Unintended Environmental Policies:

The Impact of European Agricultural Subsidies on Pollution*

Emma Du Puy†

May 17, 2026

Abstract: *I study the environmental impact of subsidies in polluting markets when producers differ in their pollution intensity. I develop a dynamic model in which the elasticity of substitution between clean and polluting inputs determines the sign of the correlation between firms' efficiency and their pollution intensity. I apply it to EU agricultural subsidies, where a shift-share analysis quantifies their micro-level impacts, and model estimation reveals that productive farms pollute more. Subsidy designs that alter market selection are consequently found to reduce aggregate pollution. Counterfactuals show that while taxes maximize welfare, subsidies achieve environmental gains without the large consumer and producer surplus losses of taxation.*

JEL Codes: L11, Q18, Q52, H23, D24
Keywords: Pollution, Environmental Policy, Agriculture, Industry Dynamics, Production Function

I thank my advisors Suresh Naidu, Pietro Tebaldi, Wolfram Schlenker, Gautam Gowrisankaran and Jeffrey Shrader for their guidance during this project. I am also indebted to Matthew Backus, Mathias Reynaert, Bernard Salanié, Sylvain Chabé-Ferret, Eric Verhoogen, Jack Willis, Allan Hsiao, Robert Metcalfe, Rodrigo Soares, Karam Kang, Sebástian Otero, Diego Restuccia, Marguerite Obolensky and Nicolas Longuet-Marx for their generous feedback, as well as Yulia Kim for her support. I am grateful for comments from numerous participants in seminars and conferences at Columbia University, the Toulouse School of Economics, the CUNY Stone Center on Socio-Economic Inequality, Sciences Po, the University of Saskatchewan, the Paris School of Economics, CIRED, NC State, the Northeastern Workshop on Energy Policy and Environmental Economics, the Bozeman Applied Economics Summer Conference, Camp Resources, NAREA, WEAI, the 8th Social Cost of Water Pollution Workshop, IIOC. All remaining errors are my own. I would like to acknowledge funding from the Columbia Center for Political Economy and the Columbia Center for Environmental Economics and Policy, and thank INSEE, the CASD and the French Ministry of Agriculture for data access. A previous version of this paper was circulated under the name "An Empirical Model of Agricultural Subsidies with Environmental Externalities".

†Du Puy: Center for Political Economy, Columbia University (td2631@columbia.edu)

1 Introduction

The design of policies that aim to reduce pollution is influenced by the characteristics of their targeted market, from the state of competition to the heterogeneity of the producers (Goulder and Parry (2008)). Just as these features shape the optimal design of environmental regulation, they also govern the environmental side effects of policies that are otherwise not explicitly environmentally focused. This paper highlights this distinct and often overlooked channel of interaction: the unintended environmental consequences of government intervention in polluting markets. Specifically, I study the environmental impact of producer subsidies in settings where active market participation is endogenous to policy, and producers differ in both efficiency and pollution intensity.

Agriculture offers an ideal setting to study these interactions. It combines substantial producer heterogeneity with extensive government support—nearly $750 billion annually (FAO (2024))—that heavily shapes market participation and distorts production. It is also an industry responsible for widespread and loosely regulated air and water pollution, which has proven difficult to abate (Tschofen et al. (2019)). Despite the scale of these expenditures, their environmental externalities remain poorly understood.

I analyze these mechanisms in the context of the European Union's Common Agricultural Policy (CAP) and the French grain market. This setting is fitting for two reasons. First, the CAP is the largest EU industrial policy, accounting for 45% of the EU budget since 1980. Second, grain production is highly reliant on both fertilizers and pesticides, two key sources of polluting chemicals. While these subsidies may generate substantial environmental externalities, they remain a major source of uncertainty for regulators (European Court of Auditors (2024)).

Subsidies affect environmental outcomes through two distinct mechanisms. First, they alter the intensive margin of production by changing farm-level input and crop choices. Second, they distort the extensive margin of market selection by retaining marginal, less productive farms. The environmental impact of this selection channel thus depends on the correlation between farm productivity and the intensity of chemical use. Where productivity is driven by resource efficiency, marginal farms tend to be more chemically intensive, making retention both economically and environmentally costly. Conversely, if productivity derives from chemical intensification, a trade-off emerges: highly productive farms are the most polluting, and subsidies that retain marginal producers may inadvertently reduce aggregate pollution.

