Dynamic treatment effects of crop insurance participation indicate positive impact on agricultural productivity

Micah Cameron-Harp[1] & Hunter D. Biram[2]

mcameronharp@ksu.edu, hbiram@uada.edu

[1]Kansas State University, Department of Agricultural Economics

[2]University of Arkansas, Department of Agricultural Economics and Agribusiness

Selected Paper for presentation at the 2023 Agricultural and Applied Economics Association Annual Meeting, Washington, D.C.; July 23-25

Abstract

As one of the primary U.S. agricultural policies, crop insurance is repeatedly studied by researchers to determine whether the changes in production behavior it induces are a source of inefficiency. One of the principle concerns is the potential for insured producers, through their "hidden-actions," to capture gains from riskier behavior while bearing a fraction of losses. In this work, we investigate whether such changes in behavior, known as moral hazard, occur using three decades of county-level yield data. Using recent estimators designed for staggered adoption settings, we find evidence that crop insurance participation leads to greater per acre yields at the county level. More importantly, the dynamic treatment effects we estimate indicate this effect grows over time. For corn, we find a statistically significant 7.59 bushel per acre increase six years after a county first reports crop insurance enrollment. The gradual improvement in yields we find over time provides valuable insights into potential mechanisms by which crop insurance may facilitate productivity gains.

1 Introduction

Crop insurance is one of the principal risk management components of U.S. agricultural policy, so understanding its impact on agricultural productivity and whether this impact is a product of moral hazard is critically important. Moral hazard refers to the changes in behavior that occur, or "hidden-actions" taken by producers, after adoption of crop insurance (Arrow, 1984). Broadly, there is evidence that firms in weather-sensitive industries, like agriculture, who use risk management strategies increase their investments and receive higher valuations (Pérez-González & Yun, 2013). However, unlike these generally positive impacts on farm finances, theory suggests the effect of crop insurance participation on productivity could be either positive or negative depending on the type of inputs and outputs (Nelson & Loehman, 1987; Ramaswami, 1993). For example, after enrolling in crop insurance, theory suggests a risk-averse producer is likely to use more risk-increasing inputs, like fertilizer, and consequently increase productivity.

In their investigation of the relationship between crop insurance and productivity, Roberts et al. (2006) use a difference-in-differences identification strategy to determine the impact of crop insurance participation on county yields in three states. Based on the positive effect of crop insurance participation on soybean and wheat yields that they find in Texas, and the lack of significant effects for Iowa and North Dakota, Roberts et al. (2006) conclude that there is limited evidence of moral hazard. Cornaggia (2013) uses a similar identification strategy, a triple differences approach, and finds a similar positive relationship between crop insurance use and yields. Interestingly, Cornaggia (2013) finds this effect is stronger for group-performance-based policies which suggests that certain program designs may be more susceptible to moral hazard.

Unlike these two prior works, Vigani and Kathage (2019) study the relationship between crop insurance participation and productivity using data from two European countries, France and Hungary. The results from the multinomial endogenous switching regression in Vigani and Kathage (2019) suggest that the effect of crop insurance on yields can be either positive or negative depending on farms' chosen risk management strategies and production characteristics. Similar to the Cornaggia (2013), Vigani and Kathage (2019) emphasize that the contract type selected plays a significant role in determining the direction of the effect. For example, in most cases, Vigani and Kathage (2019) find the effect of production-based contracts is positive.

To clarify the relationship between crop insurance participation and productivity, as well as the role that moral hazard plays, we use county-level data on per-acre yields and crop insurance participation from USDA NASS and the RMA Summary of Business datasets to create an unbalanced panel of county-level data across 41 states and 32 years. Like Roberts et al. (2006) and Cornaggia (2013), we employ a difference-in-differences identification strategy to investigate the effect of crop insurance participation on yields. However, recent econometric developments suggest a difference-in-differences approach in this empirical setting may be biased when treatment, participation in crop insurance in this context, does not occur simultaneously for all units or changes over time (Callaway & Sant'Anna, 2020; de Chaisemartin & D'Haultfoeuille, 2020; Goodman-Bacon, 2021). The RMA Summary of Business data we employ indicate insurance adoption is staggered, and we expect dynamic effects are likely given that the increase in productivity due to crop insurance participation may enable further investment into productivity-enhancing inputs.

