Measuring rotation effects in the US Corn Belt

Matthieu Stigler

August 2019

Abstract

Corn and soybeans are predominantly cultivated in a rotation scheme in the US Midwest. Understanding the extent of the rotation effect is important when analyzing the dynamics of crop choice, as large rotation benefits can make the overall area response more sluggish. Likewise, rotations play a beneficial role for the environment, allowing to save on fertiliser. Most estimates of the rotation effect come from experimental farms, which allow a clean identification on a very restricted subsample of fields. I take here another approach, using a large dataset of close to two million fields in the US Midwest, over ten years. Using observational data raises new challenges. For one, choice of field and crop is clearly not random: fields with higher fertility tend to be planted to corn more often than to soybean. This suggest that estimating rotation comparing *between* field is not the right approach, and one should seek *within* field variation. This raises however a new difficulty: by its very definition, a field in rotation was not planted to the same crop the year before, implying that we cannot use a standard difference-in-difference approach.

In this paper, I take advantage of a dataset of field-level yields over 9 states in the US Midwest, spanning ten years. I make three contributions. I firstly document the presence of selection bias, showing evidence that more fertile fields tend to be planted to corn. An interesting paradox that arises from this comparison is that fields doing always corn have higher average yields compared to those doing always rotation, despite not benefiting from the rotation effect. Second, I address formally the challenges of identifying the rotation effect in presence of non-random crop choice. I use recent results from the literature on the relationship between fixed effects and difference-in-difference estimators, and show how to adapt these to the specific context of rotation and its missing data problem. Finally, I extend the analysis to corn and soybeans fields that were previously cultivated to other crops, and show how different the identification approach is.

1 Introduction

The US Corn Belt is an important producer of corn and soybeans, concentrating two thirds of the area and volume of the US production. In this region, corn and soybeans are usually cultivated in rotation, alternating between corn and soybeans year after year on the same field. Benefits of rotation are multiple (for a survey, see Porter et al., 1997; Farmaha et al., 2016). On one hand, changing crop from year to year reduces pest, given that most parasites are specific to a given crop, and cannot survive more than a year without their crop. Further, soybeans fix nitrogen, reducing the need for fertilizer. As a result, a corn field planted in rotation will have a higher expected yield, as well as a reduced need for fertilizer, than if corn had been planted the previous year.

The presence of this rotation effect has several important consequences for the dynamics of agricultural supply, as well as for the environmental impacts of agriculture. Choice of crop by the farmer is usually dictated by a comparison of respective expected profit. With rotation effects, the farmer takes additionally into account that despite a possible lower return on the second crop, choosing that lower productivity crop might be profitable on the long term. At the aggregate level, rotation effects alter the responsiveness of crop choice to prices. Intuitively, price movements need to be larger to induce a field following rotation to be affected to say corn only. But on the other side, the same

Figure 1: Rotating versus non-rotating fields in Iowa

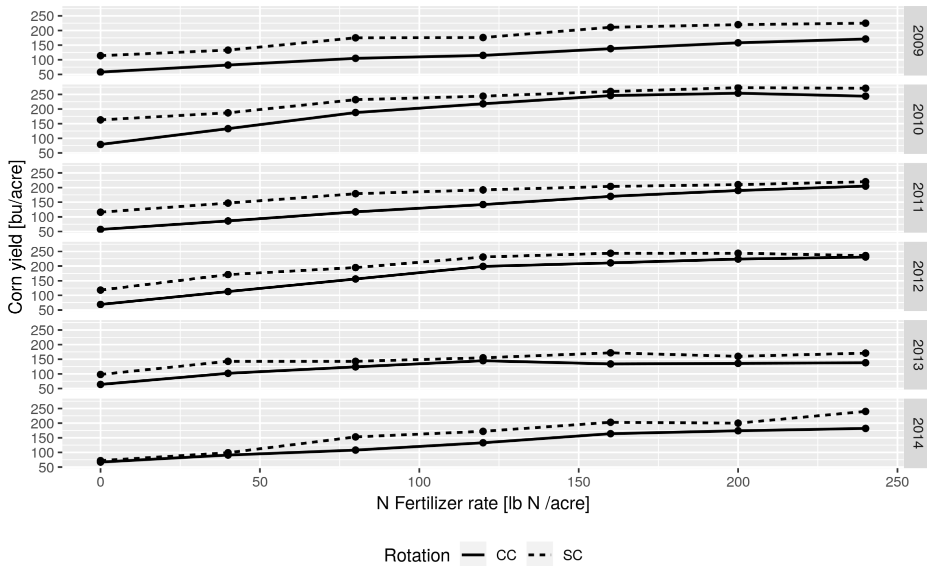

Source: Sawyer (2014)

price increase will induce soy-only fields to switch to rotation. Hendricks et al. (2014b) show also that rotation effects makes the long-term area response weaker than the short-term one, contrary to the traditional model of partial adjustment (Nerlove, 1956; Nerlove and Bessler, 2001). Furthermore, as I show in Section **??**, strong rotation effects suggest that an expansion of area made by foregoing rotation introduces fields of lower productivity. This implies that yield response observed at the aggregate level is weaker than observed at the field-level, due to this composition effect. On the environmental side, rotation acts as a natural fertilizer, cutting down the need for chemical fertilizer. This reduces in turn the risk of leaching and thereby potentially reduces damage to water streams and algal blooms. Hendricks et al. (2014a) show how increases in corn prices induce a reduction in the number of rotating fields, and estimate the impact on the hypoxic zone in the Gulf of Mexico. Lastly, estimates of the rotation effect are extensively used to formulate crop choice and nitrogen use recommendations by extension agencies (Morris et al., 2018).

Various studies have sought to estimate the effect of rotation, using either experimental plots or observational data. Porter et al. (1997) study a small sample of experimental plots in Minnesota and Wisconsin, and report a 15% increase in corn yields when previously cropped with soybeans, with some heterogeneity where lower-yielding fields see a higher increase, close to 25%. Similar numbers were obtained for soybeans. Another set of experimental data, used by Hennessy (2006) and Livingston et al. (2015), suggests rotation effects increasing yields for both crops by 25%. Based on a similar dataset from experimental plots by the same Iowa University Sawyer (2014, 2009), Figure 1 compares the yields of rotating (SC, dashed) and corn-corn (CC) fields, at different levels of fertilizer rate, over multiple years. It is clear from the picture that rotating fields always enjoy a higher yield. This rotation effect tends to decrease at higher levels of fertilizer/yield. Fertilizer rates on this experiment were set ranging from 0 to 240, which do not correspond to levels that would be typically found in practice. A farmer following the "economically optimal N rate" recommendations of 160 [lb N /acre] for rotating fields, and 200 [lb N/acre] for non rotating ones during that period (Sawyer, 2014) would obtain an average rotation effect of 15%. This number, averaged over 2009 to 2014, exhibits however a large variance, with a minimum rotation effect of 2% in 2010 and maximum of 33% in 2009.

Seifert et al. (2017) use a rich observational dataset obtained from the USDA containing around

Table 1: Average shares of crop sequences

Crop Lag	Crop	Percentage
Corn	Corn	15
Corn	Soy	40
Soy	Corn	39
Soy	Soy	5

Source: own computation from CDL. Data is the 2008-2017 average of the year-to-year shares.

120'000 fields in the U.S. Midwest.[1] They find lower rotation effects compared to studies based on experimental plots. Their favourite estimate for corn is 4.3% (6.7 [bu/acre]) for corn and 10% for soybeans (4.7 [bu/acre]). (Farmaha et al., 2016) use data from farmer-reported yields in Nebraska, over from four to eight years. They find a 2-5% (3-9 [bu/acre]) rotation effect for corn, and a 5% 2-years rotation effect for soybeans, comparing the sequence CCS to SCS. They note that their estimate of the rotation effect for corn is lower than results from experimental plot, and suggest that it might be due to the fact that their sample contain mostly high-yield fields, for which the rotation effect is smaller.

The discussion of the rotation effect so far was focused on the previous crop. It is possible however that the dependence of current yields on previous crop choice goes beyond one year, i.e. yields exhibit more than a one-year memory. Most of the studies indicate that corn has only a one year memory: there is no noticeable difference between a $\langle SSC \rangle$ and a $\langle CSC \rangle$ sequence. On the other side, soybeans is usually found to exhibit a two-year memory, where fields with two previous years of corn ($\langle CCS \rangle$) giving higher yields than fields with one year of corn and soybeans previously ($\langle SCS \rangle$) (Farmaha et al., 2016; Porter et al., 1997).

In this paper, I revisit these estimates using a novel dataset of satellite-based yield estimates, for more than one million fields, in 9 states of the Corn Belt, over ten years. This dataset covers a large proportion of fields in each county, and hence reflects natural variations in field conditions not necessarily captured by previous studies. I revisit also the methods used to estimate the rotation effect, arguing that previous studies do not capture the causal effect of rotation, but a mixed estimate including selection effects. I use recent developments in the panel data literature to wipe out selection effect and to account for heterogeneity.

2 Stylized facts on rotation

The data on crop choice used here is the USDA Cropland Data Layer (CDL) from Boryan et al. (2011). It is published every year, started in 2000 with a few states in the Corn Belt, and covered the entire US from 2008 onwards. Spatial resolution is 30m. Classification accuracy in the US Corn Belt is rather high, with a claimed accuracy of close to 90% for corn and soybeans. A few studies have used this dataset to study rotation patterns. See Chapter II, Section **??** for more details on the data. Several studies used the CDL dataset to analyse rotation patterns, with the general finding that monocropping increased over the 2000-2010 period, see Chapter II, Section **??** for a survey.

In this study, I analyze corn and soybeans during 2008 to 2017 period. This still represents a large part of the total cropland, 95% for the 3I states (Iowa, Illinois and Indiana), and 70% on average in the six remaining ones. Table 1 shows the average one-year-ahead cropping pattern between corn and soybeans, for fields doing always either corn or soybeans. A large part (80%) of the fields are under rotation from one year to the other. Corn-corn sequences ($C - C$) are much more frequent than the soy-to-soy ones ($C - C$).

The previous table shows typical cropping patterns, averaging year to year shares. But how many of the fields end up in a long-term $\langle CCCCCCCCCC \rangle$ sequence over the ten years? Figure 2 shows the

[1] The six states in the study are Iowa, Illinois, Indiana, Minnesota, Nebraska and South Dakota.

