

Heterogeneous Yield Impacts from Adoption of Genetically Engineered Corn and the Importance of Controlling for Weather

Jayson L. Lusk, Jesse Tack, and Nathan P. Hendricks

July 23, 2018

Abstract: Concern about declining growth in crop yields has renewed debates about the ability of biotechnology to promote food security. While numerous experimental and farm-level studies have found that adoption of genetically engineered crops has been associated with yield gains, aggregate and cross-country comparisons often seem to show little effect, raising questions about the size and generalizability of the effect. This paper attempts to resolve this conundrum using a panel of United States county-level corn yields from 1980 to 2015 in conjunction with data on adoption of genetically engineered crops, weather, and soil characteristics. Our panel data contain just over 28,000 observations spanning roughly 800 counties. We show that changing weather patterns confound simple analyses of trend yield, and only after controlling for weather do we find that genetically engineered crops have increased yields above trend. There is marked heterogeneity in the effect of adoption of genetically engineered crops across location partially explained by differential soil characteristics which may be related to insect pressure. While adoption of genetically engineered crops has the potential to mitigate downside risks from weeds and insects, we find no effects of adoption on yield variability nor do we find that adoption of presently available genetically engineered crops has led to increased resilience to heat or water stress. On average, across all counties, we find adoption of GE corn was associated with a 17 percent increase in corn yield.

Keywords: biotechnology, climate change, genetic engineering, GE

JEL Codes: Q16, O47, C23

Although agriculture has historically experienced one of the highest rates of productivity growth in the U.S. economy (Jorgenson, Gollop, Fraumeni, 1987), there is evidence that agricultural productivity growth is beginning to slow (Alston, Andersen, and Pardey, 2015; Alston, Beddow, and Pardey, 2009; Ray et al., 2012). The decline in productivity growth has coincided with concerns about food price spikes, social instability, food insecurity, population growth, drought, and climate change (Bellemare, 2015; Ray et al., 2013; Roberts and Schlenker, 2013; Schlenker and Roberts, 2009; Tack, Barkley, and Nalley, 2015). This confluence of problems has prompted interest in determining whether certain technologies can promote gains in crop yields, and none has been more controversial than biotechnology.

Many previous studies have investigated whether adoption of genetically engineered (GE) crops has increased yield (e.g., see reviews in Fernandez-Cornejo et al., 2014; Klümper and Qaim, 2014; National Academies of Sciences, Engineering, and Medicine (NASEM), 2016), and the consensus from the micro-level data and experimental studies is that adoption of GE crops, particularly insect-resistant Bt varieties targeting the corn borer, have generally been associated with higher yield. However, ample skepticism remains, with high profile popular publications purporting that GE crops have failed to live up to their promise of yield increases (e.g., Foley, 2014; Gurian-Sherman, 2009; Hakim, 2016).

A variety of factors might explain the divergence in views about the yield effects of GE crops, but one of the main issues is that adoption of GE crops does not appear to have had much effect on trend yields when investigating national-level yield data (Duke, 2015) nor do yield trends appear much different in developed countries that have and have not adopted GE varieties (Heinemann et al., 2014). As the NASEM (2016) put it (p. 66), “the nation-wide data on maize, cotton, or soybean in the United States do not show a significant signature of genetic-engineering

technology on the rate of yield increase.” This raises the question of whether the yield-increasing effects of GE crops observed in particular locations and experiments generalize more broadly, and if so, does the impact on crop yields vary spatially?

In this paper, we show that simple analyses of yield trends mask important weather-related factors that influence the estimated effect of GE crop adoption on yield. Our analysis couples county-level data on corn yields from 1980 to 2015 and state-level adoption of GE traits with data on weather variation and soil characteristics. Using state-level adoption data does not induce measurement error bias because state-level aggregate adoption is necessarily uncorrelated with the deviation of a particular county’s adoption from the state-level aggregate. Using state-level adoption data does induce serial correlation of the error term which we address with two-way clustering.

A number of important findings emerge from our analysis. First, changes in weather and climatic conditions confound yield effects associated with GE adoption. Without controlling for weather, adoption of GE crops appears to have little impact on corn yields; however, once temperature and precipitation controls are added, GE adoption has significant effects on corn yields. Second, the adoption of GE corn has had differential effects on crop yields in different locations even among corn-belt states. However, we find that ad hoc political boundaries (i.e., states) do not provide a credible representation of differential GE effects. Rather, alternative measures based on soil characteristics provide a broad representation of differential effects and are consistent with the data. In particular, we find that the GE effect is much larger for non-sandy soils with a larger water holding capacity. Overall, we find that GE adoption has increased yields by approximately 18 bushels per acre on average, but this effect varies spatially across counties ranging from roughly 5 to 25 bushels per acre. Finally, we do not find evidence that adoption of

GE corn led to lower yield variability nor do we find that current GE traits mitigate the effects of heat or water stress.

The adoption of GE crops does not necessarily imply that farmers perceived yield benefits because there are several other benefits associated with the adoption of GE crops—primarily through reducing the cost of production. The non-yield benefits have come in the form of labor savings, reduced insecticide use, and improved weed and pest control which has facilitated the ability to adopt low- and no-till production methods, alter crop rotations, and utilize higher planting densities (Chavas et al., 2014; Fernandez-Cornejo et al., 2014; Klümper and Qaim, 2014; Perry, Moschini, and Hennessy, 2016; Perry et al., 2016). Revealed preferences of US farmers indicate producer benefits over and above the substantially higher price of GE corn relative to conventional (Shi, Chavas, and Stiegert, 2010). The rapid adoption of GE corn by farmers also provides evidence of these benefits. GE corn was first grown commercially in the US in 1996. In just four years, a quarter of the corn acres were planted with a GE trait and in less than ten years, adoption had spread to more than half the US corn acres. In 2016, 92 percent of US corn acres were planted with GE corn, with 81 percent of the total GE corn acreage being planted with “stacked” varieties that are both insect resistant and herbicide tolerant. It is also important to recognize that GE crops can increase production through expansions of area planted to the crop (i.e., the extensive margin) because greater yields and lower costs of production provide incentives to expand crop production (Barrows, Sexton, and Zilberman, 2014).

Nonetheless, discussion of yield impacts of GE crops remains at the forefront of public discussions about whether and to what extent biotechnology can contribute to food security and help mitigate the effects of climate change. In response to the finding that GE adoption does not appear to alter national-level yield trends, NAESM (2016) recommended research (p. 66),

“should be conducted that isolates effects of the diverse environmental and genetic factors that contribute to yield.” Our objective here is to help fill this gap in the literature.

The next section reviews some of the research on the yield effects of GE crops, and we delineate our contribution to the literature. The third and fourth sections discuss the data and methods, followed by the presentation of results. The last section concludes.

Background

GE crops currently on the market do not increase yield *per se*. However, they can reduce the gap between actual and potential yield by reducing the adverse effects of weeds and insects (NAESM, 2016). It is also possible that crops with GE traits can reduce yields if introduced into less productive varieties not ideally suited to a particular growing region (Shi, Chavas, and Lauer, 2013).