To address these empirical challenges, we use an estimator which accommodates staggered adoption and dynamic treatment effects: the DID_l estimator from (de Chaisemartin & D'Haultfoeuille, 2020, 2022). In our initial analyses, we use the same treatment variables as in Roberts et al. (2006). The first is a binary indicator which takes a value of one for all years after a county first reports insurance being used on a crop-by-crop basis, and the second is an indicator variable equal to one in the first year a county reports a producer enrolling in “buy-up” coverage. To test the robustness of our initial results, we construct two additional treatment variables which capture the extent of crop insurance participation. The first, referred to as the enrollment-based participation (EBP) variable, records the percent of a county's planted acres enrolled in a crop insurance program for a given crop and year. The second, termed the liability-based participation (LBP) variable, is the percent of a county’s maximum potential liability covered by participating producers.

The primary contribution of this work is our estimation of dynamic treatment effects, or the impact that participation in crop insurance has on yields over time. While we report the average treatment effects, we focus on the dynamic treatment effects because they provide insights into the adjustments producers make over time. The dynamic effects of the effect on corn yields generated using the DID_l estimator and the indicator for any insurance being purchased, for instance, begin one year after a county’s producers first report using crop insurance and peak six years after treatment at a statistically significant 7.59 bushel per acre increase. The smooth upward trend in positive dynamic treatment effects we find suggests that moral hazard in this case may not involve riskier decision making. Instead, we posit that the income stabilizing effect of crop insurance participation allows producers to make greater capital investments in yield improving technologies. We conclude by outlining future directions of research which will clarify the mechanism driving our empirical results.

2 Data and variable construction

2.1 Crop yields and insurance participation data

We draw on county-level planted acreage and yield data from USDA-NASS and county-level data on insured acres and purchased liability from USDA-RMA Summary of Business spanning 1989-2020 for corn, soybeans, and wheat. We consider two measures of crop insurance participation which capture two unique sources of variation. EBP is constructed by taking a ratio of the insured acres to planted acres for a given county-crop-year combination and gives the crop insurance participation decision at the extensive margin. LBP is constructed by taking the ratio of the purchased liability to the maximum available liability for a given county-crop-year combination and provides a measure of crop insurance participation at the extensive and intensive margins. In other words, both EBP and LBP capture the decision to enroll in any insurance (i.e., extensive margin), but LBP also captures the decision to purchase any buy-up coverage into higher coverage levels (i.e., intensive margin) as noted by Goodwin, Vandeveer, and Deal (2004) and Connor and Katchova (2020).

The construction of EBP is straightforward using the raw data for planted and insured acres described above, but LBP must be constructed by using raw data on purchased liabilities and by calculating the maximum available liability. The maximum available liability is be calculated by taking the product of an expected price[1], expected yield, planted acreage, and the highest coverage level available. We only consider crop insurance participation for individual plans of crop insurance (i.e., Yield Protection and Revenue Protection), so the highest coverage level available is the 85% coverage level. Daily new crop futures prices during planting months on all three crops were retrieved using a Bloomberg terminal (Bloomberg, 2022), and annual measures for futures

[1] We calculate the expected price following RMA's practice outlined in Yu, Smith, and Sumner (2017).

prices were calculate by following Yu, Smith, and Sumner (2018) and used as a measure for expected price. Expected yields were calculated by running a regression of county-level yields on state fixed effects and state-specific linear trends and using the predicted values as detrended expected yields. Therefore, we assume linear changes in yield technology, and time-constant variables such as soil types are the same across counties in each state but allow for this heterogeneity across states. Crop-specific summary statistics for the yield and crop insurance participation variables are displayed in Table 1.

2.2 Temperature and soil moisture data

Daily data on counties' maximum and minimum temperatures were drawn from the Parameter-elevation Regressions on Independent Slopes Model repository maintained by Oregon State University (PRISM Climate Group, 2014), and the soil moisture data for the top ten centimeters of the soil profile were provided by NASA's SpoRT-LIS project (Kumar et al., 2006, 2008). We filter these data to only include days between March and the end of August, and then we create annual exposure variables for each county as in Schlenker and Roberts (2009). For soil moisture, expressed as the fraction of the soil volume filled with water, we construct six exposure bins. The first five variables indicate the days spent in each 0.08 cm3/cm3 interval beginning at zero and ending at 0.40 cm3/cm3, and the final variable indicates the number of days where soil moisture was greater than 0.40 cm3/cm3. The first of the temperature exposure variables aggregates all days with temperatures below 0 degrees Celsius, and the remaining five bins represent the days spent in each 10-degree interval ranging from zero to 50 degrees Celsius. Table 1 contains summary statistics for the temperature and soil moisture exposure variables calculated using the corn-specific panel dataset.

