

Urban, D. W., Sheffi ld, J., and Lobell, D. B. (2015). The impacts of future climate and carbon dioxide changes on the average and variability of us maize yields under two emission scenarios. *Environmental Research Letters*, 10(4):045003.

U.S. Department of Agriculture (USDA). 2022. Partnerships for Climate - Smart Commodities. Accessed December 07, 2022. https://www.usda.gov/climate-solutions/climate-smart-commodities

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Yoder, L., Houser, M., Bruce, A., Sullivan, A., and Farmer, J. (2021). Are climate risks encouraging cover crop adoption among farmers in the southern Wabash River Basin?. *Land Use Policy*, 102:105268.

Zhang, P., Zhang, J., and Chen, M. (2017). Economic impacts of climate change on agriculture: The importance of additional climatic variables other than temperature and precipitation. *Journal of Environmental Economics and Management*, 83:8–31.

Zulauf, C., and Brown B. (2019). Cover crops, 2017 U.S. Census of Agriculture. *Farmdoc Daily* 9(135):1–3.

Table 1: Summary Statistics for Weather Variables

Variable	Nbr. of Obs.	Min	Median	Mean	Max	Std. Dev.
PRCP_PREPLANT	3,313,207	0.600	4.300	4.891	15.567	2.569
PCRP_EARLY	3,313,207	0.780	4.879	4.850	9.264	1.282
PCRP_MID	3,313,207	0.803	4.525	4.470	8.254	1.103
PCRP_END	3,313,207	0.804	4.072	4.045	7.163	0.842
TMAX_PREPLANT (˚C)	3,313,207	12.67	17.95	18.06	23.62	2.050
TMAX_EARLY (˚C)	3,313,207	19.74	23.00	23.14	27.79	1.423
TMAX_MID (˚C)	3,313,207	21.17	24.47	24.64	29.86	1.488
TMAX_END (˚C)	3,313,207	22.08	25.29	25.44	30.23	1.451
TMIN_PREPLANT (˚C)	3,313,207	0.883	5.33	5.40	10.37	1.613
TMIN_EARLY (˚C)	3,313,207	7.97	11.20	11.23	15.74	1.284
TMIN_MID (˚C)	3,313,207	9.36	12.84	12.86	17.30	1.319
TMIN_END (˚C)	3,313,207	10.43	13.53	13.56	18.03	1.252
VP_PREPLANT	3,313,207	561.3	933.3	937.0	1,290.7	114.33
VP_EARLY	3,313,207	960	1,421	1,424	1,882	123.35
VP_MID	3,313,207	1,077	1,575	1,579	2,076	137.34
VP_END	3,313,207	1,127	1,635	1,639	2,130	136.14
VPD_PREPLANT	3,313,207	-91.25	16.31	31.38	270.41	55.637
VPD_EARLY	3,313,207	-192.36	-90.12	-77.29	457.18	71.20
VPD_MID	3,313,207	-251.76	-113.27	-99.66	694.45	84.68
VPD_END	3,313,207	-244.84	-106.07	-87.98	823.19	86.45

Table 2: Summary Statistics for the Cover Crop, Treatment, and Soil Quality Variables

Variable	Nbr. of Obs.	Min	Median	Mean	Max	Std. Dev.
Cover Crop Usage (*CC*)	3,313,207	0	0	0.264	1	0.441
TMAX_END3	2,503,014	0	0	0.168	1	0.374
TMAX_END4	2,104,253	0	0	0.086	1	0.280
TMAX_END5	1,596,890	0	0	0.051	1	0.220
Yield (bu/ac)	3,313,207	10	64.16	61.74	100	12.244
SSURGO *WHC*	3,313,207	0	201	200.4	742	73.560
SSURGO *NCCPI*	3,313,207	0	0.614	0.592	0.988	0.204
Year: 2008	385,978					
Year: 2009	424,215					
Year: 2010	398,761					
Year: 2011	507,363					
Year: 2012	219,451					
Year: 2013	449,740					
Year: 2014	392,450					
Year: 2015	535,249					

Table 3: Overlap-Weighted Average Treatment Effects (ATEs)

	Three Consec.	Four Consec.	Five Consec.
Overlap-Weighted ATE	0.008***	0.013***	0.040***
	(0.0010)	(0.0017)	(0.0028)
Other Weather Controls	Yes	Yes	Yes
Year Effects	Yes	Yes	Yes
Number of Observations	2,503,014	2,104,253	1,595,890