Figure 1. Trend in National U.S. Corn Yield and Adoption of GE Corn from 1980 to 2016. Blue circles represent observed yields, solid black line is linear yield trend, and dashed red line is percent of corn acres planted with GE corn; data are from USDA National Agricultural Statistics Service.

Figure 1 shows the national trend in US corn yield and the adoption of GE corn from 1980 to 2016. The figure suggests, in the words of Duke (2014, pg., 653), “yields have continued to increase at the same rate as before introduction [of GE crops].” Leibman et al. (2014) similarly investigated aggregate yields and found after adoption of GE corn a small (0.5 bushes/acre) trend increase; however, no tests of statistical significance were performed. These sorts of aggregate comparisons make no attempt to control for potentially confounding factors such as weather which could have coincidentally been worse in the 1980s before the adoption of GE corn. Controlling for weather in a national-level trend analysis is difficult due to the nonlinear impacts of weather on yield and highly spatially heterogeneous weather conditions within the country. This motivates the use of disaggregate data to test the impact of GE adoption on yields.

These aggregate investigations can be contrasted with the large literature from agronomic experimental studies that attempt to hold constant many factors such as location and germplasm. Nolan and Santos (2012) summarize the results of more than 30 such studies mainly published between 2000 and 2003. None of the reviewed studies report a statistically significant negative effect associated with Bt GE corn, and nearly all reported positive yield effects associated with Bt GE trait, with yield gains as high as 19 percent. In their analysis, Nolan and Santos (2012) combined datasets from multiple experiments conducted by ten different state agricultural extension services from 1997 to 2009. They found, after controlling for weather, agronomic inputs, management, and soil characteristics that planting of Bt GE corn led to yield gains of around 14 bushels/acre, although when the only GE trait present was herbicide tolerance, yield was unaffected or slightly negative. Despite finding that yield was affected by location, weather,

and soil characteristics, the authors did not investigate whether these factors interacted with the GE effect (i.e., whether GE yield gains were higher or lower in different locations, in different weather patterns, or in different soils).

As shown by Shi, Chavas, and Lauer (2013), however, there is likely ample heterogeneity in the effects of GE adoption on mean yield and on yield variance. Using experiment data from agricultural experiment stations in Wisconsin from 1990 to 2010, Shi, Chavas, and Lauer (2013) found GE traits had variable effects on corn yields, depending on the type of GE trait introduced and how long the trait had been used in production, with mean yields significantly increasing relative to conventional non-GE corn for some traits (namely Bt targeted at the European corn borer) but not others (namely, herbicide-tolerant only GE corn and GE corn with Bt targeted only at corn rootworm). Additional analysis of the same data by Shi et al. (2013) suggests some of the yield gains attributable to GE hybrids were a result of improvements in non-GE germplasm and the wider availability of higher quality germplasm. However, regardless of the GE trait analyzed, the authors found a consistent effect on yield variance, with GE crops reducing the variance of corn yields. The authors conclude that GE crops have helped farmers reduce their risk exposure.

As was the case in Nolan and Santos (2012), Shi, Chavas, and Lauer (2013) did not investigate whether the yield effects of GE traits were affected by location, management practices, soil type, etc. However, there are reasons to believe the potential yield effects of GE adoption are not uniform across location or time. Currently available GE traits rely on Bt to provide protection against the European corn borer and/or corn rootworm and/or are tolerant to certain herbicides (primarily glyphosate). While there are fewer agronomic reasons to suggest herbicide tolerance would convey significant yield benefits, insect resistance can plausibly lower the gap between potential and realized yield. As discussed in Nolan and Santos (2012),

conventionally applied insecticides only provide 60 percent to 80 percent protection against corn borer and rootworm, whereas Bt provides near 100 percent protection. As such, the effect of Bt GE corn relative to conventional corn depends on pest pressure. It has long been known that corn borer and corn rootworm pressures are affected by soil characteristics and by weather (e.g., Beck and Apple, 1961; Huber, Neiswander, and Salter, 1928; Turpin and Peters, 1971; MacDonald and Ellis, 1990), and prior research has hinted at the fact that yield effects of Bt corn might depend on soil characteristics via their effects on insect populations (Ma, Meloche, and We, 2009). Pest pressure is also likely to vary spatially according to the density of corn production, which depends on soil and climatic conditions.

In a paper most similar to the present inquiry, Xu et al. (2013) used aggregate, non-experimental data and found that GE adoption led to a 19.4 bushel/acre increase in corn yields in the central corn belt (Illinois, Indiana, and Iowa). What explains the contrast between the apparent lack of impact of GE adoption in aggregate trend yields shown in figure 1 and the results from Xu et al. (2013)? There are a variety of possibilities. For example, Xu et al. (2013) look at county (rather than national) yields, and they control for confounding factors related to weather and fertilizer use. However, it is unclear from Xu et al. (2013) what the impacts are of ignoring these factors. Moreover, these authors only considered limited geographic heterogeneity (they only explored central corn belt to non-central corn belt) and they did not consider other factors like soil characteristics or how weather and soil characteristics may influence GE-adoption effects on yield. In addition, the authors did not consider the effects of GE adoption on yield variability.

Another confounding factor that exists when exploring national yield trends is the fact that the number of acres planted to corn has increased significantly over the same period of time

that GE traits have been adopted. For example, in ten years from 1980 to 1989 prior to adoption of GE corn, 75.7 million acres of corn were planted on average each year in the US. By contrast, in the most recent ten-year period from 2007 to 2016, during a period of near full adoption of GE traits, on average 91.2 million acres of corn were planted each year in the US, a 20.5 percent increase. Some of the acreage expansion is a result of GE adoption as GE traits have increased the viability of continuous corn (planting corn after corn rather than rotating with soybeans) (Chavas et al., 2014), a practice which has historically been associated with significant yield drag (Gentry et al., 2013). Ethanol policies, among other factors, also led to dramatic increase in corn prices over the period of GE corn adoption, which both increased the prevalence of continuous corn (Hendricks et al., 2014) and led to expansion of corn onto acres which would previously have been considered marginal lands. Combined, these factors suggest national corn yields would have been higher in recent years had it not been for the expansion of corn acreage.

Data

We utilize a large panel of roughly 28,000 yield observations spanning 819 counties from 1980-2015. We chose 1980 as the starting point for the time series as this gives us a roughly equal number of years pre- and post-GE adoption, which started in 1996. Roughly 13,000 (45 percent) of the yield observations correspond to the pre-GE period. These data were collected via USDA National Agricultural Statistics Service (NASS) Quick Stats and correspond to total production divided by harvested acres in each county. As in Xu et al (2013), we omit any county where (i) more than 10 percent of harvested cropland is irrigated or (ii) yield data was reported for less than two-thirds of the pre-GE years or two-thirds of the post-GE years. Figure S1 shows that

there exists extensive cross-sectional and temporal variation of yields. Note that all tables and figures with a leading “S” are contained in the accompanying supplementary material.