Table 1: Summary statistics for yield, crop insurance participation, and weather variables. Statistics for the temperature and soil moisture variables were calculated using the corn specific dataset.

	Mean	Median	S.D.	Min.	Max.
Corn					
Yield (bushels/acre)	121.07	121.00	40.88	0.00	270.20
Any insurance indicator	0.87	1.00	0.33	0.00	1.00
Buy-up indicator	0.87	1.00	0.34	0.00	1.00
EBP	6.82	5.25	7.59	0.00	100.00
LBP	5.88	4.09	10.23	0.00	100.00
Soybeans					
Yield (bushels/acre)	37.30	37.10	11.20	0.70	80.40
Any insurance indicator	0.90	1.00	0.30	0.00	1.00
Buy-up indicator	0.90	1.00	0.30	0.00	1.00
EBP	6.90	5.38	7.51	0.00	100.00
LBP	5.18	4.11	5.53	0.00	100.00
Wheat					
Yield (bushels/acre)	47.29	45.50	18.00	0.00	153.80
Any insurance indicator	0.82	1.00	0.38	0.00	1.00
Buy-up indicator	0.82	1.00	0.39	0.00	1.00
EBP	50.53	18.51	46.38	0.00	100.00
LBP	49.65	14.98	47.03	0.00	100.00
Temperature exposure (days)					
Below zero	8.15	5.83	7.91	0.00	55.96
1-10°C	26.67	27.82	11.31	0.00	90.33
11-20°C	54.33	54.93	11.45	6.76	115.87
21-30°C	78.03	78.30	16.41	6.22	133.06
31-40°C	16.78	13.13	14.19	0.00	95.60
>40°C	0.04	0.00	0.46	0.00	24.60
Soil moisture exposure					
.01-.08 cm^3/cm^3	0.68	0.00	5.46	0.00	172.00
.09-.16 cm^3/cm^3	14.53	0.00	26.92	0.00	181.00
.17-.24 cm^3/cm^3	51.60	43.00	41.77	0.00	177.00
.25-.32 cm^3/cm^3	88.36	96.00	44.50	0.00	182.00
.33-.40 cm^3/cm^3	26.34	14.00	30.02	0.00	183.00
>.41 cm^3/cm^3	2.48	0.00	6.10	0.00	58.00

3 Empirical Approach

In this paper, we utilize a Difference-in-Differences identification strategy to recover the effect of crop insurance participation on productivity. At its core, this approach compares the change in bushels per acre in counties who adopted crop insurance to those who did not (Bertrand et al., 2004). The crucial parallel trends assumption requires that the change in yields over time for these two groups of counties, the adopting and control counties, would have been the same in the absence of crop insurance. Given the national dataset we employ in this work, and the 30-year time frame, there are a number of reasons why this assumption may not be tenable. To demonstrate potential areas of concern, and our means of addressing them, we use the following reduced form model representing the relationship between crop yields and crop insurance participation:

$$Y_{i,t} = \gamma_t + \Phi X_{i,t} + \Psi Z_i + \beta^{DD} D_{i,t} + u_{i,t}. \qquad \text{(Eq. 3.1)}$$

In Equation 3.1, $Y_{i,t}$ is the yield of county i in year t, the dependent variable of interest. The independent variable we are concerned with is $D_{i,t}$, the insurance participation variable, and the coefficient β^{DD} is the effect to be estimated. Aside from $u_{i,t}$, the idiosyncratic error term, the remaining variables represent potential confounding factors. The γ_t term is a year-specific fixed effect representing nationwide factors influencing yields, such as technological innovations or policy changes. The matrix of time-varying covariates, $X_{i,t}$, represents the potential for local weather conditions to cause differential trends between counties over time. Similarly, Z_i is an equivalent matrix for time-invariant covariates, like soil characteristics, which could lead certain counties to outperform others regardless of the role of crop insurance. In the following sections, we describe how the two estimators we employ attempt to address each of these concerns.