Figure 2: Recursive crop choice for all fields doing exclusively C or S

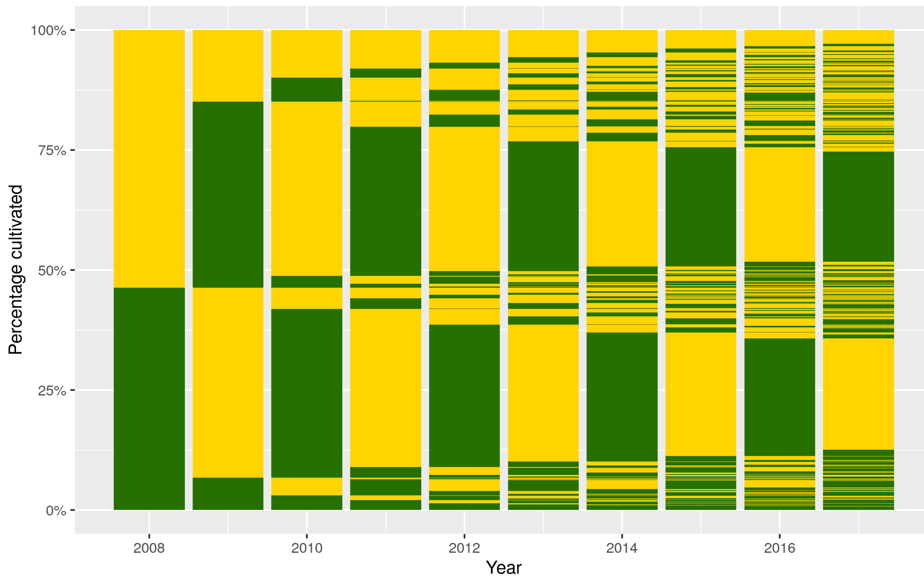

Source: own computation, from CDL.

conditional sequences of corn-soy choice for all fields in the sample, starting in 2008. The first year shows the share of corn and soybeans fields, with 48% soy and 52% corn. The second year shows the corn and soy choice, conditional on the previous year. Numbers of rotating fields here are very close to the ones in Table 1 (which are averages over the whole 2008-2017 period), with 80% of rotating fields. The year after, 2010, follows with the same logic, showing the crop choice in 2010 conditional on the crop choice in 2009, itself conditional on 2009. We see now that the share of fields rotating since 2008 decreases slightly, becoming 70%. Going till the last year 2017, the share of fields that always rotated, over the whole period of 2008-2017, is still at 46%. The share of fields practising monocropping over the whole period shrinks to 3% for corn, and to a very small 0.7% for soybeans.

The fact that there is a very large proportion of fields doing *always rotation* as well as *always corn* is a bit of a puzzling phenomenon. At first sight, crop choice on such fields seems to be fully price insensitive, going against models relating crop choice to price variations (Hennessy, 2006; Hendricks et al., 2014b). It is however interesting to note that in Livingston et al. (2015)'s model, the simple *always rotating* strategy turns out to be only slightly dominated by a much more complex strategy based on dynamic linear programming. This suggests a possible explanation of the always rotating strategy as a simple rule of thumb for farmers. An alternative explanation would be that on these fields, the rotation effect is strong enough to overcome any large deviation in the ratio of the price of corn to soybean. According to the same explanation, the fields producing *always corn* would be fields with a very low rotation effect, or a low soy yield. As the explanation relies on a purely unobservable component, it makes it difficult to test in practice. The field-level dataset here does not prove very useful to disentangle alternative explanations, as we observe here only fields, not farms. This prevents us from investigating whether always rotating patterns are farmer or field specific.

Figure 3 shows the percentage of fields practising always rotation in each county, as compared to all other fields doing exclusively corn or soybeans.

The presence of a large proportion of fields practising always or never rotation has interesting consequences in terms of estimating the rotation effect. These are two subsets on which rotation is either always or never observed. If we are to estimate rotation effects based on *within-field* variation (i.e., using plot fixed effects), these subsets do not experience any variation, and are dropped from the analysis. It is only when we are incorporating *between-field* comparisons, either by adding time

Figure 3: Location of fields always rotating over 2008-2017

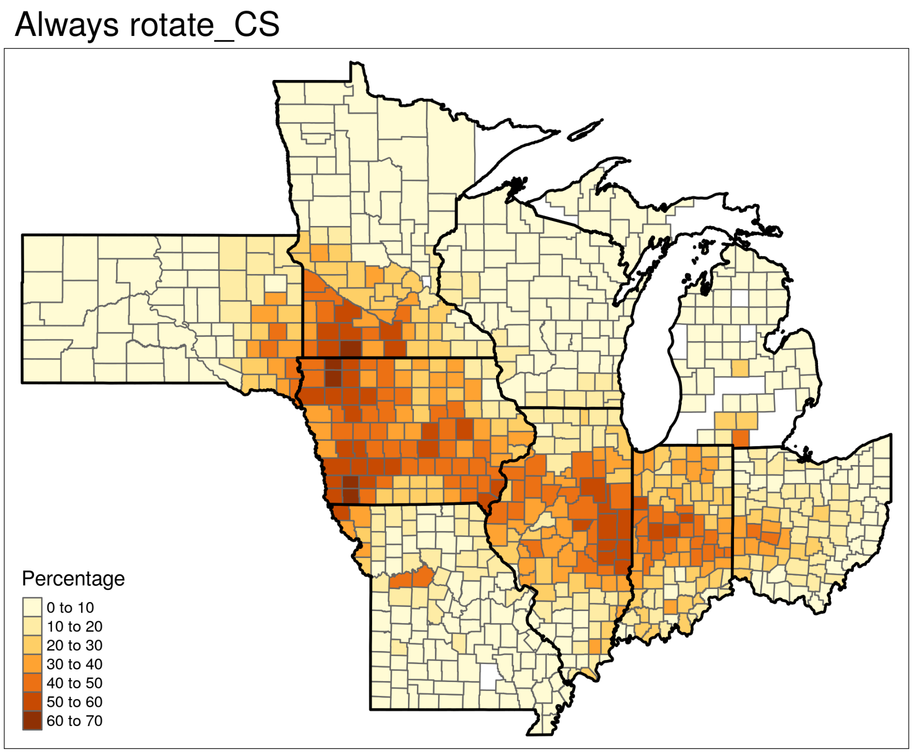

Source: own computation, from CDL.

Figure 4: Average yields for given rotation sequences

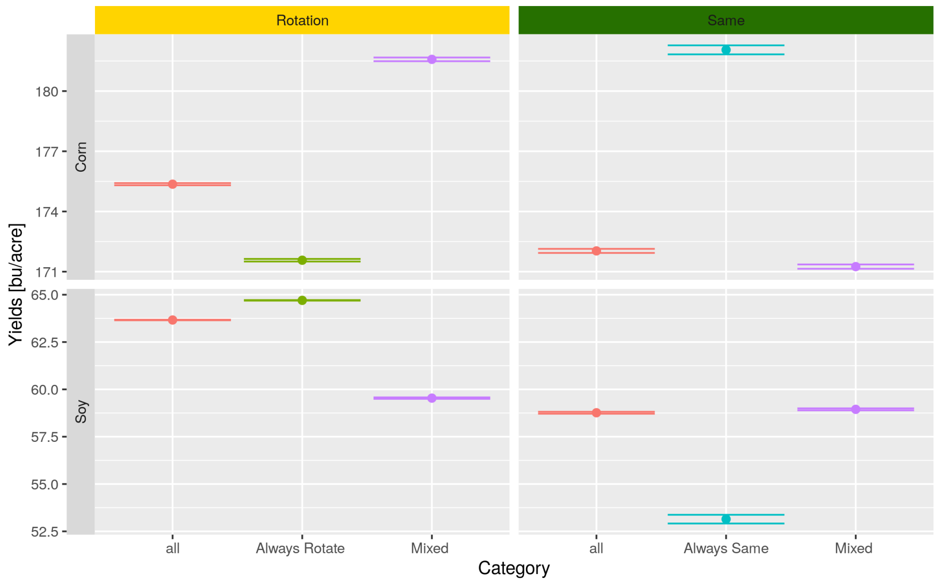

Note: Points represent averages, and bars 95% confidence intervals. Source: own computation, from CDL.

controls, or by removing fixed effects at the plot level, that these subsets can contribute to the analysis. This raises however the question of how comparable are the always and never rotaters compared to the other fields doing *sometimes* rotation.

Comparing average yields on rotating and non-rotating sequences among different subsets, Figure 4 shows an interesting paradox. The figure shows average yields for rotating sequences and non-rotating ones (right panel, *Same*), for each crop. The first number in red computes these means over the whole sample of fields. A given field might contribute to the rotating mean, the non-rotating one, or both. We see here that the rotation effect is about 4 [bu/acre] for corn, and about 5 [bu/acres] for soy. As we discussed previously, there is however a large part of fields who always rotate, and hence only contribute to the first mean. Likewise, the never-rotaters are only counted in the non-rotating mean. The next numbers in the figure show average yields depending on the long-term status of the fields, i.e. whether fields are always, never or sometimes rotating. Surprisingly, if one were to compute the rotation effect using only the always and never rotaters, we would find a negative effect, of -12 [bu/acre]! On the other side, the same comparison for soybeans leads to a larger effect than previously, of about 12 [bu/acre]. This basically means that we are comparing fields with very different environmental and managerial conditions. A more sound comparison would be based on restricting the sample to fields that experience both a rotating and non-rotating sequences (denoted *Mixed*). In the Mixed subsample, comparing Average yields in Rotation to average yields in non-rotating sequences leads now to a positive estimate for corn as we as for soybeans. Interestingly, the effect shrinks dramatically for soybeans.

What the analysis here suggests is that there appears to be an important selection effect in the choice of the rotation sequence, where highly productive fields tend to be planted to corn only, and lower productivity fields see more rotation. For soybeans, fields with soybeans only have the lowest average yields. Of course, by merely observing that fields with soybeans only have the lowest yields, we cannot distinguish whether this is caused by a location effect (low fertility fields are planted to soy) or by the rotation effect (doing soy only leads to very low yields).

To shed more light on the relationship between productivity and crop choice, Figure 5 shows the distribution of the fields average yields, over the subsample of fields that did exclusively corn or

Figure 5: Distribution of average corn and soybean yields for each field

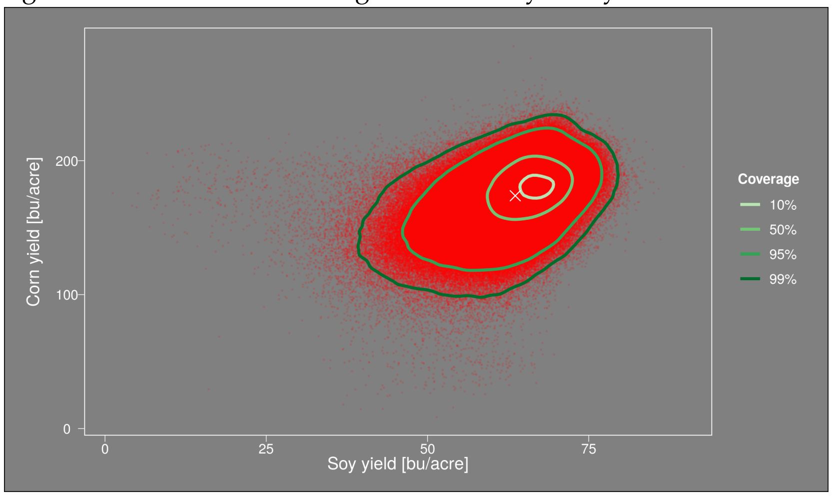

Note: data shows the 2D density estimation of the distribution of corn and soy means for each field.

soybeans. What is shown is the average corn yield, and average soy yields, for all fields that planted at least once corn and once soy. There is a clear positive correlation between the yields of the two crops: a field that is good for corn is likely to be also good for soybeans. There are some small pockets with very low corn yet high soy yields (or the converse), but these are likely to be due to measurement error (most of them have only one year planted to corn or soy, leading to very imprecise means). The correlation coefficient between corn and soy yields is 44%.

Figure **??** shows for the same dataset the number of years planted with corn (left panel) and number of years planted in rotations (right panel). The left panel shows that *good fields* then to be planted more often to corn than to soy. Conversely, fields with both low corn and low soy yields (on the bottom left) tend to be more often planted to soybeans. Interestingly, the highest yielding corn fields are not necessarily the ones planted very often to corn,[2] but have rather close to 4-5 years of corn. Highest soy yields on the other side are reached on fields that are most often planted to corn. This is probably due to the fact that the three years sequences C-C-S are the most beneficial for soybeans. Turning now to the number of CS/SC rotations in the right panel, we see that the highest corn yields are reached with the largest number of CS/SC rotations. On the other side, highest soy yields are obtained with only a few rotations, which suggests again that C-C-S rotations are more beneficial than the S-C-S ones.