The analysis proceeds in two steps, drawing on rich administrative data from the French grain market. First, I leverage the end of EU price support as a natural experiment to provide reduced-form evidence on the impact of subsidies on chemical use, pollution, and market participation. Second, I estimate a dynamic model of the grain market to decompose

these effects and evaluate counterfactual policies. The French setting allows me to combine a representative panel of producers with the universe of farms in the Agricultural Census, enabling a precise decomposition of intensive and extensive margin responses.

The end of European agricultural price support is a reform undertaken by the EU under international pressure during the GATT Uruguay Round. This 1992 reform replaced heterogeneous price support with more uniform land subsidies, simultaneously lowering per-unit subsidy support and altering incentives for both crop and input choices. In this analysis, I leverage the reform's differential impact across crops using a shift-share design. Specifically, I use farm-specific pre-reform crop shares as instruments for exposure to the reform. Controlling for the determinants of crop choice, this approach mimics a difference-in-differences design that isolates causal effects using the conditional heterogeneity of shares (Goldsmith-Pinkham et al. (2020)). I find that a one-standard-deviation increase in exposure to the reform reduced farm profits by 31% and chemical use by 24%. I also find evidence for accelerated market consolidation: the same exposure reduced the number of active farms by one per municipality. Finally, using satellite data to track algal blooms, I document that these adjustments translated into a 10% decline in observed water pollution.

The second part of the paper develops a model of the French grain market to quantify the channels through which subsidies affect chemical use and to assess their welfare implications. The model builds on the dynamic single-agent industry framework of Hopenhayn (1992) and captures the main drivers of farm chemical use. Here, farms differ both in total factor productivity (TFP) and chemical use efficiency, and crops vary in their chemical reliance. Finally, farms are multiproduct producers that accumulate capital over time, which impacts the relationship between overall productivity and exit decisions.

The estimation proceeds in three steps. First, I recover iso-elastic demand curves using weather shocks as supply shifters. Second, I estimate the production parameters. This step faces two identification challenges: producers have unobserved heterogeneity in both TFP and chemical-use efficiency, and while farms are also modeled as multiproduct, all inputs in the data except for land are observed at the farm rather than the farm-crop level. I address the endogeneity issue stemming from the unobserved farm type by inverting first-order conditions for profit maximization and specifying a process for this type. In so doing, I recover the moments needed to both identify the production parameters and to recover the distribution of unobserved productivity shocks, building on Doraszelski and Jaumandreu (2018). To account for multiproduct producers, I follow De Loecker et al. (2016) and restrict estimation to single-product farms, for which the input-output allocation is known. I control for the induced estimation bias with an extension of the correction developed by De Loecker et al. (2016). Third, I use a simulated method-of-moments estimator to recover the dynamic costs of capital adjustment, entry, and incumbency.

The estimation yields a central empirical result: an elasticity of substitution between land and chemicals of 1.47 induces a positive correlation between farm productivity and chemical intensity. The substitution pattern is intuitive if one thinks of organic farming, which replaces high chemical use with more intensive land use. This correlation implies that marginal incumbents, whose market participation decision is impacted by subsidies, will pollute less than infra-marginal farms on average. Thus, subsidies that keep marginal producers in the market reduce aggregate pollution intensity while also lowering production efficiency. The estimated dynamic parameters then govern the strength of this selection effect by determining the volatility of productivity and the friction of exit.

Policy Implications: I use the model to analyze optimal policy design. First, I compare how different subsidy instruments interact with the intensive and extensive margins of pollution. A yearly lump-sum payment to low-pollution producers achieves environmental gains by retaining marginal, low-pollution farms on the market—a pure selection effect. In contrast, land subsidies operate through an intensive margin: by raising the relative price of chemicals, they induce substitution toward land use. Lump-sum payments reduce chemical use intensity more effectively, but are less effective at lowering aggregate chemical use because they stimulate a larger expansion in production.

Second, borrowing the French Government's valuation of agricultural chemical pollution (CGDD (2011)), I identify three welfare-maximizing instruments conditional on design: a 28k€ lump sum payment, a 10% homogeneous land subsidy, and a differentiated land subsidy for wheat and non-wheat crops (10%, 40%). They reach a 1%, 1.4%, and 2.2% decline in total chemical use respectively, and a marginal value of public funds of 1.03, 2.9, and 1.6. Across these designs, the differentiated land subsidy performs best, delivering larger environmental gains and a small positive effect on consumer surplus. However, these subsidies present as a poor substitute for the optimal chemical tax, which achieves a 54% reduction in aggregate chemical use. Some tension also exists, as the tax delivers far greater environmental gains at substantial cost to both producers and consumers.