3.1 Two-way fixed effects

The two-way fixed effects (TWFE) results in our study use the following estimator produced by applying the within transformation to Equation 3.1:

$\ddot{Y}_{i,t} = \gamma_t + \Phi\ddot{X}_{i,t} + \beta^{DD}\ddot{D}_{i,t} + \ddot{u}_{i,t}$; where $\ddot{x}_{i,g,t} = x_{i,g,t} - \bar{x}_{i,g}$. (Eq. 3.2)

The two-way fixed effects estimating equation above addresses the possibility of nationwide, annual shocks using the year fixed effect, γ_t, and the within transformation accounts for the impact of time-invariant characteristics, Z_i. Including the time-varying covariates, now time demeaned in $\ddot{x}_{i,t}$, implies that we are using a conditional version of the parallel trends assumption. The conditional parallel trends assumption necessary for TWFE with time-varying covariates to identify the effect of insurance participating requires any differences between trends in counties' counterfactual outcomes be explained by a linear model in $\ddot{x}_{i,t}$ (de Chaisemartin and D'Haultfoeuille 2020).

In our context, the conditional parallel trends assumption for TWFE allows counties to have differences in counterfactual yield outcomes, so long as they are explained by changes in their annual temperature and soil moisture conditions. However, even if this conditional parallel trends assumption holds, the TWFE estimates of β^{DD} will be biased in this empirical setting. As mentioned previously, crop insurance participation did not occur simultaneously across the country. In such staggered adoption settings, TWFE estimates do not recover the average treatment effect because they make faulty comparison using the post-adoption behavior of early adopters as control observations (Callaway & Sant'Anna, 2020; de Chaisemartin & D'Haultfoeuille, 2020; Goodman-Bacon, 2021).

3.2 de Chaisemartin and D'Haultfoueille (2022)

Our preferred estimator, from de Chaisemartin and D'Haultfoeuille (2022), addresses the possibility of making misleading comparisons between early and later adoption counties by comparing outcomes between cohorts. In our setting, a cohort is comprised of all counties who first report a producer using crop insurance in year g. Like in de Chaisemartin and D'Haultfoeuille (2022), we define $N^{1}_{t,\ell}$, as the number of counties who first report crop insurance participation ℓ years before t, and we let N^{nt}_{t} represent how many counties have yet to report crop insurance participation at time t. We then employ the following estimator to recover the average and dynamic treatment effects of crop insurance participation:

$$\widehat{DID}^{X}_{t,\ell} = \frac{1}{N^{1}_{t,\ell}} \sum\nolimits_{\forall\, i \in\, g\, =\, t-\ell} \left(Y_{i,g,t} - Y_{i,g,t-\ell-1} - (X_{i,g,t} - X_{i,g,t-\ell-1})'\hat{\theta}_{0}\right) -$$

$$\frac{1}{N^{nt}_{t}} \sum\nolimits_{g=t+1} \left[\sum\nolimits_{i=1}^{N_{g,t}} \left(Y_{i,g,t} - Y_{i,g,t-\ell-1} - (X_{i,g,t} - X_{i,g,t-\ell-1})'\hat{\theta}_{0}\right)\right]. \qquad \text{(Eq 3.3)}$$

Equation 3.3 compares the change in crop yields from time $t - \ell - 1$ to t for counties who first reported crop insurance participation ℓ years ago with counties who have not reported any participation by t. As such, it prevents making erroneous comparisons between counties who adopt insurance later in the time series and early adopters. Similar to TWFE, differences in yield outcomes due to temperature and soil moisture are addressed in Equation 3.3 by regressing the change in yields within the control group, $\left(Y_{i,g,t} - Y_{i,g,t-\ell-1} \mid g > t\right)$, on the change in its time-varying covariates, $\left(X_{i,g,t} - X_{i,g,t-\ell-1} \mid g > t\right)$, and time fixed effects. Defining $\hat{\theta}_{0}$ as the coefficients resulting from this regression for $\left(X_{i,g,t} - X_{i,g,t-\ell-1}\right)$, Equation 3.3 removes the effect of covariate driven differential trends by subtracting the counterfactual change in yields predicted for a county based on its covariates from its observed change in yields.