*** indicates $p < 0.01$, ** indicates $p < 0.05$, and * indicates $p < 0.1$

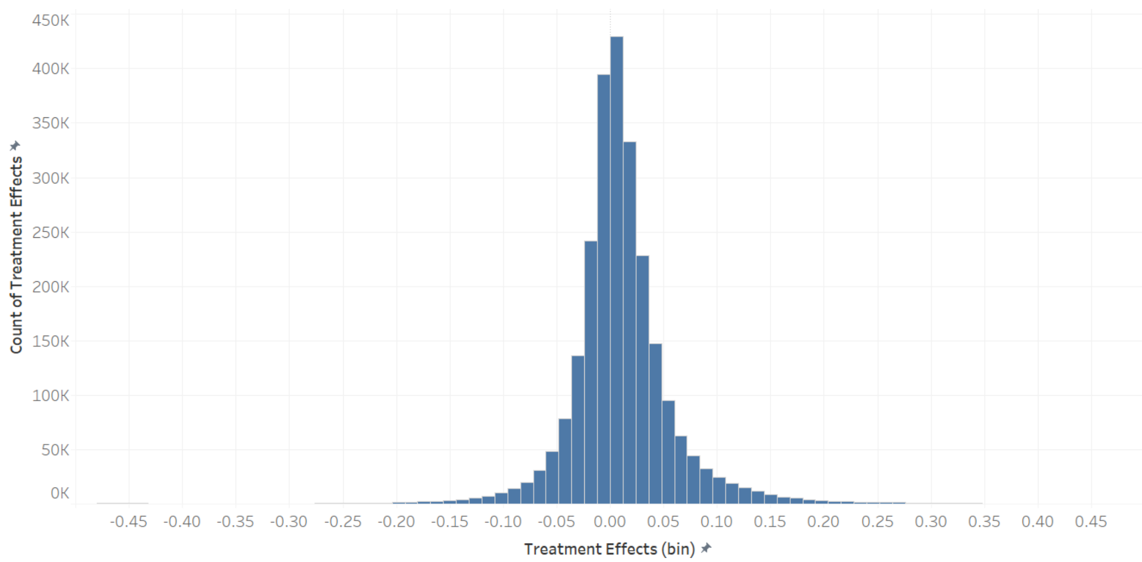

Figure 1: Histogram of the unweighted distribution of treatments effects for 3 consecutive above-average TMAX_END

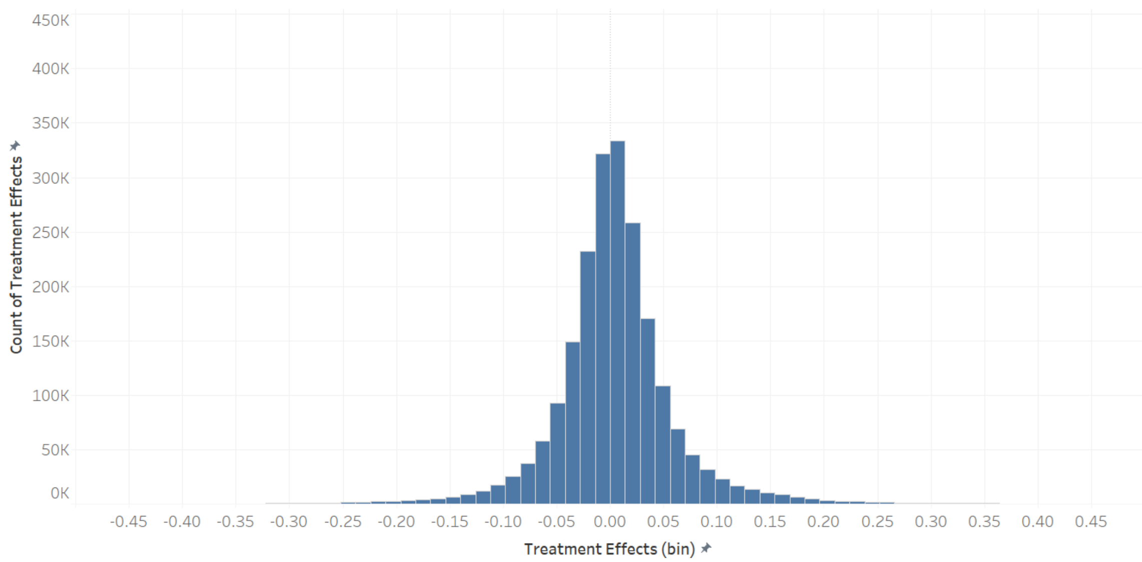

Figure 2: Histogram of the unweighted distribution of treatments effects for 4 consecutive above-average *TMAX_END*