As discussed above, we would expect that fields having a clear advantage for one crop versus the other (i.e. either lying on the sides of the distribution) would have a higher tendency to cultivate one crop versus the other. There is weak evidence that this is the case when looking at the right panel of Figure **??**. We see that regions in the center tend to have more rotations than those on the outside. The pattern is however not very regular, with some important exceptions on the north-east region.

We move now to investigate in more detail the relationship between number of years planted to corn and yields. Figure 7 shows the distribution of (average field) yields distinguishing whether a field was planted once, twice etc to corn. Fields planted more often to corn have higher yields than those planted only fewer times to corn. This goes against the intuition that doing more rotation increases the yields, and is probably explained by the non-random location effect. Interestingly, we see that the same pattern happens with soybeans, whose yields are also increasing with fields that are more often planted to corn. This on the other side is explained both by the field selection effect, as well as the rotation effect, as the more often corn is planted, the more likely soybeans will be on a

[2]Remember we do not observe the always planted to corn on this graph as they would not have a soy value.

Figure 6: Average number of corn and rotation per cell

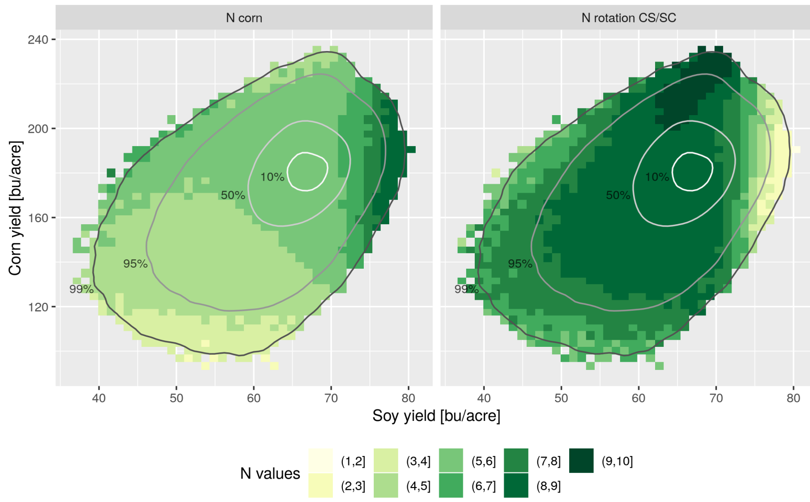

Note: data represents a simple binning estimator on the grid of corn and soy yields. Contours show the distribution of the data, the line with 99% indicating 99% of the data is within this contour.

beneficial C-S or C-C-S rotation.

Among fields planted the same number of years to corn, there is variation in the number of years they were planted in rotation. As an example a field planted 5 years to corn could be planted as C-C-C-C-C-S-S-S-S-S (0 rotations) or C-S-C-S-C-S-C-S-C-S (5 rotations).[3] Figure 8 disaggregates further the yields, calculating average yields separately depending on the number of rotations and the number of years planted to corn. We see that now, controlling for the number of years planted to corn, fields with more rotations have higher yields. The rotation difference shows a slight tendency to decrease for fields with higher yields (higher number of years planted to corn). Turning to soybeans in Figure 9, the difference in yields by number of rotations is even clearer. Except for a single case (4 years planted to corn) doing more rotations is always and strongly associated to higher yields. Interestingly, we see that unlike corn the effect seems to increase with higher yields/higher number of years planted to corn.

This section highlighted a few stylised facts on rotation choice and field production in the US Corn Belt. To summarise these finding, I showed that, in the subset of fields doing exclusively corn or soybeans: 1) There is a large proportion of fields doing always $\langle CS \rangle$ rotation : 46% of the fields did always $\langle CS \rangle$ over ten years. 2) Fields doing always corn $\langle C \rangle$ have a higher yield than fields doing always rotation $\langle CS \rangle$. On the other side, fields doing always soybeans $\langle S \rangle$ have lower yields than their always rotating counterpart $\langle CS \rangle$. 3) The previous finding suggests that better fields are used more often to corn, while lower quality ones are planted more often to soybeans. This claim is corroborated by two further findings. 4) First, looking at fields by years planted to corn, we see a clear increase in corn yields, but also an increase in soy yields. 5) Second, looking at the correlation between averages of corn and soy yields for each field, we find a positive correlation (44%). Findings 4) and 5) seem to support the model by Hendricks et al. (2014b) of a one-dimensional field quality gradient, with low fertility fields doing more soy, medium fertility ones doing rotation and and high fertility ones doing more corn. 6) Finally, I show that controlling for the number of years planted to corn, sequences with more rotations are associated with higher yields.

These stylised facts indicate that there are important location effects in the choice of rotating or not, which raise important challenges for estimation. A further issue with the analysis made above

[3] In fact, there are many more possible configurations, see Figure 20 in the Appendix for an illustration.

Figure 7: Yields by number of years planted to corn

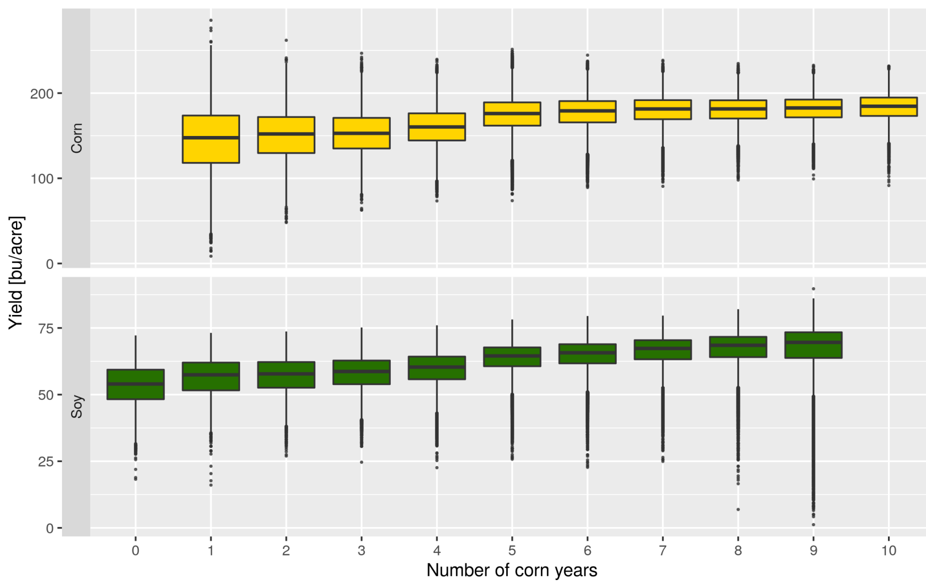

Figure 8: Yields by number of corn years and number of rotations

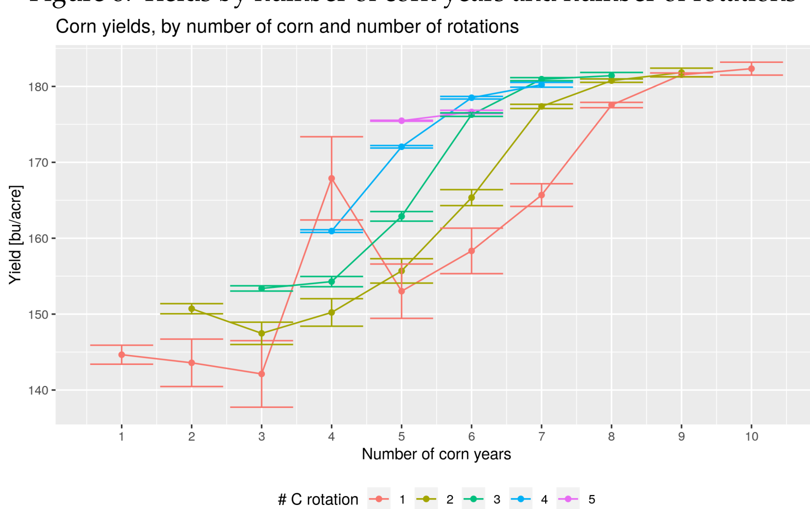

Figure 9: Soy yields by number of corn years and number of rotations

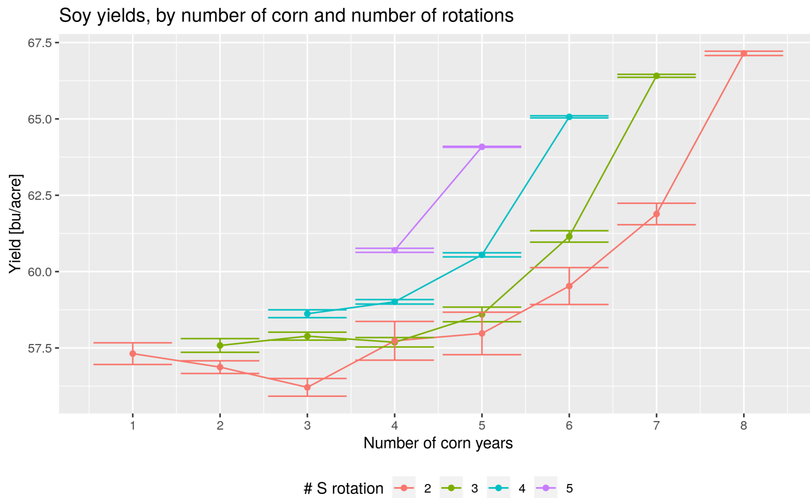

is that it does not take into account state specificities, or difference in years and weather. The next section discusses formally both points, investigating how to estimate rotation effects when there are important location effects.

3 Estimation

In this section, I discuss briefly how to conceptualize the effects of rotation in Section 3.1, then define the parameter of interest in Section 3.2, and finally discuss the various way to estimate the parameter in Section 3.3.

3.1 Conceptual model of rotation

Rotation effects in production functions have been modelled in various ways, with the main differences being in the way the rotation effects are taken into account, and in the way dynamics are introduced. An early strand of literature used mathematical programming to derive optimal rules in a static framework, see El-Nazer and McCarl (1986) or Musser et al. (1985). Dynamic programming methods based on Bellman equations have been used, see Thomas (2003) on the general case of crop choice in presence of nitrogen carry-over, or Livingston et al. (2015); MacEwan and Howitt (2011) specifically on rotations. Although interesting, these methods have the drawback that they do not lead to closed-form estimators. On the other side, the framework of Hennessy (2006) provides a clear modelling framework of rotation effects amenable to direct estimation, that I adopt here.

Hennessy (2006) considers two effects of rotation, the *input saving* effect α and the *yield boost* effect β. The input saving effect α arises from nutrient carry-over from the previous period(s), and is assumed to be perfectly substitutable with chemical fertiliser. Total amount of nutrient N for crop is equal to the sum of chemical fertiliser F and the input-saving effect α, $N = F + \alpha$. The input-saving effect α depends on the type of crop succession, which we will write as α^i_j, i.e. when crop i follows crop j. It is assumed that in absence of rotation, this effect is zero, $\alpha^i_j = 0 \, \forall j = i$. The second effect of crop rotation is the yield *boost effect* β^i_j (for crop i following crop j), which is assumed to enter additively. These two elements lead to the following yield production function for crop i:

$$Y^i(F|j) = y^i(F + \alpha^i_j) + \beta^i_j$$

Embedded in this representation is the assumption that both the input-saving α^i_j and yield boost β^i_j effects do not depend on previous level of nutrient N_{t-1} or on actual level of fertiliser F_t. While this restrictive assumption departs from the nitrogen carry-over literature[4], it has the advantage of alleviating the need for dynamic programming tools and to model the crop choice decision in a more intuitive way. An important implication of the perfect substitutability assumption between input saving α and chemical fertiliser F is that the optimal fertiliser level with rotation $\tilde{\tilde{F}}$ is simply the optimal level without rotation $\tilde{F}$ minus the input saving α, $\tilde{\tilde{F}} = \tilde{F} - \alpha$. As a consequence, the total amount of nutrient N is the same, and the difference in yields is simply β, i.e. $y^{\langle SC \rangle} - y^{\langle C \rangle} = \beta^C_S$. A second consequence of the perfect substituability between chemical fertilizer and input-saving is that profit with rotation $\tilde{\tilde{F}}^i_j$ is equal to the profit without rotation $\tilde{\pi}^i_j$ plus yield boost β plus efficiency saving α: $\tilde{\tilde{\pi}}^i_{j \neq i} = \tilde{\pi}^i_i + p^i \beta^i_j + w\alpha^i_j$, with p^i the price of crop i and w the price of fertilizer.