The approach to addressing differential trends displayed in Equation 3.3 requires a conditional parallel trends assumption also identical to that necessary for the TWFE estimator. For the $DID_{t,\ell}^{X}$ estimator to recover the effect of crop insurance participation, any differential trends in counterfactual outcomes must be explainable using a linear model in $X_{g,t} - X_{g,t-\ell-1}$. However, given the heterogeneity amongst U.S. growing regions, we anticipate that this assumption may not hold. As such, we replicate the $\widehat{DID}_{t,\ell}^{X}$ estimation using the non-parametric matching feature of de Chaisemartin et al.'s *did_multiplegt* STATA package such that outcomes are only compared within USDA-ERS Farm Resource Regions (F.R.R.). For the dynamic treatment effects, we estimate effects ten years following, and three years prior to, the year in which insurance participation begins.

4 Results and Discussion

In Table 2, we present the estimated effects for both the TWFE and DID_l estimators. While the DID_l and two-way fixed-effects estimators both use county and year fixed effects, and control for differences in local weather, the resulting estimates differ in sign. For example, two-way fixed effects estimation suggests that crop insurance participation has a significant, negative effect on corn yields of -1.67 or -1.41 bushels per acre using the any insurance and "buy-up" insurance indicator variables, respectively. The results from the DID_l estimator, in contrast, indicate a significant, increase of 3.40 or 3.51 bushels per acre. We see a similar difference in signs when comparing the TWFE and DID_l results for soybeans, although the difference between the two estimators' results is smaller in magnitude. For wheat, in contrast, there is no difference in sign and little difference in the magnitude of the TWFE and DID_l results. The impact of using non-parametric matching displayed in Table 2 also differs when comparing the results for wheat with those of the other two crops. For corn and soybeans, using non-parametric matching by

USDA-ERS Farm Resource Regions widens the disparity between the TWFE and DID_l estimates. But for wheat, using non-parametric matching reduces the difference between TWFE and DID_l estimates.

One concern with the estimates for soybeans and wheat in Table 2 is the failure of the placebo test we observe. For both the soybean and wheat DID_l estimates, the p-value for a test that the pre-treatment effects of crop insurance in the three periods prior to the first recorded instance of participation are jointly zero is always less than 0.10. While this is not conclusive evidence that the conditional parallel trends assumption does not hold (Roth, 2022), we believe it warrants interpreting the results for soybeans and wheat with caution. As such, we only display the results for corn when using our continuous measures of insurance participation in Table 3 and the dynamic treatment effects in Figure 1. In Table 3, the similar disparity between TWFE and DID_l estimates we observe for both the enrollment-based and liability-based participation measures, E.B.P. and L.B.P. respectively, indicates that our results for corn are robust to the specification of our dependent variable. Then, in Figure 1, the reason for the difference between TWFE and DID_l estimates is readily apparent. If the post-adoption behavior of early adopters is used as a control for late adopters of crop insurance, the increase in yields experienced by early adopters will result in a negative treatment effect estimate for later adopters in comparison. For both variables, the effect of participating in crop insurance peaks at a nearly 7 bushel per acre increase 6 years after a county first records one of its constituent producers enrolling in crop insurance.

Table 2: Average effect of crop insurance participation on county yields (bushels per acre) with 95% confidence intervals in parentheses. For all estimators, standard errors are clustered at the county level. The $DID_{t,\ell}^{X}$ estimates represent the average effect over the ten years following a county's first year participating or reporting "buy-up" enrollment.

	Estimator		
	TWFE	$DID_{t,\ell}^{X}$	$DID_{t,\ell}^{X}$
Corn			
"Any" indicator	-1.67 (-3.26, -0.08)	3.40 (1.00, 5.80)	3.91 (0.68, 7.14)
"Buy-up" indicator	-1.41 (-2.98, 0.15)	3.51 (0.55, 6.50)	3.69 (0.79, 6.60)
Year fixed effects	Yes	Yes	Yes
Weather covariates	Yes	Yes	Yes
F.R.R. matching	No	No	Yes
Soybeans			
"Any" indicator	-1.23 (-1.66, -0.80)	0.93 (0.11, 1.76)	1.09 (0.24, 1.95)
"Buy-up" indicator	-1.23 (-1.66, -0.80)	0.94 (0.01, 1.88)	1.07 (0.19, 1.96)
Year fixed effects	Yes	Yes	Yes
Weather covariates	Yes	Yes	Yes
F.R.R. matching	No	No	Yes
Wheat			
"Any" indicator	0.12 (-0.49, 0.73)	0.85 (-0.23, 1.92)	0.79 (-0.25, 1.82)
"Buy-up" indicator	0.02 (-0.59, 0.64)	0.74 (-0.20, 1.67)	0.53 (-0.38, 1.44)
Year fixed effects	Yes	Yes	Yes
Weather covariates	Yes	Yes	Yes
F.R.R. matching	No	No	Yes