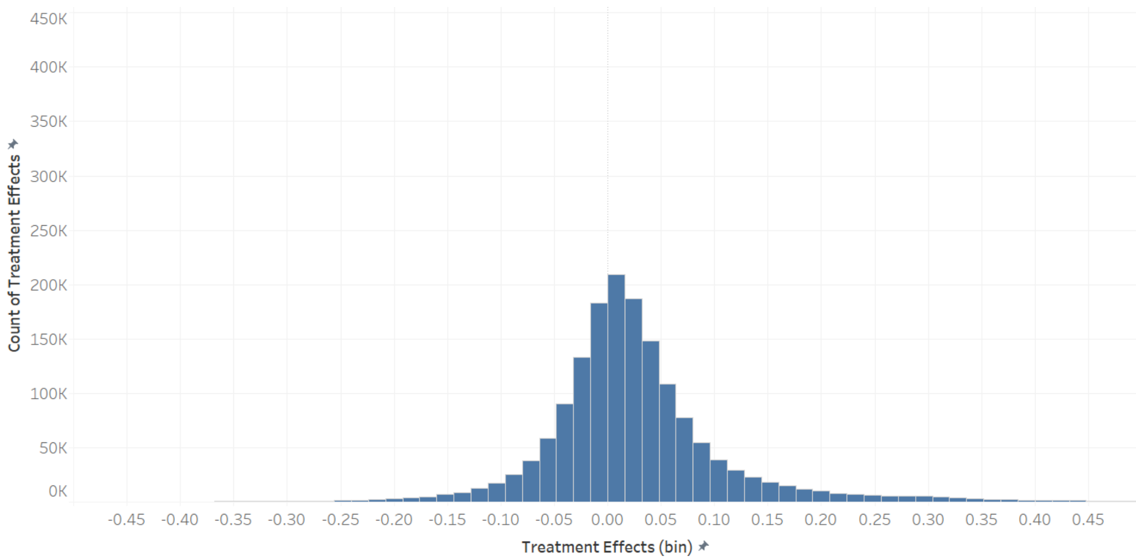

Figure 3: Histogram of the unweighted distribution of treatments effects for 5 consecutive above-average *TMAX_END*

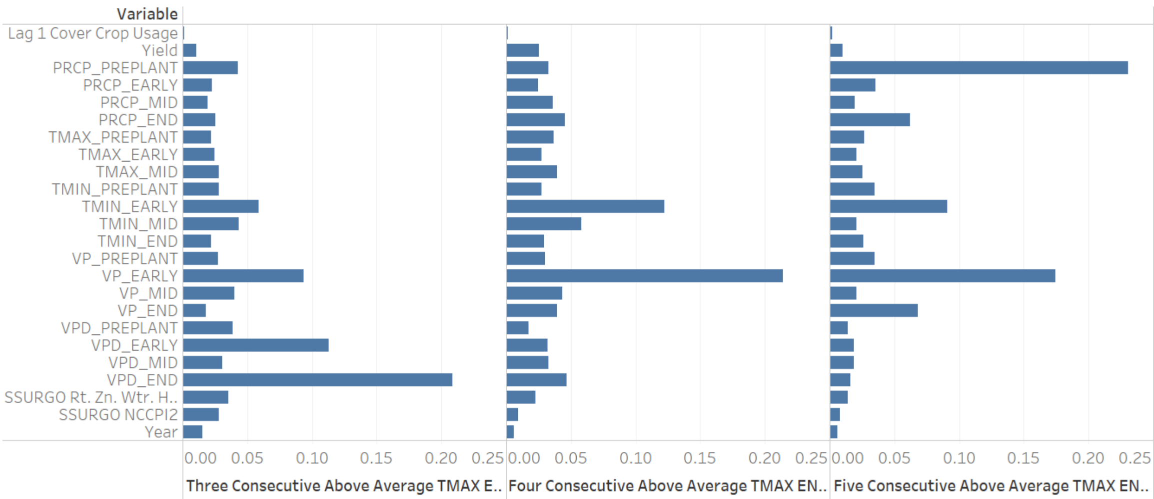

Figure 4: Variable Importance from the Causal Random Forests for each Weather Event Treatment Effect