Analyzing the farmer's crop choice and rotation sequence decision is particularly easy under the perfect substitutability assumption. To simplify, let us assume a world with two crops, corn and soybeans, and a stationary environment with fixed prices p^C, p^S and w. Under a one-year memory, the farmer needs to compare only three crop-rotation combinations $\langle C \rangle$, $\langle S \rangle$ and $\langle CS \rangle$, over a two-year period:

- $\pi(\langle C \rangle) = p^C \tilde{y}^C - w\tilde{F}^C$
- $\pi(\langle S \rangle) = p^S \tilde{y}^S - w\tilde{F}^S$
- $\pi(\langle SC \rangle) = \frac{1}{2}\left(\pi(\langle C \rangle) + p^C \beta^C_S + w\alpha^C_S\right) + \frac{1}{2}\left(\pi(\langle S \rangle) + p^S \beta^S_C + w\alpha^S_C\right)$

Without rotation, the farmer will choose C over S if $\pi(C) > \pi(S)$. Let us assume for now that $\pi(C) > \pi(S)$ holds in general. With rotation, the choice is now between $\langle C \rangle$ and $\langle CS \rangle$, and the condition to choose $\langle C \rangle$ over $\langle CS \rangle$ becomes: $\pi(\langle C \rangle) > \pi(\langle S \rangle) + p^C \beta^C_S + w\alpha^C_S + p^S \beta^S_C + w\alpha^S_C$. That is, corn profit needs not only be higher than soy profit, but also higher than soy profit and the various rotation gains (evaluated at market prices), β^C_S, β^S_C, α^C_S and α^S_C. The condition to choose $\langle C \rangle$ over $\langle CS \rangle$ then becomes:

$$\tilde{y}^C > \frac{p^S}{p^C}\tilde{y}^S + \frac{1}{p^C} w(\tilde{F}^C - \tilde{F}^S) + \beta^C_S + \frac{w}{p^C}(\alpha^C_S + \alpha^S_C) + \frac{p^S}{p^C}\beta^S_C \tag{1}$$

Symmetrically, choosing $\langle S \rangle$ over $\langle CS \rangle$ will require:

$$\tilde{y}^C < \frac{p^S}{p^C}\tilde{y}^S + \frac{1}{p^C} w(\tilde{F}^C - \tilde{F}^S) - \left(\beta^C_S + \frac{w}{p^C}(\alpha^C_S + \alpha^S_C) + \frac{p^S}{p^C}\beta^S_C\right)$$

Figure 10 illustrates the choice of decision in the corn and soy yield space. The data is from the 2018 crop budget for Central Illinois (Schnitkey, 2018) which reports prices, yields and fertiliser costs with and without rotation. The straight line represents the decision boundary in absence of rotation comparing $\langle C \rangle$ over $\langle S \rangle$, that is the line $\tilde{y}^C = \frac{p^S}{p^C}\tilde{y}^S + \frac{1}{p^C} w(\tilde{F}^C - \tilde{F}^S)$. Assuming that the farmer knows the corn and soy yields of her field, she will choose $\langle C \rangle$ is the point is above the line, or $\langle S \rangle$ if below. The two black points correspond to two fields described in the report, one described as low productivity and the other one as high. Without rotation, this farmer would set up for $\langle C \rangle$. The dashed line represent choice under rotation, taking into account the intercept shift from the rotation benefits $\beta^C_S + \frac{w}{p^C}(\alpha^C_S + \alpha^S_C) + \frac{p^S}{p^C}\beta^S_C$ as in Equation 1. This illustrates the point before that profit from $\langle C \rangle$ needs not only to be higher than profit from $\langle S \rangle$, but also higher than the profit from $\langle CS \rangle$. In this case, the yield-boost effect of 5% (shown by the arrow) and input-saving effect of 8% (only for corn) are strong enough to put both fields within the $\langle CS \rangle$ zone.

[4] Thomas (2003) uses for example a specification similar to $\alpha^i_j = m^j(n_{t-1})$.

Figure 10: Choice of rotating sequence

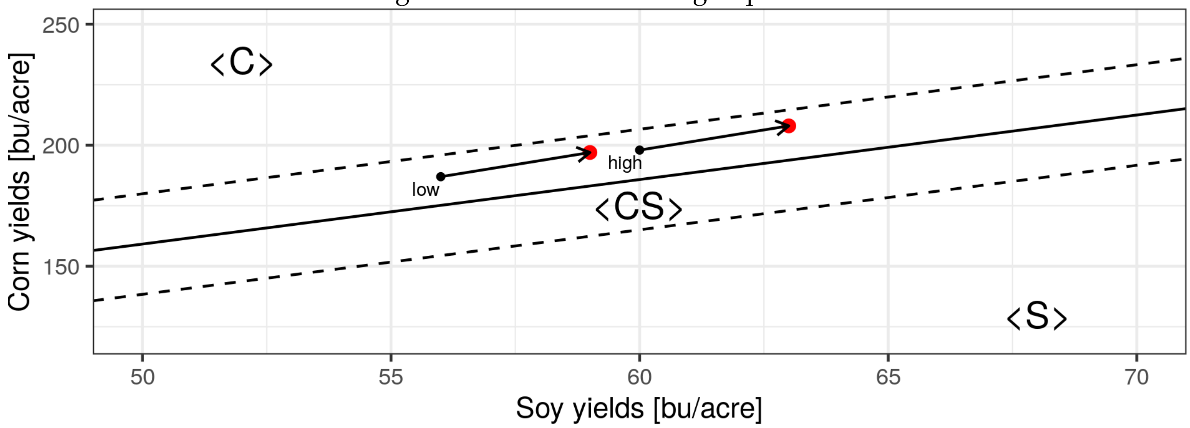

3.2 Parameter of interest

Estimating the rotation effect raises an interesting econometric challenge. For one, the preliminary analysis in the previous section shows that there are important selection effects in the fields choosing rotating sequences. Second, the rotation effect raises a fundamental problem of missing data, where observing a rotation implies that we are not observing yields on the preceding year. As a consequence, standard identification strategies such as the difference-in-difference are not directly applicable.

What is the rotation effect? Or more precisely, which rotation effect are we interested in? Let us denote by $y(N, c_{t-1}, M)$ the yield response function using nitrogen level N, previous crop choice c_{t-1}, and management practice M. In the experimental field approach, M and N are fully controlled, allowing to estimate the conditional average rotation effect (CARE):

$$CARE(n, m) = E[y(N, 1, M) - y(N, 0, M) | N = n, M = m]$$

This is the effect shown in Figure 3, comparing yields at every level of fertilizer. In practice, there is more interest in reporting the effect at the so-called economic optimal nitrogen rate (EONR), which is the point on the yield curve where the slope is equal to the ratio of output to input prices p^{out}/p^{in}. Owing to the input-saving effect, the EONR is different for rotating and non-rotating fields, so that what I call here the conditional average rotation effect based on price CAREP reads:

$$\begin{aligned} CAREP &= CARE(m, p^{out}, p^{in}) \\ &= E[y(EONR_1(p^{out}/p^{in}), 1, M) - y((EONR_0(p^{out}/p^{in}), 0, M) | M = m] \end{aligned}$$

In the simple model of Section 3.1, the CAREP corresponds to the yield-boost effect β, while the difference in $EONR_1 - ENOR_0$ corresponds to the input-saving effect α.

With the dataset we have, we do not observe fertilizer application, nor management practices, though we still observe output prices. We hence cannot condition on N or M, implicitly averaging over them. The quantity we can identify shall be simply called the average rotation effect conditional on prices (AREP):

$$AREP(p^{out}, p^{in}) \equiv E_{n_1, n_0, m} \left[E[y(1, n_1(p^{out}/p^{in}), m) - y(0, n_0(p^{out}/p^{in}), m)] \right]$$

Under the assumption that farmers maximize profit and adjust their fertilizer level to the EONR, this estimate recover also the yield-boost β. The main difference between the experimental CAREP and the observational AREP is that CAREP controls the field/management variables as well as fertilizer level, while the AREP averages over field and management variables.

3.3 Identification strategies

Studies based on experimental plot data use a simple *between* plots strategy, either running t-tests or linear regression controlling for various covariates. This is justified when plots are sufficiently similar to be compared, which is at the discretion of the experimenter. With observational data, this approach may not hold anymore. As the simple analysis in Section 2 showed, dependence between the rotation choice and potential yield can bias the comparison.

The traditional way to address the fact that fields with different qualities might opt for different cropping patterns is to use field-level fixed effects α_i:

$$y_{it}^C = \alpha_i + \beta D_{it} + \varepsilon_{it} \tag{2}$$

where D_{it} is the treatment variable of interest, in our case whether or not there was rotation. With fixed effects, the identifying variation is the *within-plot* variation in and out of rotation. That is, we are only considering *switching* fields that do at least one monocropping sequence and one rotating sequence (the *Mixed* category shown in Section 2). The causal effect corresponds to the average treatment on the treated (ATT):

$$ARET = E[y(1) - y(0)|switcher]$$

A possible threat for the pure within-plot estimation is that one does not control for year differences. A situation where "good" plots are rotated during good years while "bad" plots are kept under monocropping would bias our estimate of the rotation effect. One way to address this is to control for observable variables, such as weather. The more common solution is to add year fixed effects λ_t in a two-way fixed effects model:

$$y_{it}^C = \alpha_i + \lambda_t + \beta D_{it} + \varepsilon_{it}$$

Several recent paper discuss the causal interpretation of the two-way fixed effects model (Imai and Kim, 2019; Goodman-Bacon, 2018; de Chaisemartin and D'Haultfoeuille, 2019). It is well known that with just two time periods, the two-way fixed effects estimator is numerically equivalent to the difference in difference (DiD) estimator (see Angrist and Pischke, 2008). Under the parallel trend assumption, the DiD estimator recovers the treatment on the treated at the second time period, i.e. on the group that experienced both a treated and non-treated period. Let D_{it} denote the treatment variable (taken here in a general context, the specific case of rotation will be discussed later on), with $D = 1$ indicating treatment, $D = 0$ otherwise. The parallel assumption states that the treated group ($D_{t_1} = 0$ and $D_{t_2} = 1$) would have followed the same trend in output as the control group ($D_{t_1} = 0$ and $D_{t_2} = 0$), if it had not been treated in the second period: ${}^{\text{0-}\tilde{0}}/\!\!/_{\text{0-0}} \equiv E\left[\Delta Y_t(0)|0-1\right] = E\left[\Delta Y_t(0)|0-0\right]$. The ${}^{\text{0-}\tilde{0}}/\!\!/_{\text{0-0}}$ notation indicates that $0 - \tilde{0}$ is the counterfactual trend of the 0-1 group. Under this assumption, the DID compares the difference in trend between in-switchers ($D_{i1} = 0$ and $D_{i2} = 1$) and non-treated ($D_{i1} = 0$ and $D_{i2} = 0$) outcomes, and will be denoted henceforth DiD_{00}^{01}.