The limiting factor for the cross-sectional (spatial) representation of the data is the availability of GE adoption data. We utilize the same NASS data as Xu et al (2013), which reports GE adoption at the state-year level for 13 states: Illinois, Indiana, Iowa, Kansas, Michigan, Minnesota, Missouri, Nebraska, North Dakota, Ohio, South Dakota, Texas, and Wisconsin. These data were first recorded in 2000, for all but North Dakota and Texas which started recording in 2005, several years after adoption had already started to occur in some areas. We interpolate missing data using predictions from a generalized linear model with a binomial family and a logit link function. A pooled model with state fixed effects provides similar predictions as separate models for each state. We use the latter here. Our interpolation procedure follows the seminal work of Griliches (1957) who modeled the diffusion of hybrid corn seed as logistic growth. Figure S2 provides the observed adoption data in addition to the predictions from both models used to interpolate missing values. Table S1 provides summary statistics for both the observed and observed-plus-interpolated GE adoption rate variables. Figure S3 provides a spatial map of the in-sample counties studied in the analysis.

While our analysis focuses on identifying the effects of GE adoption on corn yields, it is likely that one must control for several sources of confounding factors in practice. For example, if the period of adoption post 1996 coincided with an abnormal run of good or bad weather conditions, then failure to control for weather could bias the estimate of the GE effect. Recent evidence suggests that this can be an important consideration for crop yield analyses (Tack, Barkley, and Nalley, 2015). We use the same weather data as in Schlenker and Roberts (2009) updated to 2015 to control for the influence of weather on corn yields. Daily outcomes on

minimum and maximum temperatures at the county level are interpolated within each day using a sinusoidal approximation, and are then used to construct three degree-day variables: between 0 and 10°C, 10 and 29°C, and above 29°C. Along with cumulative precipitation, these variables are aggregated across March-August. Figure S4 shows that there is extensive variation both cross-sectionally and over time for these variables.

Soil characteristic data are from the gSSURGO (Gridded Soil Survey Geographic) database created by NRCS (Natural Resources Conservation Service) and were also used in Hendricks (2016). Soils are aggregated to the county level using only the area in the county classified as cropland according to the National Land Cover Database. One measure of soil quality that we consider is water holding capacity (measured in mm), which is the total volume of plant-available water that the soil can store within the root zone and is calculated as a weighted average across the county. Figure S7 provides a spatial map of the water holding capacity across counties, which has a wide range from near zero to just over 300 mm and a sample average value of 216 mm. Another measure we consider is a grouping of soil types based on the soil texture of the county. This is calculated as the dominant soil texture classification within the county, and includes nine different soil types: clay, clay-loam, loam, loamy-sand, sand, sandy-loam, silt-loam, silty-clay, and silty-clay-loam. Figure S8 provides a spatial map of these soil types by county.

Empirical Model

We assume that the effect of GE adoption on corn yields is identified using the regression model

$$y_{ist} = \alpha + \delta A_{it}^* + f(\mathbf{x}_{ist}, \boldsymbol{\gamma}) + \varepsilon_{ist}$$

where y_{ist} are corn yields (bu/acre) in county i , state s , and year t . The variable A_{it}^* is the unobserved GE adoption rate at the county-year level and is measured as the fraction of acreage planted to GE varieties. The parameter of interest is δ , which measures the effect of GE adoption on corn yields. We discuss the implications of only observing adoption rates at the state-year level in the next section. We also include a vector of control variables \mathbf{x}_{ist} that include county-level fixed effects, state-level trend variables, and weather variables that are measured at the county-year level.

Our identification comes from two different sources of variation. First, the effect is identified from differences in yield over time as GE adoption has increased, controlling for state-specific yield trends and weather. Second, the effect is identified from differences in yields from the county-specific average - adjusted for the state-specific trend - between counties that had different levels of GE adoption.

One potential concern is that GE adoption could be endogenous—counties with larger increases in yield adopted GE more rapidly. Our primary source of variation in GE adoption—before and after the introduction of GE technology—is exogenous because the introduction of GE varieties was driven by supply of the technology rather than farmer demand. A more likely source of endogeneity is that counties with greater (or smaller) increases in yields had more rapid early adoption. To test if this endogeneity affects our results, we present a robustness check where all periods during diffusion are omitted from the sample so that we only exploit yield differences before and after the introduction of the GE technology. An alternative strategy is to include year fixed effects so that our source of identification is the change in yields between counties with different changes in adoption rates of GEs. We do not include year fixed effects in our preferred specification because it precludes the use of exogenous pre-adoption data as a

counterfactual for post-adoption data. Including year fixed effects also exploits a narrow source of variation in GE adoption.

State-Level Adoption Rates

One of the main concerns in identifying δ is that data on GE adoption are only available at the state level, thus the variable that we observe is A_{st} . This can be cast as a non-classical measurement error problem where the true unobserved measure A_{it}^* is related to our observed measure A_{st} by

$$A_{it}^* = A_{st} + v_{it}$$

where v_{it} denotes the difference between the county-specific adoption and the state-level aggregate adoption. Substituting this expression into the regression model and rearranging gives

$$y_{ist} = \alpha + \delta A_{st} + f(\mathbf{x}_{ist}, \gamma) + u_{ist}$$

$$u_{ist} = \delta v_{it} + \varepsilon_{ist}.$$

The error term u_{ist} is a composite random variable. Note that this source of measurement error is non-classical in the sense that it does not induce bias in our estimate of δ because the measurement error is uncorrelated with the observed state-level adoption. The state-level aggregate adoption in each year is uncorrelated—by definition—with the deviation of a particular county's adoption from the state-level aggregate.

However, the measurement error does induce serial correlation of the error terms on a subsample of the data. Some counties are likely to lead or lag the state-level adoption rate in all periods of technology diffusion, resulting in serial correlation of the errors. Note that measurement error is only a concern during the period of adoption when $A_{st} \in (0,1)$ since A_{it}^* and

A_{st} are necessarily both equal to zero prior to adoption and both equal to one at full adoption. Thus, the measurement error induces serial correlation in the errors only when both periods are in the adoption phase of the data.

We use the two-way clustering approach of Cameron, Gelbach, and Miller (2012) to account for multiple sources of correlation in the errors. The first dimension of clustering is accomplished by using a county-by-adoption-phase grouping scheme such that each county is split into two groups, one when GE adoption equals zero and another when adoption is positive. This clustering accounts for potential serial correlation resulting from the measurement error of the state-level adoption variable. The second dimension of clustering is by year in order to account for the presence of spatial correlation in the errors (ε_{ist}) driven by the spatial similarity of residual weather shocks not accounted for in the model across counties within each year. We interpret this approach as being robust to heteroscedasticity, spatial correlation of the error terms across counties, and serial correlation of the errors within each county both within and outside of the adoption period.