Note – F.R.R. matching indicates non-parametric matching based on USDA-ERS Farm Resource Regions was used.

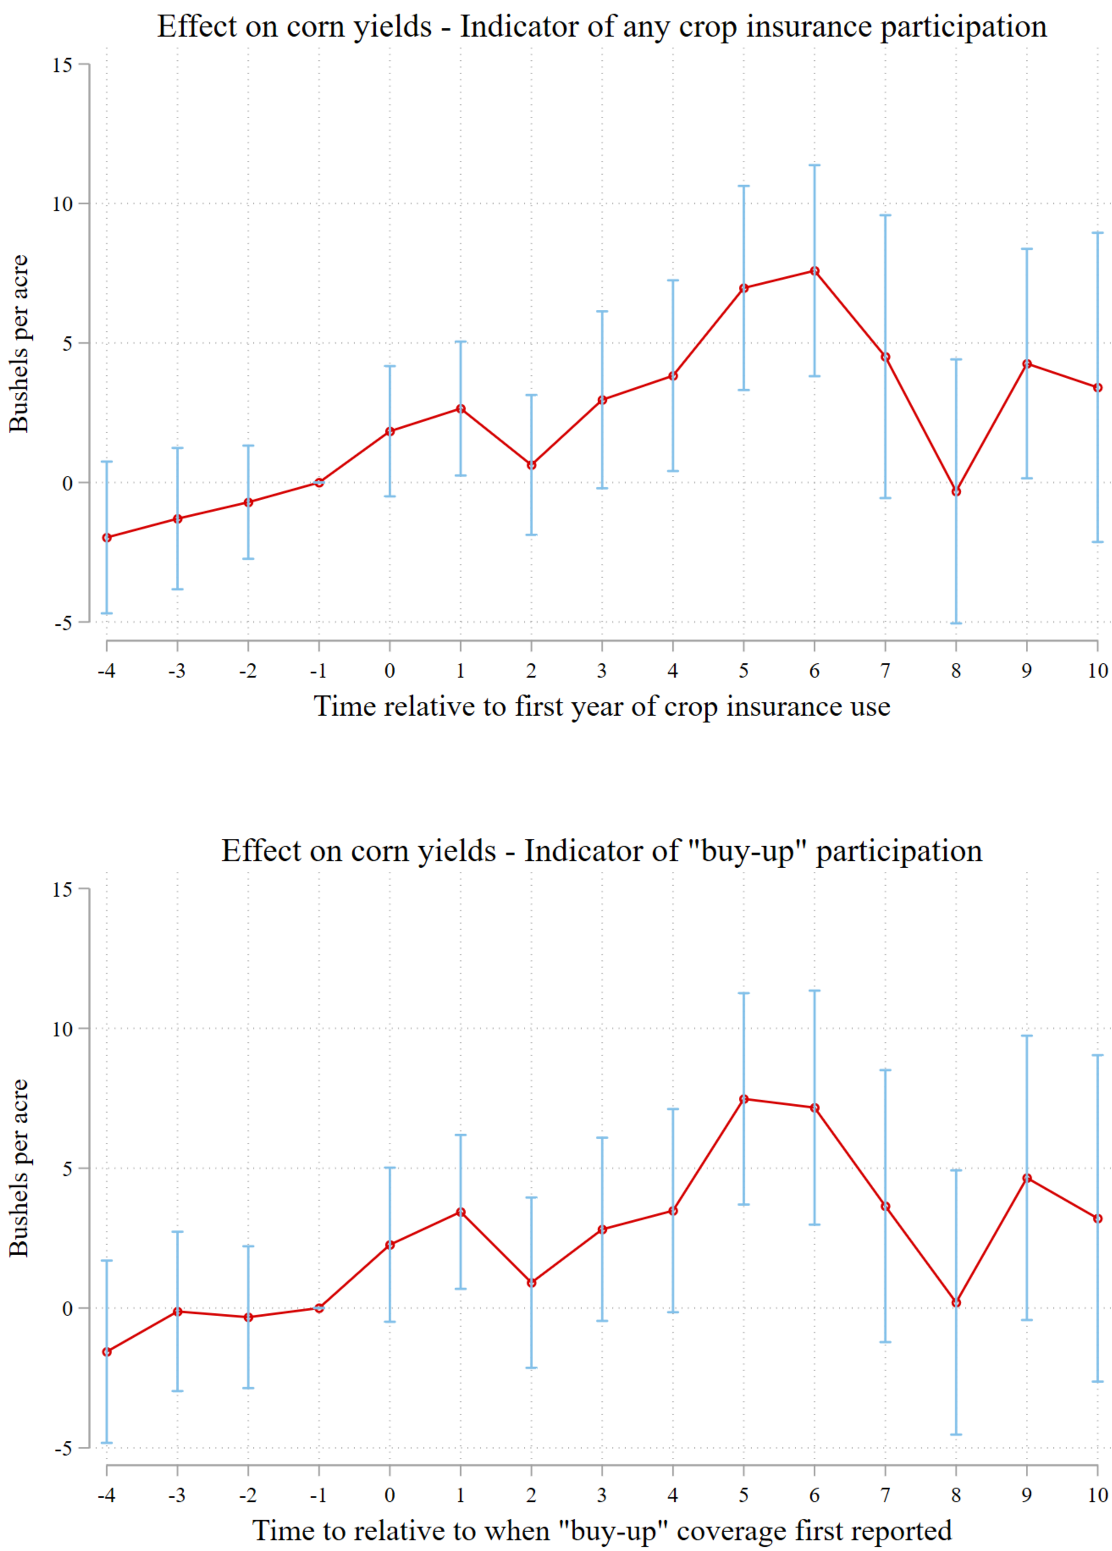

Figure 1: Dynamic treatment effects of crop insurance participation on corn yields using indicator of any crop insurance participation (top panel) and "buy-up" enrollment (bottom panel).

Table 3: Average effect of crop insurance participation on county yields (bushels per acre) with 95% confidence intervals in parentheses. For all estimators, standard errors are clustered at the county level. The $DID_{t,\ell}^{X}$ estimates represent the average effect over the ten years following a county's first year participating.

	Estimator	
	TWFE	$DID_{t,\ell}^{X}$
Corn		
E.B.P.	-0.05	0.55
	(0.09, -0.00)	(0.14, 0.97)
L.B.P.	-0.75	-0.00
	(-0.83, -0.68)	(-0.44, 0.44)
Year fixed effects	Yes	Yes
Weather covariates	Yes	Yes
F.R.R. matching	No	Yes

Note – E.B.P. is the measure of enrollment-based participation, and L.B.P. is the liability based measure of insurance participation. F.R.R. matching indicates non-parametric matching based on USDA-ERS Farm Resource Regions was used.

5 Conclusion

Without question, understanding the impact of crop insurance on agricultural productivity is an important means of determining whether moral hazard, or "unseen-actions" by enrolled producers, is a source of inefficient government expenditure. However, clarifying this relationship also serves to deepen our understanding of how producers adjust their behavior following participation in new programs or adoption of novel technologies. In this paper, we use county level data on crop yields and a variety of measures of crop insurance participation to study this relationship in the U.S. agricultural context.

We find that crop insurance participation led to an increase in county level crop yields, and that this increase grew in magnitude over the five years following a county's first reported instance of crop insurance participation. While this may indicate producers' make increasingly risky decisions following enrollment in crop insurance, the gradual increase in productivity we find for corn could just as feasibly be the result of gradual investments into productivity improving capital. As such, further research is necessary to determine whether the dynamics we observe are due to the changes in input use or the income-stabilizing effect of crop insurance and any consequent changes to capital investments (Girao et al., 1974).

More broadly, we think this work provides a useful empirical example of when staggered adoption and dynamic treatment effects can cause two-way fixed effects to produce statistically significant, yet incorrectly signed, estimates. By presenting both the two-way fixed effects and DID_l estimators, we illustrate their relative advantages as well as when their necessary assumptions are met within an empirical context.