With two time periods, we can in some situations observe two more groups, the *out-switchers* ($D_{i1} = 1$ and $D_{i2} = 0$) and *fully treated* ($D_{i1} = 1$ and $D_{i2} = 1$). This opens the possibility to use up to four different DiD estimators. For one, we can estimate the ATT in period 2 for the out-switchers (1-0), comparing them to the fully treated (1-1). Under the parallel trend assumption that out-switchers would have followed the same trend as fully-treated had they not been untreated in period 2 (${}^{\text{1-}\tilde{1}}/\!\!/_{\text{1-1}}$), the DiD_{11}^{10} recovers the same ATT, but for the out-switcher group. Going further, we can also compare the in-switcher (0-1) to the fully treated (1-1). The DiD_{11}^{01} will recover the ATT at period 1, if one is willing to assume now that the 0-1 would have followed the same trend as the 1-1, had they been treated in period 1 (${}^{\tilde{1}\text{-1}}/\!\!/_{\text{1-1}}$). Table 2 summarises the different estimators and the identifying assumption required to recover a causal effect. All of these estimators can be implemented either using a difference in means, or using a two-way fixed effects model.

Table 2: Causal effects with multiple groups

Group	Estimator	Control group	Assumption	Causal effect
in-switchers (0-1)	DiD^{01}_{00}	0-0	${}^{0\text{-}\tilde{0}}/\!/_{0\text{-}0}$	$ATT(0\text{-}1, T=2)$
	DiD^{01}_{11}	1-1	${}^{\tilde{1}\text{-}1}/\!/_{1\text{-}1}$	$ATT(0\text{-}1, T=1)$
out-switchers (1-0)	DiD^{10}_{11}	1-1	${}^{1\text{-}\tilde{1}}/\!/_{1\text{-}1}$	$ATT(1\text{-}0, T=2)$
	DiD^{10}_{00}	0-0	${}^{\tilde{0}\text{-}0}/\!/_{0\text{-}0}$	$ATT(1\text{-}0, T=1)$

When there are multiple periods, what is a two-way FE β^{FE_2}estimating? Or, more importantly, what should it estimate? A possible causal parameter of interest would be the average of ATT over groups and time periods, i.e. $ATT \equiv \sum_g \sum_t \frac{N_{gt}}{N} ATT(G = g, T = t)$ where N_{gt} refers to the number of treated units in group g at time t, and N is the grand sum.

de Chaisemartin and D'Haultfoeuille (2019) show that the two-way FE estimates a weighted average of ATTs, but with weights different from the population weights N_{gt}/N. Worse, the weights can be negative, so that the sign of the β^{FE_2} can be even different from the sign of the general ATT. Imai and Kim (2019) suggest a multi-period estimator, estimating an ATT for each two-year period, and averaging over the periods. They consider however only the DiD^{01}_{00} estimator, so that their estimator does not cover all the switching units.[5] de Chaisemartin and D'Haultfoeuille (2019) suggest a similar approach, including this time also DiD^{10}_{11}. Both these papers require the parallel trend assumptions to hold for each sub-period. Goodman-Bacon (2018) provides maybe the most comprehensive treatment, showing that the β^{FE_2} is an average of the multi-period ATT and the group trends. This clarifies the content of the β^{FE_2} in case of a failure of the parallel trends assumption.

In this paper, I follow the approach of de Chaisemartin and D'Haultfoeuille (2019), adapting their multi-period estimator for the specific case of rotation effects. Estimating the rotation effect raises the additional challenge that we never get to observe an untreated-treated (0-1) sequence directly. Indeed, by its very nature, observing a rotation means that the previous crop was not the same. It is hence not possible that the previous period $t-1$ was under monocropping. On the other side, we can observe directly a treated-untreated sequence (1-0). To see this, let us write crop choice as c_{it}, with $c_{it} = 1$ for corn and 0 for soybeans. Let D_{it} be now the rotation treatment dummy, with $D_{it} = 1$ indicating rotation: $D_{it} = \mathbb{1}(c_{it} \neq c_{it-1})$. In the standard difference-in-difference framework, we are comparing outcomes from 0-1 treatment units with outcomes from 0-0 control ones. This cannot happen with rotations: having a rotation implies that previous's year crop was different, so that we do not observe the previous outcome. What we observe is rather a $0 - \emptyset - 1$ sequence, the $\emptyset$ value in the middle indicating that soy was planted that year. This implies that to apply the DiD idea we need to take (at least) a two years difference, comparing differences in sequences with and without rotation at period 2 and 4. In other words, we are comparing $\overline{C}C\overline{S}C$ to $\overline{C}C\overline{C}C$. It should be noted that while the first sequence contains 3 C values and the second 4, we are only using two of them, the second and the fourth. The parallel trend assumption becomes here the assumption that the (counterfactual) trend in absence of rotation at time 3 would be the same as the observed trend for never treated (CCCC).

The CCSC sequence is not the only sequence that experiences *within* variation. For the out-switcher category (1-0 in standard DID), the situation is even more complicated. For one, there are two different cropping sequences with out-switching (i.e. rotation preceding monocropping), $\overline{S}C\overline{C}C$ or $CSCC$, corresponding to corn rotation sequences 1-0-0 and $\emptyset$ -1-0. Furthermore, for the first one, we can estimate up to three different ATT (at period 2, 3 or 4) while for the second one we can only estimate one (at period 3). Most of the sequences require either the always or never takers as control group, although for some of them the control group turns out to be a switcher group.

[5] This would make sens in an event-study framework though, as no 1-0 is observed.

Table 3: Corn cropping sequences and identification

Sequence	Share	Rotations	Description	Switching?	Control	Periods considered	Period identified
SCSC	61%	1-∅-1	Always rotate				
CCSC	9%	0-∅-1	in-switcher at 4	✓	0-0-0 (never rotate CCCC)	2 and 4	4
					1 ∅ 1 (always rotate SCSC)	2 and 4	2
SSSC	1.3%	∅-∅-1	only one				
CSSC	3%						
CSCC	8%	∅-1-0	out-switcher at 3	✓	∅-0-0 (CCCC)	3 and 4	3
SSCC	0.4%						
SCCC	4%	1-0-0	out-switch at 3 or 4	✓	1-∅-1 (always rotate SCSC)	2 and 4	4
					0 - 0 - 0 never rotate CCCC	2 and 4	2
					∅-1-0	3 and 4	3
CCCC	13%	0-0-0	Never rotate				

1 Note: Column *Sequence* indicates the 4 years sequence. Column *Rotations* indicate whether a rotation is done in eachof the last three period. 1 stands for rotation, 0 stands for no rotation, and ∅ stands for not observed. Column *Control* indicate the control sequence, and the last two columns indicate which periods in the sequence are considered for the difference-in-difference, and for which period the ATE is identified.

Table 3 shows all possible sequences with four periods for corn.[6] The table shows the sequences, and their notation in terms of rotation. As an example, the sequence S-C–S-C (always rotate) reads 1-∅-1 in terms of rotation: the second value in the sequence was in rotation (SC), the third is unobserved for corn (CS, written as ∅), and the fourth is rotation again (SC). The table indicates which are the switching sequences, i.e. the sequences that have at least once a rotation SC and once a monocropping sequence CC. These are the sequences that exhibit within variation and used to identify the ATT. The column control indicates for the switching sequences which sequences can be used as control in the diff-diff estimator. The columns *period considered* indicate which period will be used to form the DiD estimator, while column *period identified* indicates at which time the ATT is identified. As an example, for the CCSC sequence, —coded as 0-∅-1 and denoted as *in-switcher*— we can either use *never-rotaters* to identify the effect at time four, or *always-rotaters* to identify the effect at time 2. Going back to the general discussion about difference-in-difference estimation with two periods and four groups, this case corresponds to the first two rows of Table 2, where DiD^{01}_{00} identifies $ATT(0\text{-}1, t_2)$ while DiD^{01}_{11} identifies $ATT(0\text{-}1, t_1)$. The only difference here is that for rotations, we need to take the difference between time 2 and 4 in a four sequence period, while for the standard DiD we use a period of two years and take a simple difference. Table 3 reveals more exotic situations in the rotation context. For one, the out-switchers (i.e. experiencing first 1 then 0) can be found in three different sequences. Two of the sequences are equivalent under one-year memory. The DiD here would be between period 3 and 4, identifying the effect at period 3. For the last sequence, S-C-C-C, we can identify up to three different ATT.

If more than four time periods are available, the same analysis is done on a rolling basis. This implies that for each year, we can have multiple estimators for the same quantity. This is not specific to the rotation framework: as we saw for the DiD with two periods and multiple groups, we can estimate ATT at each period. Adding one more year adds two more identification for the groups that switch in both period. This leaves up to 4 estimates for three periods. Let us take the group 0-1-0 as example. We can estimate its ATT(T=2) either with $DiD^{01\cdot}_{00\cdot}$ (use periods 1 and 2), or with $DiD^{\cdot 10}_{\cdot 11}$ (use periods 2 and 3). This suggests that we have over-identifying restrictions, that can serve as a test for time homogeneity.

To be valid, the DiD discussed above all require a specific version of the parallel trend hypothesis. For the CCSC sequence using CCCC as control, we are assuming that $E[C_4 - C_2 | CCSC] = E[C_4 -$

[6] The result is symmetric for soybeans.

$C_2|CCCC]$, where $\tilde{C}_4$ indicates the counterfactual value of C_4 had it been preceded by corn (i.e. had it not had the rotation treatment) instead of soybeans. The parallel trend assumption is an identifying assumption that cannot be tested. In practice, it is customary to run *placebo* tests, that test for parallel trend in situations where treatment and control units have the same treatment status. In the standard DiD framework, this require using pre-intervention years, where both groups have treatment status 0. In the rotation case, we do not necessarily need more data: it turns out that for the specific CCSC sequence, we can actually build a placebo test from the sequence itself, without additional data. Indeed, we are comparing CCSC to CCCC, yet are only considering values at time 2 (both CC) and 4 (SC versus CC). The first sequence CC is common to both, and hence can be used as a pre-trend test.

To summarize this section, I showed here methods to estimate rotation effects with observational data. A key element to take into account with observational data is the non-randomness of crop choice and field characteristics. I showed here the paradox that higher fertility soils are used more often for corn than soybeans. As a result, fields doing always rotation turn out to have a higher average yields compared to those doing always rotation. A first approach to address this issue is to use plot-level fixed effects. This will control for field characteristic, but identification can be threatened if there is correlation between rotation choice and specific good/bad years. A usual way to address this second concern is to add year fixed effects. Resulting two-way fixed effect estimators can be interpreted as doing an average of difference-in-difference estimators. In the rotation case however, the interpretation is more complicated, as rotating a field implies that last year's yield was not observed. I show how to extend the standard DiD in this case, and suggest placebo tests. In the next section, I proceed to the estimation, and discuss the results.

4 Results

In this section, I proceed to the analysis itself of the various estimators suggested above. I start in Subsection 4.1 by using the plot fixed-effects estimators. I move then in Subsection 4.2 to the two-way fixed effects and the multi-year estimator.

Throughout this section, the analysis is made using the subset of fields that never cultivated something else than corn or soybeans over the 2008–2017 period. Out of our initial sample of fields doing at least once corn or soybeans (1.6 mio fields), this leaves us with a sample of close to 800'000 fields.[7] Among these 800'000 fields, most are in Illinois (27%), Iowa (25%), Indiana (16%) and Minnesota (12%).