Importance of Controlling for Weather

It is worth noting that the need to control for weather is important if the exposure differs between the pre- and post-GE subsamples in the data (or if weather was relatively good or bad in counties that adopted more rapidly). In theory, if one were to observe a large enough frequency of weather outcomes in both periods such that average weather exposures were similar, then one would not need to control for its influence. In practice, weather may bias coefficients in samples where the number of time periods is not large. We investigate this possibility by comparing precipitation and extreme heat exposure (degree days above 29°C) in both periods, i.e. pre and

post 1996. For each county, we calculate the percentage difference in the average precipitation and extreme heat variables across periods and report these values in Figures S5 and S6. Precipitation was roughly 4 percent higher in the post-GE period on average across counties. However, this masks a large amount of heterogeneity as county-level differences ranged from -15 to 21 percent. Similarly, the occurrence of extreme heat exhibited a large amount of heterogeneity as differences as large as -63 and 22 percent spanned an average value of -20 percent. This suggests that controlling for weather is important, and must be done at a local level. It would not likely be properly accounted for using spatially aggregated measures of weather shocks at the regional or national level, nor using crude measures such as year fixed effects.

Heterogeneous Yield Response

We also consider models of the form

$$y_{ist} = \alpha + \delta_g A_{st} + f(\mathbf{x}_{ist}, \boldsymbol{\gamma}) + u_{ist}$$

where we are now allowing the parameter of interest δ to vary across different subsections of the data represented by groupings g . We interact the GE adoption variable with the weather variables that are a subset of the variables in \mathbf{x}_{ist} to investigate whether GE varieties are more or less susceptible to certain weather outcomes. We also consider several models of cross-sectional heterogeneity, each based on a different assignment of the county to a group. The first utilizes a grouping based on the state each county is in, while the other two assign each county to a group based on measures of soil quality. The first measure of soil quality defines groups based on the percentiles of the observed water holding capacity variable: 0-10th, 10th-25th, 25th-50th, 50th-75th, 75th-90th, and 90th-100th. The second defines groups based on the dominant soil texture in each county.

Results

We report estimates for three classes of models in this section. The first section provides estimates for a class of models that assume a homogeneous GE effect across counties. The second class of models maintains this homogeneity assumption but reports estimates for models of the higher order moments of the corn yield distribution. The final class of models relaxes the homogeneity assumption and allows the GE effect to vary across weather outcomes and county groupings. All models are estimated using OLS with standard errors clustered using the two-way approach discussed above.

Homogeneous GE Effect

Table 1 reports parameter estimates for five models that sequentially include additional control variables. In the absence of any controls, the estimated GE effect is 43 bushels per acre. However, it is clear that this estimate is confounded by the absence of a trend variable which, when included, changes the estimate to -8 bushels per acre. This sensitivity is expected as the increase in GE adoption has coincided with many other production innovations that have increased productivity over time. Failure to account for this source of variation in the data confounds the estimate. The next model adds county fixed effects to the model, and the estimate becomes positive but is not statistically significantly different from zero at conventional levels. In addition, the estimate of the time trend parameter is also sensitive to the inclusion of these fixed effects as it decreases from 2 to 1.5 bushels per acre per year. It is clear that controlling for county-specific time-invariant yield

Table 1. Regression results: impacts of GE adoption on corn yield (bushels per acre)

	Model 1	Model 2	Model 3	Model 4	Model 5
--	----------------	----------------	----------------	----------------	----------------

Variables					
GE adoption rate	43.36***	-7.648	6.547	18.15***	18.26***
	[7.116]	[11.17]	[12.15]	[6.546]	[6.748]
Time trend		2.000***	1.506***	1.008***	0.943***
		[0.393]	[0.413]	[0.206]	[0.208]
Precipitation (mm)				1.652***	1.596***
				[0.363]	[0.363]
Precipitation squared (mm ²)				-	-
				0.0146***	0.0142***
				[0.00328]	[0.00327]
Degree Days 0-10°C				0.0216	0.0181
				[0.0262]	[0.0263]
Degree Days 10-29°C				0.0159	0.0176
				[0.0155]	[0.0143]
Degree Days above 29°C				-0.590***	-0.591***
				[0.0868]	[0.0808]
County fixed effects	N	N	Y	Y	Y
State-specific trends	N	N	N	N	Y
R-squared	0.171	0.234	0.649	0.781	0.792
Out of Sample RMSE (% Reduction)	--	3.89	32.8	46.9	48.8
Observations	28,628	28,628	28,628	28,628	28,628
Counties	819	819	819	819	819
Years	36	36	36	36	36

Notes: The reported coefficient estimate for the time trend variable under model 5 is the simple average of the state-specific estimates. The out-of-sample prediction comparison reports the percentage reduction in the root-mean-squared prediction error (RMSE) for each model compared to the baseline model 1 that does not include any control variables. Each model is estimated 1,000 times, where each iteration randomly selects 80 percent of the sample observations. Relative performance is measured according to the accuracy of each model's prediction for the omitted 20 percent of the data. Two-way clustered standard errors by year and county-adoption are reported in square brackets. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels.

drivers such as soil quality is an important consideration for estimating productivity gains. The next model includes the precipitation and weather variables, and again we see a sensitivity of the estimates as the GE effect increases and becomes statistically significant at the 1 percent level.

Note that adding the weather variables nearly triples the estimate of the GE effect and reduces the standard error by half. Adding the weather variables also reduces the trend coefficient to 1 bushel per acre per year, implying that the trend increase in yield was much smaller before the introduction of GE varieties after we account for weather conditions.

The final model allows the time trend parameters to vary across states and we find that the GE estimate has stabilized across this additional generalization. Both the in-sample fit and the out-of-sample prediction accuracy of the models steadily increases as we include additional control variables in the model. Further supporting this finding that control variables matter, we reject the null of equality for the state-specific trend estimates (p-value 0.000) and reject the null that the weather variable estimates are jointly zero (p-value 0.000). Thus, under an assumption of a homogenous GE effect, we find that the introduction of GE corn has increased yields by approximately 18 bushels per acre. This represents a roughly 17 percent increase in yields relative to the five-year average yield prior to the introduction of GE traits.

We also estimated alternative specifications that exploit different sources of variability in GE adoption. First, we removed all observations during the diffusion process—defined as adoption greater than 0 and less than 0.85—so that we only exploit variation in GE adoption before and after the introduction of the technology. This specification exploits variation in GE adoption that is more plausibly exogenous, but we reduce the sample size to 17,424. The effect of GE adoption on yield is larger at 28 bushels per acre (p-value 0.013) but the standard error increases to 11.24.

Our second specification adds year fixed effects to the main specification. The effect of GE adoption on yield decreases to 8 bushels per acre and lacks statistical significance (p-value 0.555) with a standard error of 13.92. As mentioned previously, the specification with year fixed effects exploits a small amount of cross-sectional variation in GE adoption which leads to an imprecise estimate.

We also consider specifications that only use a subset of years in the data to better understand our primary source of identification. When we omit years for which the national

adoption rate exceeded 50 percent (i.e. 2005-14), the effect of GE adoption is 31 bushels per acre (p-value 0.013). If we instead omit years when national adoption is between 0 and 50 percent (i.e. 1996-2004), the effect of GE adoption is 24 bushels per acre (p-value 0.028). However, when we omit years when national adoption was zero (i.e. 1980-1995) the effect of GE adoption drops substantially to 2 bushels per acre (p-value 0.79; standard error 9.1). These results indicate that our estimate of the effect of GE adoption in our preferred specification is identified primarily from the inclusion of pre-adoption data that permits a credible estimation of yield trends prior to adoption. Overall, these alternative specifications provide evidence that endogenous selection into early- vs late-adoption cohorts does not substantially bias the results.