References

Arrow, K. J. (1984). *The Economics of Agency* (p. 30). Stanford Univ CA Institute for Mathematical Studies in the Social Sciences.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). HOW MUCH SHOULD WE TRUST DIFFERENCES-IN-DIFFERENCES ESTIMATES? *Quarterly Journal of Economics*, *119*(1), 249–275.

Bloomberg L.P. (2022) Retrieved March 9, 2022 from Bloomberg terminal.

Callaway, B., & Sant'Anna, P. H. C. (2020). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

Connor, L., & Katchova, A. L. (2020). Crop Insurance Participation Rates and Asymmetric Effects on U.S. Corn and Soybean Yield Risk. *Journal of Agricultural and Resource Economics*, *45*(1), 1–19. https://doi.org/10.22004/AG.ECON.298431

Cornaggia, J. (2013). Does risk management matter? Evidence from the U.S. agricultural industry. *Journal of Financial Economics*, *109*(2), 419–440. https://doi.org/10.1016/j.jfineco.2013.03.004

de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, *110*(9), 2964–2996.

de Chaisemartin, C., & D'Haultfoeuille, X. (2022). *Difference-in-differences estimators of intertemporal treatment effects* (p. No. w29873). National Bureau of Economic Research.

Girao, J. A., Tomek, W. G., & Mount, T. D. (1974). The Effect of Income Instability on Farmers' Consumption and Investment. *The Review of Economics and Statistics*, *56*(2), 141. https://doi.org/10.2307/1924433

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014

Goodwin, B. K., Vandeveer, M. L., & Deal, J. L. (2004). An Empirical Analysis of Acreage Effects of Participation in the Federal Crop Insurance Program. *American Journal of Agricultural Economics*, *86*(4), 1058–1077. https://doi.org/10.1111/j.0002-9092.2004.00653.x

Kumar, S. V., Peters-Lidard, C. D., Eastman, J. L., & Tao, W. K. (2008). An integrated high-resolution hydrometeorological modeling testbed using LIS and WRF. *Environmental Modelling and Software*, *23*(2), 169–181. https://doi.org/10.1016/j.envsoft.2007.05.012

Kumar, S. V., Peters-Lidard, C. D., Tian, Y., Houser, P. R., Geiger, J., Olden, S., Lighty, L., Eastman, J. L., Doty, B., Dirmeyer, P., Adams, J., Mitchell, K., Wood, E. F., & Sheffield, J. (2006). Land information system: An interoperable framework for high resolution land surface modeling. *Environmental Modelling and Software*, *21*(10), 1402–1415. https://doi.org/10.1016/j.envsoft.2005.07.004

Nelson, C. H., & Loehman, E. T. (1987). Further Toward a Theory of Agricultural Insurance. *American Journal of Agricultural Economics*, *69*(3), 523–531. https://doi.org/10.2307/1241688

Pérez-González, F., & Yun, H. (2013). Risk management and firm value: Evidence from weather derivatives. *Journal of Finance*, *68*(5), 2143–2176. https://doi.org/10.1111/jofi.12061

PRISM Climate Group. (2014). *PRISM daily data*. Oregon State University. https://prism.oregonstate.edu

Ramaswami, B. (1993). Supply Response to Agricultural Insurance: Risk Reduction and Moral Hazard Effects. *American Journal of Agricultural Economics*, *75*(4), 914–925.

Roberts, M. J., Key, N., & O'Donoghue, E. (2006). Estimating the extent of moral hazard in crop insurance using administrative data. *Review of Agricultural Economics*, *28*(3), 381–390. https://doi.org/10.1111/j.1467-9353.2006.00303.x

Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, *4*(3), 305–322. https://doi.org/10.1257/aeri.20210236

Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(37), 15594–15598. https://doi.org/10.1073/pnas.0906865106

Vigani, M., & Kathage, J. (2019). To Risk or Not to Risk? Risk Management and Farm Productivity. *American Journal of Agricultural Economics*, *101*(5), 1432–1454. https://doi.org/10.1093/ajae/aaz020

Yu, J., Smith, A., & Sumner, D. A. (2018). Effects of crop insurance premium subsidies on crop acreage. *American Journal of Agricultural Economics*, *100*(1), 91–114. https://doi.org/10.1093/ajae/aax058