4.1 Plot fixed-effect analysis

I start here by showing the results from a panel approach using plot-level fixed effects. Table 4 shows the coefficient β^{FE_1} from the fixed-effect panel estimator, versus a model without the α_i (*pooled* estimators):

$$y_{it}^C = \alpha_i + \beta D_{it} + \varepsilon_{it}$$

D_{it} is the rotation dummy, indicating whether the field was in rotation ($c_{i,t} \neq c_{i,t-1}$) or not. We estimate two different pooled estimators. The first one—*Pool switchers*—is using the sample of switcher fields that had at least once a rotation and once no rotation. These are the fields that are used in the FE1 estimation. The second pool estimator—*Pool all*—, uses all fields, including the always rotaters and always corn. The regression confirms the intuition we obtained from Figure 4: doing a within or between/pooled analysis has dramatic consequences on the estimates. We see that when not controlling for anything (*Pool all*), we obtain a negative rotation value for corn and strong positive value for soy. This arises from the fact that better fields are more often planted to corn, and worse fields

[7]The total number of fields which did always either corn and/or soybeans in our sample is actually higher, but not all fields have yields estimates from SCYM for every year, mainly due to cloud issues.

Table 4: Simple FE1 and pooling estimation

	Corn			Soy		
	FE1	Pool switchers	Pool all	FE1	Pool switchers	Pool all
Rotation	10.57^{***}	5.68^{***}	-0.81^{***}	0.51^{***}	1.73^{***}	6.24^{***}
	(0.06)	(0.06)	(0.04)	(0.03)	(0.03)	(0.02)
Num. obs.	1838471	1838471	4239420	899550	899550	3546032
Num N obs	287324			150854		
Num T obs (ave)	6.40			5.96		
N. variables	1	2	2	1	2	2
Mean dep	177.24	177.24	175.33	59.14	59.14	62.94

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Standard errors clustered at the state-year level.

Table 5: Within estimator with weather covariates

	Corn FE1	Corn FE1 weather	Corn weather	Soy FE1	Soy FE1 weather	Soy weather
Rotation	10.57^{***}	6.20^{***}	1.86^{***}	0.51^{***}	1.00^{***}	2.71^{***}
	(0.06)	(0.36)	(0.50)	(0.03)	(0.19)	(0.17)
Num. obs.	1838471	3716937	3716937	899550	2463737	2463737
Num N obs	287324	709650		150854	490297	
Num T obs (ave)	6.40	5.24		5.96	5.02	
N. variables	1	90	91	1	90	91
Mean dep	177.24	171.06	171.06	59.14	58.23	58.23

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

are planted more often to soy. In other terms, we are comparing different field qualities rather than assessing the rotation effect. Estimating the pool estimator on the subset of switching fields leads to a higher coefficient for corn, and a lower one for soybean, removing part of the location bias. Finally, using the fixed-effect estimator accentuates the sign reversal, giving even higher values for corn and even lower for soybeans. Surprisingly, soy has now a quite small estimated rotation effect, at 0.5 [bu/acres], which seems to go against findings in the literature suggesting that rotation is particularly beneficial for soybeans.[8]

The next table, Table 5, shows how the fixed-effect and pool estimators change when weather covariates are included. The set of weather covariate is a very rich set of minimum and maximum temperatures, precipitation and vapour pressure deficit. I use monthly values, over eleven months, and include squared terms (see Section **??** in Chapter II). Adding weather controls does have an effect for both corn and soybean. It reduces importantly the corn coefficient, while somehow increases the soy one. The fact that adding weather covariates changes the FE1 estimator suggests that rotation choice is not totally random across years. The last column for each crop (*corn weather* and *soy weather*) is the pooled estimator, this time controlling for weather. The surprising result that the pool estimator indicates a negative effect for corn is now mitigated. Estimates are however quite different from any of the FE ones, suggesting that controlling only for weather, as is done in some studies, is not enough to control for the location bias.

As a final check, I investigate whether heterogeneity is important, and whether it is likely to change our average estimates. Gibbons et al. (2019) show when there is heterogeneity in the β_i, fixed-effect estimators neglecting the heterogeneity do not necessarily deliver an estimate close to the ATE. To test whether this might happen in my sample, I estimate slope parameters at the county level, as well as at the state level. Averaging these leads to results that are very close to the estimate assuming

[8]Porter et al. (1997) reports 13% yield benefit on experimental fields, and Seifert et al. (2017) find 10% based on observational fields. Schnitkey (2018) documents as well effects of about 3 [bu/acres], i.e. 5%, on fields in Illinois in 2018.

Figure 11: Rotation effect at the county level, by state

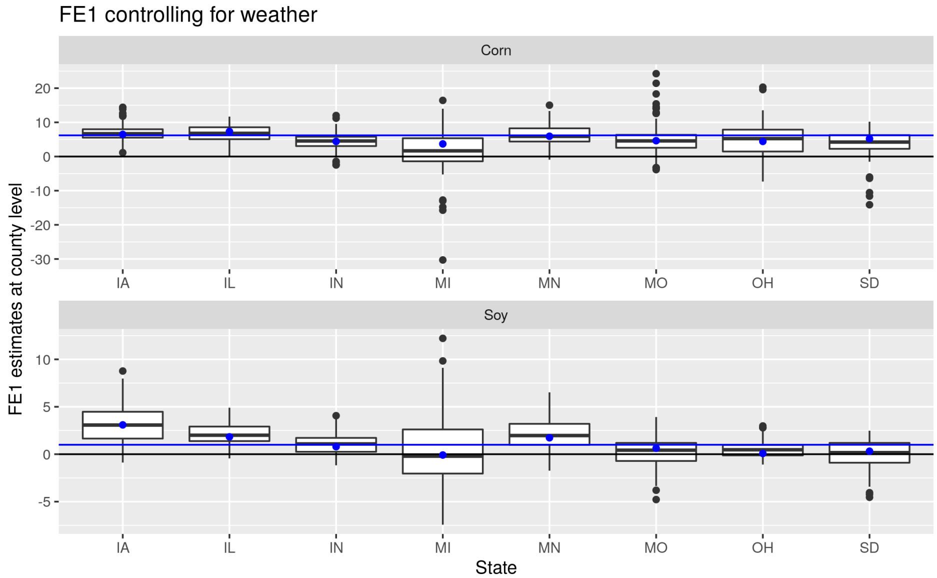

1 Note: Line in blue shows the value when estimating one single coefficient. Dots in blue show the average when estimating one coefficient per state.

full homogeneity. This suggests that the average *within* estimate is not subject to heterogeneity bias. This does not suggest however that there is no heterogeneity in the estimates. Figure 11 reveals difference among states. The figure shows the county-specific coefficients in a boxplot for each state. For corn, a few estimates appear negative, but there are either outliers, or from the state of Michigan (MI), which has very few fields cultivating exclusively corn and soybeans in my sample.[9] On the other side, for soybeans, the lower quartile of the county-specific coefficients are negative in a few states (MO and SD). This is a result that will deserve further investigation.

4.2 Plot and time fixed effects

I move now to the analysis adding time fixed effects besides the plot-level ones. The time fixed effects are first estimated globally, then by state, by agricultural zone (so-called MLRA), and finally by county.[10] Figure 12 shows the resulting coefficients, along with their standard errors. The first coefficient corresponds to the plot fixed effects β^{FE_1} discussed in the previous section. We see that adding time fixed effects has the same effect as controlling for weather: for corn, the coefficients are 7 or 6.3, very close to the value of 6.39 obtained in Table 5. Likewise, for soybeans, we see also an increase in the estimates with respect to the FE1 estimator. The change is however not very strong, and decays quickly as the time effects are estimated at finer scales.

4.3 Multiperiod DiD estimates

As argued above, causal interpretation of the two-way fixed-effects coefficients is not straightforward, and due to the weighting scheme of the two-way fixed effect, they might be far from the causal parameter of interest, the time- and group-average ATT. To investigate this, I implement the multi-DiD estimator discussed in 3.3. I focus on two specific treatment groups, the *in-switchers* SCCC and

[9]There are 4000 of these fields in our sample, among which only 2800 have variation in cropping/non-cropping patterns. The second smallest state, South Dakota, has ten times more fields both in total (42'000) as well as for fields exhibiting variation (21'567 fields).

[10]More precisely, I interact the year dummies with state or county indicators.

Figure 12: Coefficient froms one-way and two-way fixed-effect models

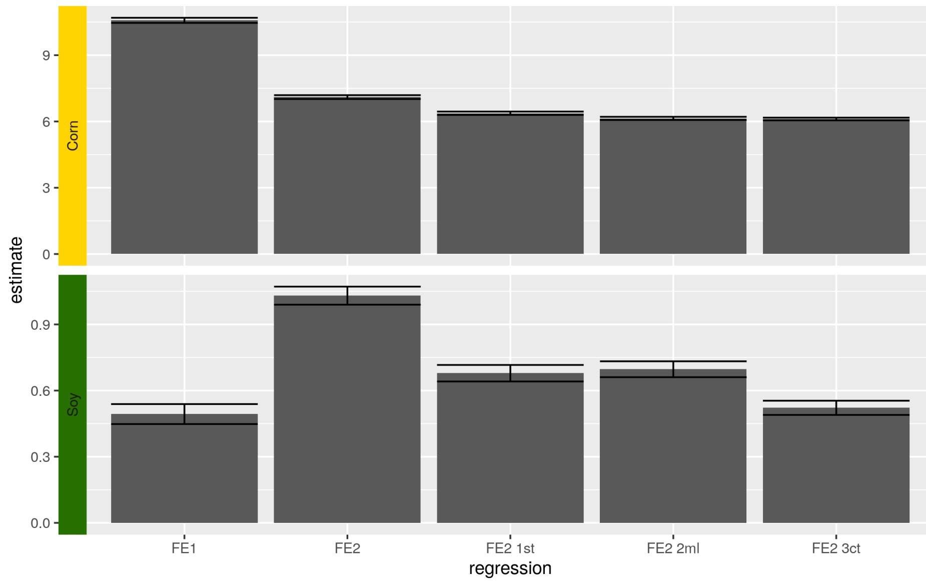

out-switchers CCSC. For these two groups, I seek to estimate the effect at the fourth period, rolling over periods of four years (refer to Table 3). I then investigate the over-identifying restrictions for a same group. Finally, I implement a placebo test for the *in-switchers* SCCC subgroup, comparing the difference in the two first CC values of $\overline{CC}SC$ to those of the control group, the $\overline{CC}CC$.

Figure 13 illustrates a specific sequence in 2008–2011 for in-switchers CCSC, compared to non-rotaters CCCC. The in-switchers, shown with a dashed line, are found to have lower yields than the always rotaters for the first two periods. This corresponds to the location bias, where fields of higher quality are more likely to be used intensively for corn. The DiD itself compares differences for each group between 2009 and 2011. We see that despite being lower in 2009, fields that did rotate in 2010 have now higher yields in 2011. Under the assumption of parallel trends, the difference can be attributed to the rotation effect.