We also consider robustness to a model that includes solar radiation as an additional control as recent literature has suggested that this can confound trend estimation for US corn (Tollenaar et al. 2017), but find little concern here as the GE adoption estimate is 17.5 bushels per acre (p-value 0.022) under this alternative. Finally, it is possible that GE adoption coincided with large expansions and/or retractions of acreage to/from marginal land within each county. While we cannot observe this directly in these data, we can omit annual observations for each county that are more/less than one standard deviation away from county-specific sample averages of harvested acreage. This effectively produces a subset of data (25 percent fewer observations) that experiences smaller year-to-year acreage variations, and we find a similar estimate of 14.8 bushels per acre (p-value 0.012).

Higher Order Moment Effects of GE

We next consider whether GE adoption has influenced the higher order moments of the corn yield distribution using a moments-model approach (Antle 1983, 2010; Just and Pope, 1978).

Specifically, we estimate both the variance and skewness of the yield distribution using the squared and cubed residuals from the preferred model from the previous section (Table 1, model 5). These transformed residuals are then regressed on the same set of covariates as the preferred model. Under expectation of the dependent variable, these models provide linkages between the GE adoption and control variables on the variance and skewness of the yield distribution.

The parameter estimates for these models are reported in Table S2. We find no evidence that GE adoption has affected the variance or skewness of the yield distribution as the estimates are not statistically significant at standard levels. However, taken in conjunction with the results revealing increases in mean yields, our results suggest that GE adoption has led to a reduction in yield risk as it has increased yields without a proportionate increase in the standard deviation; that is, the coefficient of variation has decreased. Importantly, the estimate for the time trend variable implies an increase in yield variance over time, suggesting that it is an important control variable for studies to consider when estimating yield risk implications of GE. Although not reported here, when the time trends are dropped from the model, the estimate of the GE effect on yield variance becomes positive and significant at the 10 percent level (p-value 0.096).

Heterogeneous GE Effect

Results for joint hypothesis tests for the heterogeneous models are reported in Table S3, where the p-values correspond to a null hypothesis of a homogenous GE effect. The first three models explore interactions between the weather variables and GE adoption. We find no evidence of these interactions for precipitation alone, temperature alone, nor precipitation and temperature combined. Thus we conclude that, while ignoring weather severely biases the estimated effect of GE corn adoption, the performance of GE varieties is not likely dependent on particular weather

outcomes occurring. One reason for this result is that the initial GE traits focused on insect resistance and herbicide tolerance, rather than developing traits to improve drought or heat tolerance. In the future, GE traits may focus more on drought and heat so that our result may not continue to hold (Marshall, 2014).

The next set of heterogeneity models that we consider assign each county to a particular group. We first use state boundaries to define the grouping and estimate the heterogeneous effects by interacting dummy variables for each state with the GE adoption variable. We fail to reject the null of a homogeneous effect at standard significance levels (p-value = 0.1117), however Figure S9 provides a spatial map of these estimates and suggests that there are potentially large differences in effects across regions. A simple average of the estimates is 19.1 and they range from 5.5 to 27.5 bushels per acre.

To further explore potential spatial heterogeneities, we assign each county to one of six groups based on the soil's water holding capacity. The groups correspond to different percentiles of the empirical distribution of observed values (1: 0-10th percentile, 2: 10-25th percentile, 3: 25-50th percentile, 4: 50-75th percentile, 5: 75-90th percentile, and 6: 90-100th percentile). We interact dummy variables for each group with the GE adoption variable, and we find evidence that this pattern of spatial heterogeneity is supported by the data as a joint hypothesis test suggests rejecting the null of a homogeneous effect at standard significance levels (p-value = 0.0000). The parameter estimates are reported in Table S4 and Figure 2 provides a spatial map of the impacts by county. The county-level estimates have an average value of 18.4 and they range from 12.5 to 25.1 bushels per acre. It is clear from the map that there exists substantial within-state variation of the GE effect that the state-specific heterogeneity model is not capable of capturing. Figure S10 plots the range of county-level GE effects within each state along with the

average value within that state, and shows that there exists a broad range of over 10 bushels per acre within most states. This insight is consistent with the state-specific model's failure to reject the null of a homogeneous effect, as the spatial heterogeneities are not being driven by ad hoc political boundaries but rather the localized growing conditions.




Figure 2. Impacts of GE Corn Adoption (bushels per acre) by Counties' Soil Water Holding Capacities. Each county is assigned to one of six groups based on the soil's water holding capacity, and a separate GE impact is estimated for each group. Estimated impacts are then binned according to values in the figure legend.

The final model that we consider further supports this insight and suggests that soil texture is also an important dimension for understanding heterogeneous GE effects. We assign each county to one of nine groups based on the dominant soil texture: clay, clay-loam, loam, loamy-sand, sand, sandy-loam, silt-loam, silty-clay, silty-clay-loam. We interact dummy variables for each group with the GE adoption variable and we find evidence that this pattern of spatial heterogeneity is supported by the data as a joint hypothesis test suggests rejecting the null of a homogeneous effect at standard significance levels (p -value = 0.0001). The parameter estimates are reported in Table S4 and Figure 3 provides a spatial map of the impacts by county. The county-level estimates have an average value of 18.3 and they range from 3.9 to 24.0 bushels per acre. We again find that there exists substantial within-state variation of the GE effect that the state-specific heterogeneity model is not capable of capturing as shown in Figure S10.

The location of greater yield impacts from GE adoption corresponds generally with the location of greater utilization of Bt traits (see Hutchison et al. 2010). As mentioned earlier in the paper, research from experimental plots has generally found greater yield benefits from Bt traits than herbicide resistant traits. There are a couple reasons why soil characteristics might be driving greater adoption of Bt traits. First, corn borer and corn rootworm pressures are affected by soil characteristics (e.g., Beck and Apple, 1961; Huber, Neiswander, and Salter, 1928; Turpin and Peters, 1971; MacDonald and Ellis, 1990). Second, areas with better soil characteristics have a greater concentration of corn production resulting in greater pest pressure.




Figure 3. Impacts of GE Corn Adoption (bushels per acre) by Counties' Soil Types. Each county is assigned to one of nine groups based on the soil's texture, and a separate GE impact is estimated for each group. Estimated impacts are then binned according to values in the figure legend.

Conclusion

There is considerable interest, both within the academic community and among the broader public, in the effects of GE crop adoption. In particular, the effect of GE crop adoption on yield has been the subject of much debate, perhaps because of relationships between yield and

environmental outcomes via land use and because of the implications for food security. Numerous experiments have found that GE traits have tended to reduce the gap between actual and potential corn yields (Fernandez-Cornejo et al., 2014; Klümper and Qaim, 2014; and NASEM, 2016); however, experimental studies often generate significantly higher yields than farmers actually experience (Lobell et al., 2009). Moreover, aggregate yield trends in the US appear, at first blush, relatively stable before and after adoption of GE corn.