The year by year DID estimates are shown in Figure 14. There are rather large variations over time for the estimates, ranging for corn from a value of 0 in 2016 to more than 15 in 2014. For soybean, some

Figure 13: Illustration of the DiD strategy, 2011

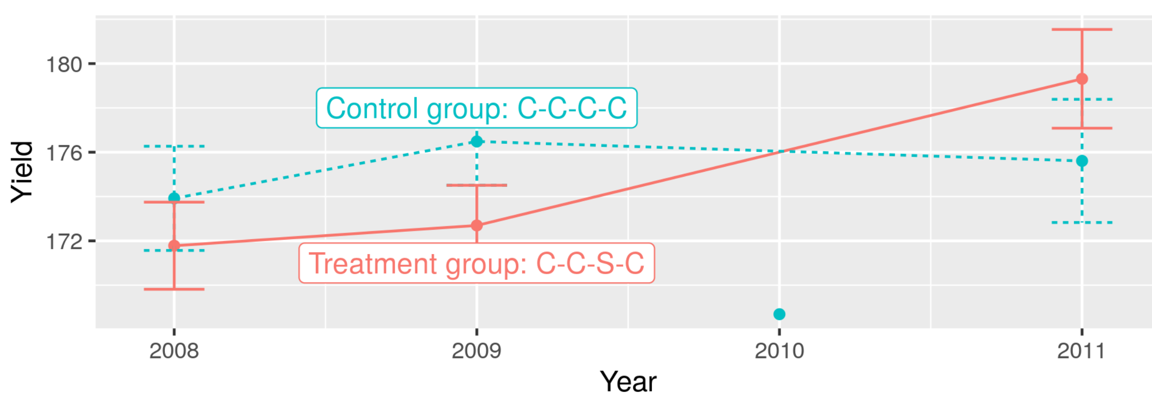

Figure 14: Multi-period DiD for rotation effect

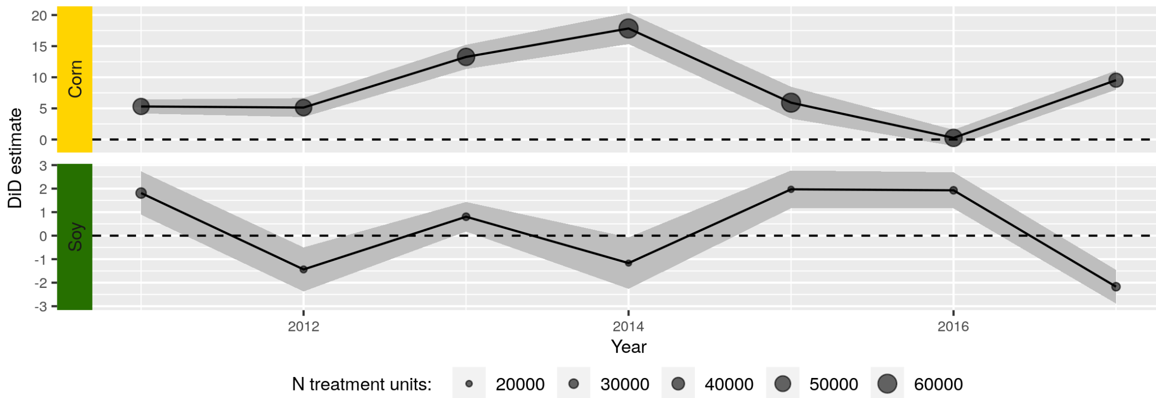

1 Note: points represent the DiD estimates for each subperiod. Confidence bands are based on standard errors clustered at the county-year level.

negative values appear for at least 3 years. The size of the points in the figure indicates the number of fields in the treatment group. This reveals that with the DID estimator, we are using a smaller part of the total sample. For soybeans, this is particularly small, as sequences with three of four soy values are not encountered often in practice. The sequences we use as control group represent 13% of the sample for corn, while they only represent 7% for soybeans.

Averaging the multi-year DiD estimators, I obtain a value of 8.31 for corn, and 0.281 for soybeans. Looking back at Figure 12 that compared β_{FE_1} and β_{FE_2} estimates, this places the multi-year DiD in-between those. In the standard DiD case with simple treatment, one would expect the multi-year DiD and β_{FE_2} to coincide if 1) there is no heterogeneity, 2) the parallel trend assumption is fulfilled (see de Chaisemartin and D'Haultfoeuille, 2019). In our case here with rotation, there are even more reasons for the estimators to be different. In particular, one should note that we are not using all the switching groups (refer to Table 3).

As discussed in Section 3.3, by varying the control groups, we can identify the effects at different timings, when considering the same subperiod. For example, for the in-switchers CCSC, comparing this group to the CCCC identifies the effect at time 4, while comparing it to the SCSC group identifies period 2. As soon as we have more than three subperiods of four years, we have multiple estimates for the same year (i.e., the effect at year 2011 can be estimated from the 2007–2011 period using CCCC, or from 2009–2013 using SCSC). Under the assumption that the groups are trajectory-independent,[11] and under the parallel trends assumptions, we should expect the two estimates to be the same. Figure 15 shows that it is generally not the case: for most periods, the estimates are different. Only for 2014 and 2015 are the estimates for corn very close.

It would be interesting to investigate why estimates are so similar during some periods and different at other. I conjecture that this is due to the parallel trend assumption holding in the subperiods related to the 2014 and 2015. Unfortunately, this conjecture is difficult to substantiate in practice. We obviously cannot test for the parallel trends themselves. Most we can do is to test for the initial trends to hold in some of the sequences. This can be done comparing the initial values of the CCSC sequence to the control sequence CCCC. Figure 16 shows the placebo tests for the two first values of the CCSC sequence, compared to the CCCC ones. A positive value indicates that the treatment $\overline{SC}CC$ increased more rapidly than the control $\overline{CC}CC$. In general, we see deviations from the parallel trends, some of which appearing significative.[12] Interestingly, there is an opposite pattern for corn and soybeans. For corn, the placebo test is in general negative: the treatment group increased more slowly, or decreased

[11] More precisely, the assumption here is that we do not need to condition on the full trajectory of a sequence.

[12] Note that here significance depends crucially on the type of standard errors used. Without clustering at the county-year level, results would appear more significative.

Figure 15: Multi-period DiD: over-identifying test

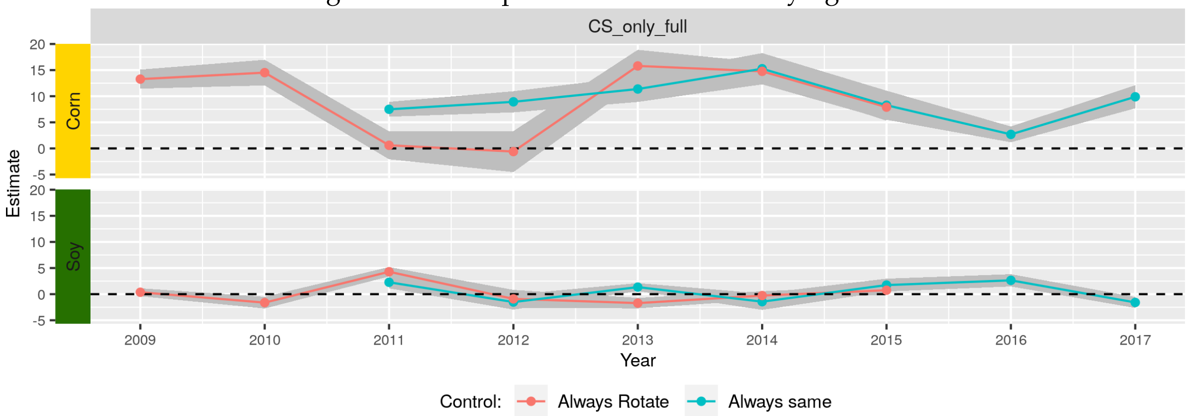

Figure 16: Multiperiod placebo DiD

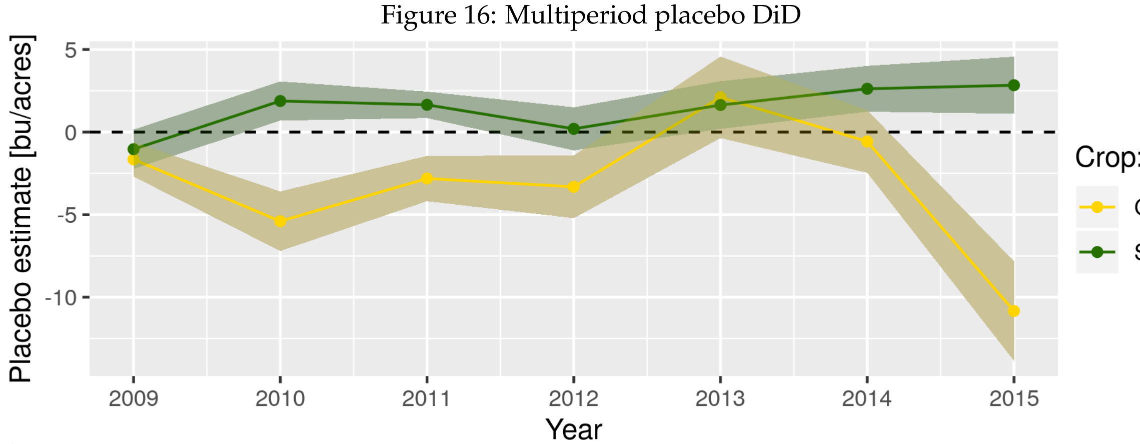

1 Note: The year refers to the end period of the placebo test: for example, the point for 2015 indicates that we were comparing the pair 2014–2015. Standard errors are clustered at the county-year level.

faster compared to the control group. For soybean, it is the opposite: the treatment group increased faster. For more details, Figure 17 in the appendix shows the two-year trend sequences for corn and soybeans.

The fact that the placebo test is rejected for many of the series definitely casts some doubt on the validity of the approach. It is too early at this point to explain why this is the case. Preliminary explorations focused on using more disaggregated time trends, doing state-specific DiD. The rationale for doing so is that we might still be comparing fields too far from each other. Results from this exercise do not suggest that placebo tests are less significative. Further steps will involve also controlling for weather in the regressions, in which case the parallel trend assumption will be done conditional on the weather covariates.

An alternative explanation to this phenomenon could be that the decision to rotate or not precisely depends on the previous values observed. Noting that in its second year, corn produced a lower values than expected, farmers opt to cultivate soybeans instead. The story would go the other way for soybeans: noting a good soy yield during the second period, farmers decide to rotate for corn. This latter story seems less plausible however, since a particularly bad soy yield would call for a rotation rather than continuing the same crop.

Table 6: Frequency of four-year rotation sequences

Crop	Sequence	Percentage
Corn	S_C_S_C	60.99
	C_C_C_C	13.13
	C_C_S_C	8.99
	C_S_C_C	8.31
	S_C_C_C	3.89
	C_S_S_C	2.96
	S_S_S_C	1.35
	S_S_C_C	0.38
Soy	C_S_C_S	73.41
	S_C_C_S	5.72
	C_C_C_S	5.36
	S_S_C_S	5.06
	S_C_S_S	4.95
	S_S_S_S	3.37
	C_S_S_S	1.68
	C_C_S_S	0.44

5 Conclusion

In this paper, I discuss methods to estimate the effect of rotating crops on the yield of each crop. I use a rich dataset comprising ten years of corn and soybean yields in the US Midwest, over nine states. I show first that fields with higher quality tend to be more often cultivated to corn than to soybeans. This phenomenon leads to counter-intuitive results. Doing a very crude comparison of means from fields doing always rotation versus fields doing always corn, I show that fields doing always corn have actually higher averages than fields doing always rotation. This paradox arises from the fact that better fields are planted more often to corn. I discuss then how to address this issue, adopting a fixed-effect approach. A plot-level fixed-effects approach solves the issue from comparing different fields, yet does not control year differences. Building on recent literature, I discuss how controlling for year fixed effects corresponds to a multiyear difference-in-difference approach. I show then that the rotation framework brings in an additional difficulty, in that the value from the previous year is never observed if there is rotation. I adapt then the standard DiD method for the rotation case, which implies to take longer differences.