This paper sought to identify whether, in fact, for corn “the nation-wide data . . . in the United States do not show a significant signature of genetic-engineering technology on the rate of yield increase,” as was indicated by NASEM (2016). Using corn yield panel data corresponding to roughly 28,000 U.S. county-years before and after adoption of GE corn, a simple model only including a time trend confirms NASEM’s assertion, as the effect of GE adoption appears, if anything, to have had a negative effect on yields. However, subsequent analysis reveals this simple model is biased. After controlling for weather and soil characteristics, and assuming a homogeneous effect of adoption, we find that adoption of GE corn has led to an approximate 17 percent increase in corn yields. We also find significant heterogeneity in the yield-effect that is not related to state-boundaries but rather to soil characteristics. On average, adoption of GE corn has led to an 18.5 bushel per acre increase in yield, but the effects range from 12.5 to 25.1 bushels per acre depending on soil characteristics. We conjecture that the variation across soil types may be related to differences in insect pressure.

While we found important soil-GE adoption interactions, there were no significant interactions related to weather. The findings suggest that the current GE traits have not led to more resilience to heat or water stresses. Moreover, while we find that the variance in corn yield has increased over time, adoption of GE corn has not lowered the variance. Nonetheless, if as our

results show, adoption of GE corn increases yield without affecting variance, the coefficient of variation on yields has fallen as a result of GE corn adoption. This suggests GE corn is less risky as, for example, the actuarially fair price of insurance to indemnify a given yield falls as the coefficient of variation falls.

Our study has a number of limitations. As we discussed, the available adoption data only exists at the state level. We showed that this produces a type of measurement error problem that can lead to serially-correlated errors – an issue we address using the two-way clustering approach of Cameron, Gelbach, and Miller (2012). There are other issues that have likely affected national-level yields, such as the move toward more continuous corn and other factors that have led to the expansion of corn acres. The adoption of GE crops may have also led to an expansion of corn acres on more marginal land that decreased county average yield. We partially addressed this problem by limiting our analysis to only those counties that reported yield data for at least two-thirds of the years before and after GE adoption, but an altogether different sort of analysis that moves from our primal production function approach to a structural model that relates planting decisions to input and output prices would likely be required to fully address the issue. To the extent that adoption of GE crops did reduce county average yields through an expansion of growing area, our estimate reflects this effect and understates the impact of GE on field-level corn yields. It would also be of interest to conduct the sort of analysis performed here using data, for example, from the European Union, where there has been little to no adoption of GE corn. Such an approach would permit a truer difference-in-difference estimate of the effect of GE corn adoption.

A final important caveat to be noted is that the estimated effects of GE corn adoption depend critically on the available GE technologies. Genetic engineering is not a single “thing.”

In the case of our data, GE corn is one of four types: herbicide tolerant, Bt corn-borer tolerant, Bt root-worm tolerant, or stacked varieties that include combinations of the previous three types. Geneticists and plant scientists are continually working on new genetic modifications that could further reduce the gap between actual and potential yield or even increase potential yields. For example, Kromdijk et al. (2016) recently genetically engineered a tobacco plant to improve the efficiency of photosynthesis, which increased potential yields by about 20%. Other research has focused on genetic pathways to increasing nitrogen utilization (McAllister et al., 2012). Whether these additional GE crop technologies can live up to the “hype” of increasing crop yields remains to be seen. However, if and when these biotechnologies arrive, it will be important to closely scrutinize whether they substantively affect real-world farm yields, just as this study has attempted to do with the first generations of GE corn.

References

- Antle, J. 1983. "Testing the stochastic structure of production: A flexible moment-based approach." *Journal of Business and Economic Statistics* 1(3):192–201.
- Antle, J. 2010. "Asymmetry, partial moments, and production risk." *American Journal of Agricultural Economics* 92:1294–1309.
- Alston, J.J., M.A. Andersen, and P.G. Pardey. 2015. "The rise and fall of U.S. farm productivity growth, 1910-2007." Working Paper, Department of Applied Economics, University of Minnesota. Staff Paper P15-02. <http://ageconsearch.umn.edu/handle/200927>
- Alston, J.M., J.M. Beddow, and P.G. Pardey. 2009. "Agricultural research, productivity, and food prices in the long run." *Science* 325(5945):1209-1210.
- Barrows, G., S. Sexton, D. Zilberman. 2014. "The Impact of Agricultural Biotechnology on Supply and Land-Use." *Environment and Development Economics* 19(6):676–703.
- Beck, S.D. and J.W. Apple. 1961. "Effects of temperature and photoperiod on voltinism of geographical populations of the European corn borer, *Pyrausta nubilalis*." *Journal of Economic Entomology* 54(3):550-558.
- Bellemare, M.F. 2015. "Rising food prices, food price volatility, and social unrest." *American Journal of Agricultural Economics* 97(1):1-21.
- Cameron, A.C., J.B. Gelbach, and D.L. Miller. 2012. "Robust inference with multiway clustering." *Journal of Business & Economic Statistics* 29(2):238-249.
- Chavas, J.P., G. Shi, and J. Lauer. 2014. "The Effects of GM Technology on Maize Yield." *Crop Science* 54(4):1331-1335.
- Duke, S.O. 2015. "Perspectives on transgenic, herbicide-resistant crops in the USA almost 20 years after introduction." *Pest Management Science* 71:652–657.
- Fernandez-Cornejo, J., S. Wechsler, M. Livingston, and L. Mitchell. 2014. "Genetically Engineered Crops in the United States." U.S. Department of Agriculture, Economic Research Service. Economic Research Report Number 162, February. https://www.ers.usda.gov/webdocs/publications/err162/43668_err162.pdf
- Foley, J. "GEs, Silver Bullets and the Trap of Reductionist Thinking." *Ensa.com*. February 25, 2014. <http://ensia.com/voices/GEs-silver-bullets-and-the-trap-of-reductionist-thinking/>
- Gentry, L.F., M.L. Ruffo, and F.E. Below. 2013. "Identifying factors controlling the continuous corn yield penalty." *Agronomy Journal* 105(2): 295-303.
- Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (4): 501–22.
- Gurian-Sherman, D. 2009. "Failure to Yield: Evaluating the Performance of Genetically Engineered Crops." Union of Concerned Scientists. April. http://www.ucsusa.org/sites/default/files/legacy/assets/documents/food_and_agriculture/failure-to-yield.pdf
- Hakim, D. 2016. "Doubts About the Promised Bounty of Genetically Modified Crops." *New York Times*, October 29. http://www.nytimes.com/2016/10/30/business/GE-promise-falls-short.html?_r=0
- Heinemann, J.A., M. Massaro, D.S. Coray, S.Z. Agapito-Tenfen, and J.D. Wen. 2014. "Sustainability and innovation in staple crop production in the US Midwest." *International Journal of Agricultural Sustainability*. 12(1)71-88
- Hendricks, N.P. 2016. "The economic benefits from innovations to reduce heat and water stress in agriculture." Working paper. Available at: <https://nphendricks.com/>