Results show that controlling for plot-level fixed effects plays a critical role for estimation. The multi-DiD approach reveals an important time heterogeneity. Turning to a placebo exercise, I find significant deviations in pre-trends between control and treatment groups. Current efforts are focusing on understanding how robust this result is, and how it can be explained. A potential explanation could be that farmers base their decision on the pre-trend, which would invalidate the result. If this turns out to be the case, I will investigate the use of instruments that affect farmer's decision to rotate. Past prices and weather can be used as instruments, as they do impact last year decision, and probably do not impact directly the current year's yields.

6 Appendix

Figure 17: Placebo sequences for in-switchers CCSC or SSCS

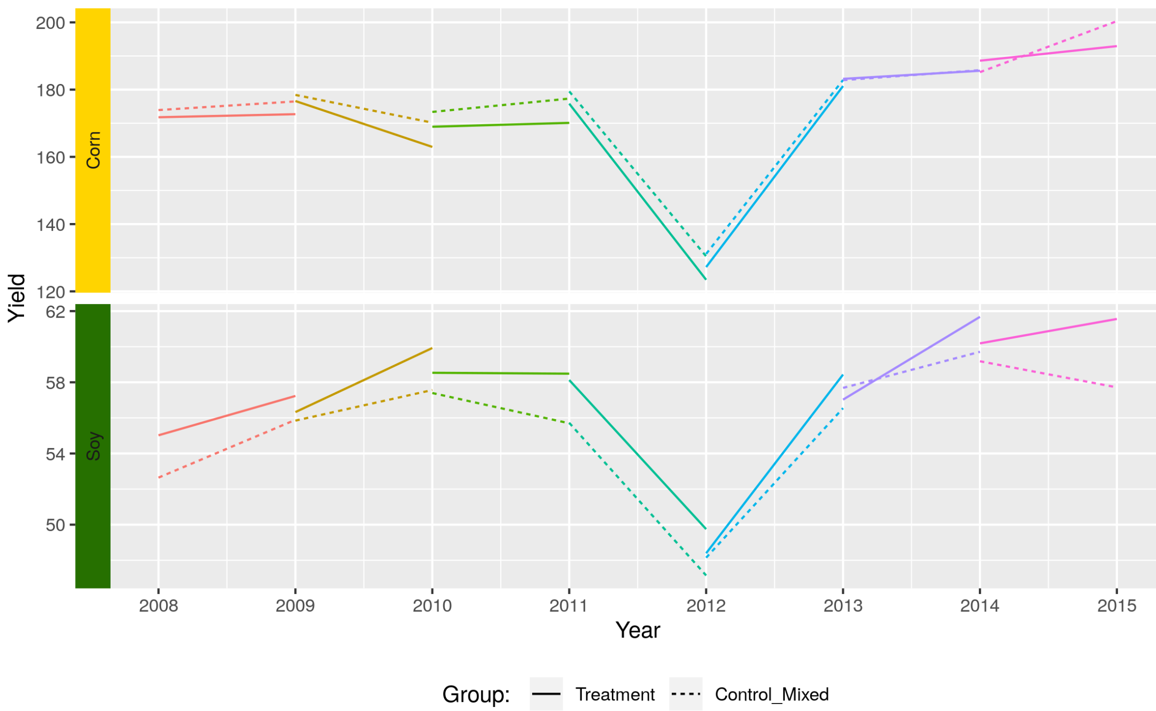

Figure 18: Recursive crop choice for all fields in the sample (2.2 mio)

Source: own computation, from CDL.

Figure 19: Share of county area for each crop

/home/matifou/Dropbox/GIS_server/output/figures/crop_location/mean_CS_Area_byCounty_rYlMpBc.

Source: USDA NASS. Share are here relative to total county area (not restricted to cropland).

Figure 20: Variation in number of rotation, by number of years planted to corn

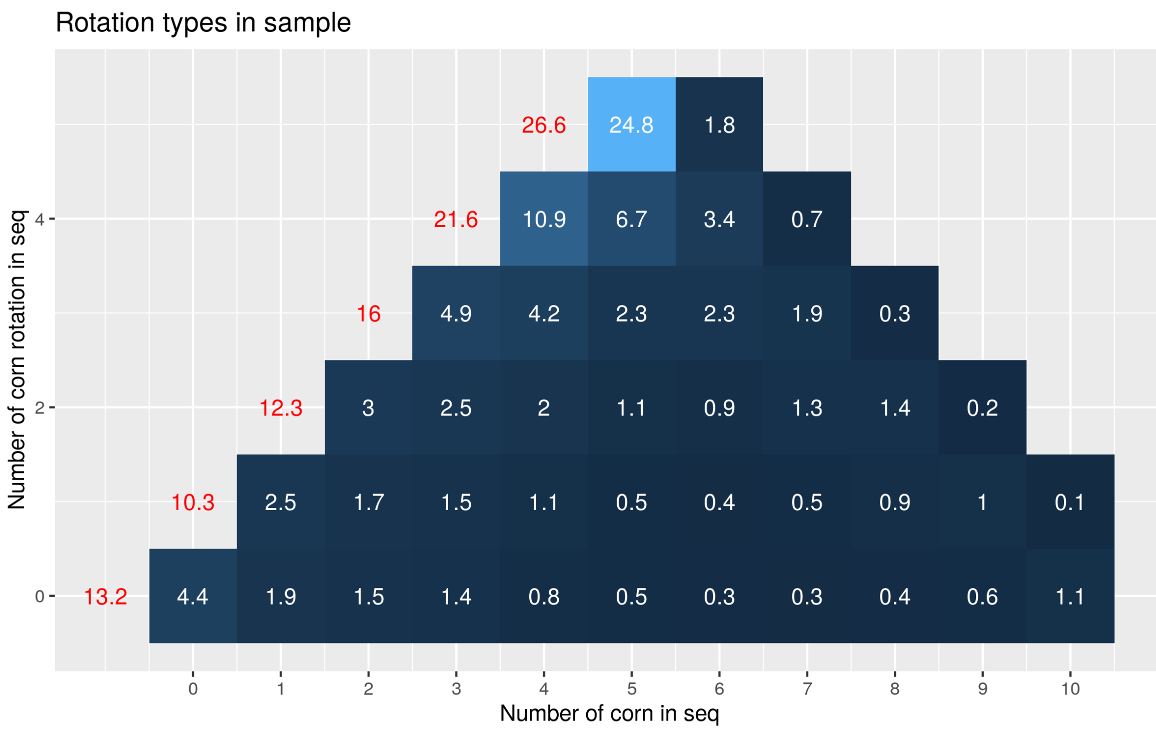

References

ANGRIST, J. D. AND J.-S. PISCHKE (2008): *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.

BORYAN, C., Z. YANG, R. MUELLER, AND M. CRAIG (2011): "Monitoring US agriculture: the US Department of Agriculture, National Agricultural Statistics Service, Cropland Data Layer Program," *Geocarto International*, 26, 341–358.

DE CHAISEMARTIN, C. AND X. D'HAULTFOEUILLE (2019): "Two-way fixed effects estimators with heterogeneous treatment effects," Tech. rep.

EL-NAZER, T. AND B. A. MCCARL (1986): "The Choice of Crop Rotation: A Modeling Approach and Case Study," *American Journal of Agricultural Economics*, 68, 127–136.

FARMAHA, B. S., K. M. ESKRIDGE, H. Y. K. G. CASSMAN, J. E. SPECHT, AND P. GRASSINI (2016): "Rotation Impact on On-Farm Yield and Input-Use Efficiency in High-Yield Irrigated Maize-Soybean Systems," *Agronomy Journal*, 108:23, 2313–2321.

GIBBONS, C., J. C. SUÁREZ SERRATO, AND M. URBANCIC (2019): "Broken or Fixed Effects?" *Journal of Econometric Methods*, 8, 1–12.

GOODMAN-BACON, A. (2018): "Difference-In-Differences With Variation In Treatment Timing," Tech. Rep. 25018, NATIONAL BUREAU OF ECONOMIC RESEARCH.

HENDRICKS, N. P., S. SINNATHAMBY, K. DOUGLAS-MANKIN, A. SMITH, D. A. SUMNER, AND D. H. EARNHART (2014a): "The environmental effects of crop price increases: Nitrogen losses in the U.S. Corn Belt," *Journal of Environmental Economics and Management*, 68, 507 – 526.

HENDRICKS, N. P., A. SMITH, AND D. A. SUMNER (2014b): "Crop Supply Dynamics and the Illusion of Partial Adjustment," *American Journal of Agricultural Economics*, 96, 1469–1491.

HENNESSY, D. A. (2006): "On Monoculture and the Structure of Crop Rotations," *American Journal of Agricultural Economics*, 88, 900.

IMAI, K. AND I. S. KIM (2019): "On the Use of Two-way Fixed Effects Regression Models for Causal Inference with Panel Data," Tech. rep.

LIVINGSTON, M., M. J. ROBERTS, AND Y. ZHANG (2015): "Optimal Sequential Plantings of Corn and Soybeans Under Price Uncertainty," *American Journal of Agricultural Economics*, 97, 855.

MACEWAN, D. AND R. HOWITT (2011): "An Empirical Model of Crop Rotations," Tech. rep., Prepared Paper for the 55 th Annual AARES Conference.

MORRIS, T. F., T. S. MURRELL, D. B. BEEGLE, J. J. CAMBERATO, R. B. FERGUSON, J. GROVE, Q. KETTERINGS, P. M. KYVERYGA, C. A. LABOSKI, J. M. MCGRATH, J. J. MEISINGER, J. MELKONIAN, B. N. MOEBIUS-CLUNE, E. D. NAFZIGER, D. OSMOND, J. E. SAWYER, P. C. SCHARF, W. SMITH, J. T. SPARGO, H. M. VAN ES, AND H. YANG (2018): "Strengths and Limitations of Nitrogen Rate Recommendations for Corn and Opportunities for Improvement," *Agronomy Journal*, 110, 1–37.

MUSSER, W. N., V. J. ALEXANDER, B. V. TEW, AND D. A. SMITTLE (1985): "A Mathematical Programming Model For Vegetable Rotations," *Southern Journal of Agricultural Economics*, 17.

NERLOVE, M. (1956): "Estimates of the elasticities of supply of selected agricultural commodities," *Journal of Farm Economics*, 38(2), 496–509.

NERLOVE, M. AND D. A. BESSLER (2001): "Expectations, information and dynamics," in *Handbook of Agricultural Economics*, ed. by B. L. Gardner and G. C. Rausser, Elsevier, vol. 1 of *Handbook of Agricultural Economics*, chap. 3, 155–206.

PORTER, P., J. LAUER, W. LUESCHEN, J. FORD, T. HOVERSTAD, E. OPLINGER, AND R. CROOKSTON (1997): "Environment affects the corn and soybean rotation effect," *Agronomy Journal*, 89, 441–448.

SAWYER, J. E. (2009): "Seasonal and Rotational Influences on Corn Nitrogen Requirements," Tech. Rep. 622, Iowa State Research Farm Progress Reports.

——— (2014): "Seasonal and Rotational Influences on Corn Nitrogen Requirements," Tech. Rep. 2046, Iowa State Research Farm Progress Reports.

SCHNITKEY, G. (2018): "Crop Budgets, Illinois," Tech. rep., Department of Agricultural and Consumer Economics, University of Illinois.

SEIFERT, C. A., M. J. ROBERTS, AND D. B. LOBELL. (2017): "Continuous Corn and Soybean Yield Penalties across Hundreds of Thousands of Fields," *Agronomy Journal*, 109, 541–548.

THOMAS, A. (2003): "A dynamic model of on-farm integrated nitrogen management," *European Review of Agricultural Economics*, 30, 439.