- Hendricks, N.P., A. Smith, and D.A. Sumner. 2014. "Crop supply dynamics and the illusion of partial adjustment." *American Journal of Agricultural Economics* 96(5):1469-1491.
- Huber, L.L., C.R. Neiswander and R.M. Salter. 1928. "The European corn borer and its environment." Wooster, OH: Ohio Agricultural Experiment Station. Bulletin 429. December.
- Hutchison, W. D., E. C. Burkness, P. D. Mitchell, R. D. Moon, T. W. Leslie, S. J. Fleischer, M. Abrahamson, et al. 2010. "Areawide Suppression of European Corn." *Science* 330: 222–25.
- Jorgenson, D., F. Gollop, and B. Fraumeni. 1987. *Productivity and U.S. economic growth*. Harvard University Press.
- Harvard Univ. Press, Cambridge, MA.; Jorgenson, D., M. Ho, and K. Stiroh. 2005. *Productivity: Information technology and the American growth resurgence*. MIT Press, Cambridge, MA
- Just, R.E., and R.D. Pope. 1978. "Stochastic specification of production functions and economic implications." *Journal of Econometrics* 7(1):67-86.
- Klümper, W. and M. Qaim. 2014. "A meta-analysis of the impacts of genetically modified crops." *PLoS One* 9(11):p.e111629.
- Kromdijk, J., Głowacka, K., Leonelli, L., Gabilly, S.T., Iwai, M., Niyogi, K.K. and Long, S.P., 2016. "Improving photosynthesis and crop productivity by accelerating recovery from photoprotection." *Science* 354(6314):857-861.
- Leibman, M., J.J. Shryock, M.J. Clements, M.A. Hall, P.J. Loida, A.L. McClerren, Z.P. McKiness, J.R. Phillips, E.A. Rice, and S.B. Stark. 2014. "Comparative analysis of maize (*Zea mays*) crop performance: Natural variation, incremental improvements and economic impacts." *Plant Biotechnology Journal* 12(7):941–950.
- Lobell, D.B., K.G. Cassman, and C.B. Field. 2009. "Crop yield gaps: their importance, magnitudes, and causes." *Annual Review of Environment and Resources* 34(1):179-204.
- Ma, B.L., F. Meloche, and L. Wei. 2009. "Agronomic assessment of Bt trait and seed or soil-applied insecticides on the control of corn rootworm and yield." *Field Crops Research* 111(3):189-196.
- MacDonald, P.J. and C.R. Ellis. 1990. "Survival time of unfed, first-instar western corn rootworm (Coleoptera: Chrysomelidae) and the effects of soil type, moisture, and compaction on their mobility in soil." *Environmental Entomology* 19(3):666-671.
- Marshall, A. 2014. "Drought-tolerant varieties begin global march." *Nature Biotechnology* 32(4):308.
- McAllister, C.H., P.H. Beatty, and A.G. Good. 2012. "Engineering nitrogen use efficient crop plants: the current status." *Plant Biotechnology Journal* 10(9), pp.1011-1025.
- National Academies of Sciences, Engineering, and Medicine. 2016. *Genetically Engineered Crops: Experiences and Prospects*. Washington, DC: The National Academies Press. doi: 10.17226/23395.
- Nolan, E. and P. Santos. 2012. "The contribution of genetic modification to changes in corn yield in the United States." *American Journal of Agricultural Economics* 94(5):1171-1188.
- Perry, E.D., F. Ciliberto, D.A. Hennessy, and G. Moschini. 2016. "Genetically engineered crops and pesticide use in US maize and soybeans." *Science Advances* 2(8):e1600850.
- Perry, E. D., G. Moschini, and D.A. Hennessy. 2016. "Testing for complementarity: Glyphosate tolerant soybeans and conservation tillage." *American Journal of Agricultural Economics* 98(3):765-784.

- Ray, D.K., N. Ramankutty, N.D. Mueller, P.C. West, and J.A. Foley. 2012. "Recent patterns of crop yield growth and stagnation." *Nature communications* 3:1293.
- Ray, D.K., N.D. Mueller, P.C. West, and J.A. Foley. 2013. "Yield trends are insufficient to double global crop production by 2050." *PloS One* 8(6):e66428.
- Roberts, M.J., and W. Schlenker. 2013. "Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate." *American Economic Review* 103(6):2265-2295.
- Schlenker, W. and M.J. Roberts. 2009. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *Proceedings of the National Academy of Sciences* 106(37):15594-15598.
- Shi, G., J.P. Chavas, and K. Stiegert. 2010. "An analysis of the pricing of traits in the US corn seed market." *American Journal of Agricultural Economics* 92(5):1324-1338.
- Tack, J., A. Barkley, and L.L. Nalley. 2015. "Effect of warming temperatures on US wheat yields." *Proceedings of the National Academy of Sciences* 112(22):6931-6936.
- Tack, J., A. Barkley, and L.L. Nalley. 2015. "Estimating yield gaps with limited data: An application to United States wheat." *American Journal of Agricultural Economics* 97(5):1464-1477.
- Tollenaar, M., J. Fridgen, P. Tyagi, P.W. Stackhouse Jr, and S. Kumudini. 2017. "The contribution of solar brightening to the US maize yield trend." *Nature Climate Change* 7(4):275-278.
- Turpin, F.T. and D.C. Peters. 1971. "Survival of southern and western corn rootworm larvae in relation to soil texture." *Journal of Economic Entomology* 64(6):1448-1451.
- Xu, Z., D.A. Hennessy, K. Sardana, and G. Moschini. 2013. "The realized yield effect of genetically engineered crops: US maize and soybean." *Crop Science* 53(3):735-745.

Supplementary Material: Heterogeneous Yield Impacts from Adoption of Genetically Engineered Corn and the Importance of Controlling for Weather

This document reports supplementary tables S1-S4 and figures S1-S10 that are referenced in the manuscript.

Supplementary Tables

Table S1. Summary statistics for raw and interpolated state-level GE adoption rates, 1980-2015

Site		Mean	Std Dev	Min	Max	Obs
All States:	Raw	0.349	0.395	0.00	0.98	419
	Interpolated	0.322	0.377	0.00	0.98	481
Illinois:	Raw	0.322	0.381	0.00	0.93	33
	Interpolated	0.294	0.370	0.00	0.93	37
Indiana:	Raw	0.289	0.366	0.00	0.88	33
	Interpolated	0.262	0.354	0.00	0.88	37
Iowa:	Raw	0.368	0.399	0.00	0.95	33
	Interpolated	0.341	0.385	0.00	0.95	37
Kansas:	Raw	0.381	0.408	0.00	0.95	33
	Interpolated	0.353	0.394	0.00	0.95	37
Michigan:	Raw	0.312	0.371	0.00	0.93	33
	Interpolated	0.284	0.359	0.00	0.93	37
Minnesota:	Raw	0.387	0.409	0.00	0.93	33
	Interpolated	0.360	0.394	0.00	0.93	37
Missouri:	Raw	0.341	0.373	0.00	0.93	33
	Interpolated	0.315	0.360	0.00	0.93	37
Nebraska:	Raw	0.388	0.413	0.00	0.96	33
	Interpolated	0.360	0.399	0.00	0.96	37
North Dakota:	Raw	0.391	0.462	0.00	0.97	28
	Interpolated	0.357	0.414	0.00	0.97	37
Ohio:	Raw	0.252	0.338	0.00	0.86	33
	Interpolated	0.228	0.327	0.00	0.86	37
South Dakota:	Raw	0.437	0.446	0.00	0.98	33
	Interpolated	0.409	0.429	0.00	0.98	37
Texas:	Raw	0.360	0.425	0.00	0.91	28
	Interpolated	0.327	0.378	0.00	0.91	37
Wisconsin:	Raw	0.319	0.369	0.00	0.92	33
	Interpolated	0.293	0.357	0.00	0.92	37

Table S2. Regression results: impacts of GE adoption on variance and skewness of corn yields

Variables	Variance	Skewness
GE adoption rate	-95.18 [166.1]	1526.6 [10927.8]
Time trend	9.086** [4.588]	-134.0 [306.5]
Precipitation (mm)	-22.46*** [7.371]	-352.3 [503.8]
Precipitation squared (mm ²)	0.187*** [0.0498]	3.344 [4.773]
Degree Days 0-10°C	0.689 [0.512]	-22.56 [50.22]
Degree Days 10-29°C	-0.375 [0.354]	-8.051 [24.25]
Degree Days above 29°C	4.955** [1.940]	104.7 [183.8]
County fixed effects	Y	Y
State-specific trends	Y	Y
R-squared	0.134	0.034
Observations	28,628	28,628
Counties	819	819
Years	36	36

Notes: The reported coefficient estimates for the time trend variable in both models is the simple average of the state-specific estimates. Two-way clustered standard errors by year and county-adoption are reported in square brackets. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels.

Table S3. Joint hypothesis tests for the heterogeneous GE effect models

Null Hypothesis	P value
Weather interaction model	
<i>All weather/GE interactions are equal to zero</i>	0.3344
<i>All precipitation/GE interactions are equal to zero</i>	0.4538
<i>All temperature/GE interactions are equal to zero</i>	0.6672
State specific GE effect	
<i>All state specific GE effects are equal</i>	0.1117
GE effect varies by soil water holding capacity	
<i>All GE effects for each group are equal</i>	0.0000
GE effect varies by soil type	
<i>All GE effects for each group are equal</i>	0.0001

Notes: The reported p-values correspond to the joint hypothesis of a homogenous GE effect using two-way clustered standard errors by year and county-adoption.

Table S4. Regression results: heterogeneous impacts of GE adoption by soil type

Variables	Model 1	Model 2
(GE adoption rate) X (<10 th p water holding capacity)	12.50* [7.070]	
(GE adoption rate) X (10-25 th p water holding capacity)	18.68*** [6.971]	
(GE adoption rate) X (25-50 th p water holding capacity)	14.12** [6.805]	
(GE adoption rate) X (50-75 th p water holding capacity)	15.69** [6.941]	
(GE adoption rate) X (75-90 th p water holding capacity)	21.84*** [6.866]	
(GE adoption rate) X (>90 th p water holding capacity)	25.13*** [7.002]	
(GE adoption rate) X (clay soil)		24.01* [12.31]
(GE adoption rate) X (clay-loam soil)		19.46*** [7.385]
(GE adoption rate) X (loam soil)		22.35*** [6.861]
(GE adoption rate) X (loamy-sand soil)		7.537 [7.329]
(GE adoption rate) X (sand soil)		3.911 [7.955]
(GE adoption rate) X (sandy-loam soil)		12.56* [7.436]
(GE adoption rate) X (silt-loam soil)		16.41** [6.934]
(GE adoption rate) X (silty-clay soil)		14.79** [7.490]
(GE adoption rate) X (silty-clay-loam soil)		22.99*** [7.382]
Simple average of GE effects	17.99*** [6.735]	16.00** [6.963]
R-squared	0.7929	0.7927
Out of Sample RMSE (% Reduction)	-50.1	-50.0
Observations	28,628	28,628
Counties	819	819
Years	36	36

Notes: Both models include a full set of controls: weather variables, county fixed effects, and state-specific linear trends. The out-of-sample prediction comparison reports the percentage reduction in the root-mean-squared prediction error (RMSE) for each model compared to a baseline model that does not include any control variables and assumes a homogeneous GE effect. Each model is estimated 1,000 times, where each iteration randomly selects 80 percent of the sample observations. Relative performance is measured according to the accuracy of each model's prediction for the omitted 20 percent of the data. Two-way clustered standard errors by year and county-adoption are reported in square brackets. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels.

Supplementary Figures





Figure S1. Spatial and temporal variation of yields. We observe yields at the county-year level, and construct boxplots for each year. Each box is defined by the upper and lower quartile, with the median depicted as a horizontal line within the box. The endpoints for the whiskers are the upper and lower adjacent values, which are defined as the relevant quartile \pm three-halves of the interquartile range, and circles represent data points outside of the adjacent values.



Graphs by State

Figure S2. Observed and interpolated values for state-level GE adoption rates. The dots represented observed data, and the lines denote interpolated (predicted) values for the rates using two different. We interpolate missing data using predictions from a generalized linear model with a binomial family and a logit link function. The “prediction pooled” model pools all states and includes state fixed effects, while the “prediction state” model estimates a separate model for each state.




Figure S3. Spatial map of counties included in analysis. In-sample counties are dark grey.




Figure S4. Spatial and temporal variation of weather variables. We observe the weather variables at the county-year level, and construct boxplots for each year. Each box is defined by the upper and lower quartile, with the median depicted as a horizontal line within the box. The endpoints for the whiskers are the upper and lower adjacent values, which are defined as the relevant quartile \pm three-halves of the interquartile range, and circles represent data points outside of the adjacent values.




Figure S5. Precipitation difference, pre- and post-GE. For each county we calculate the average of the observed cumulative growing season precipitation across years in the pre- and post-GE periods. We report the percentage change of the latter over the former here.




Figure S6. Extreme heat difference, pre- and post-GE. For each county we calculate the average of the observed cumulative growing season degree days over 29°C across years in the pre- and post-GE periods. We report the percentage change of the latter over the former here.




Figure S7. Spatial map of water holding capacity (mm). For each county we observe the total volume of plant-available water that the soil can store within the root zone.




Figure S8. Spatial map of soil types. For each county we observe the most common soil texture.




Figure S9. Impacts of GE corn adoption (bushels per acre) by state. The parameter of interest in the regression model measuring the impact of GE adoption on yield is allowed to vary by state. Estimated impacts are then binned according to values in the figure legend.




Figure S10. Impacts of GE corn adoption (bushels per acre) by counties' water holding capacity and soil texture, reported for each state. County-level estimates from figures 2 and 3 are reported by state. The vertical bars are the distance between the highest and lowest value within each state. Dots denote the average of the county-level estimates within each state.