Related Literature: My paper contributes to the study of optimal environmental policy, which has long stressed the importance of market structure and equilibrium effects (Oates (1991); Bovenberg and Goulder (1996)). The work of Carlton and Loury (1980) is particularly relevant here in showing how entry and variation in the number of active firms in a polluting market can make a Pigouvian tax fail. Recent work at the intersection of environmental economics and industrial organization also studies how environmental policies perform at equilibrium, relying on structural methods to tailor empirical models to specific markets (Blundell et al. (2020); Armitage (2023); Rafey (2023); Gowrisankaran et al. (2024)), with a growing focus on dynamic land use in agriculture (Scott (2013); Hsiao (2022); Burlig et al.

(2024); Obolensky (2025)). I contribute to this literature by developing a framework that explicitly models the dual environmental impact of agricultural subsidies: their effect on production choices within the farm, and their effect on market participation through farm exit. Disentangling these channels is critical for policy design. As my analysis shows, the optimal policy depends fundamentally on the identity of marginal incumbents, and whether the largest environmental gains come from changing incumbent production behavior or from altering the set of incumbent farms. In contrast to recent papers on dynamic land use, I also have access to rich farm-level administrative data. This allows me to model farms as multiproduct firms accumulating capital, which are important features for my analysis. The first captures how farm-level crop mix impacts aggregate chemical use, and the second helps capture the extent to which exit selects on productivity.

While most of the environmental economics literature has focused on the design of policies explicitly targeting the environment, I highlight the relevance of environmental considerations in the design of generic industrial policies within polluting markets. Previous work along these lines has largely documented cases where subsidies had unexpected negative environmental spillovers, such as fuel subsidies that increase emissions and deplete natural resource stocks (Davis (2017); Englander et al. (2023)), or trade protections for carbon-intensive industries (Shapiro (2021)). In contrast, I show that a major subsidy regime can have unintended, positive environmental consequences. Here, I focus on the well-documented case of agricultural pollution (Missirian (2020); Frank (2021); Chabé-Ferret et al. (2021); Taylor (2022); Dias et al. (2023)), measuring fertilizer pollution directly through satellite data (Taylor and Heal (2023)). My core theoretical contribution is to embed the longstanding literature on input-biased technical change in agriculture (Hayami and Ruttan (1971); Binswanger (1974); Bustos et al. (2016); Clemens et al. (2018)) within a dynamic, heterogeneous-agent industry model. As such, this approach is most similar in structure to the analysis of the declining U.S. labor share by Oberfield and Raval (2021), and allows me to model agricultural intensification and its environmental outcomes as the result of micro-level production decisions.

Finally, this paper contributes to the literature evaluating the welfare consequences of agricultural subsidies. Agricultural subsidies are traditionally understood as generating significant misallocations in the labor market, both across farms of varying productivity and across economic sectors, in turn slowing down structural transformation in developing countries (Gollin et al. (2013); Adamopoulos and Restuccia (2014)). Here, I turn my attention to the largely understudied EU CAP. In a labor misallocation context, any distortions generated by the CAP likely bear little economic significance, given the relatively small share of EU labor allocated to agriculture. Its impact on input use misallocation, however, is still a relevant question when we turn to land use. Agriculture accounts for 38% of all EU land, making it

an industry central to most contemporary environmental questions.[1] While several recent papers have focused on the extensive margin of agricultural land use—the transfer of land out of agriculture into a conservation status (Aspelund and Russo (2024); Larsen (2025) and Grupp et al. (2024) in the European context)—the intensive margin related to the sustainable use of land within agriculture is less understood. In my welfare analysis, I pit the environmental consequences of agricultural subsidies against their better-known effects on economic efficiency to test the general relevance of this environmental channel. I find that, for standard valuations of pollution damages, the magnitude of the environmental benefits from subsidies can be large enough to dominate their concurrent effects on economic efficiency, providing a new and consequential dimension to the evaluation of agricultural policy.

This paper proceeds as follows: Section 2 provides background on the chemical use in French agriculture, the Common Agricultural Policy and describes the data. In Section 3, I use a shift-share design to evaluate the consequences of the MacSharry reform. Section 4 outlines the structural model and the estimation strategy, and Section 5 presents counterfactual simulations. Finally, Section 6 concludes the analysis.

2 Data and Context

2.1 Data

My primary analysis relies on the French subset of the EU-wide Farm Accountancy Data Network (FADN). The French subset is a representative panel of approximately 8,000 commercial farms, which accounts for 95% of national agricultural production. This accounting-based dataset provides detailed annual information on crop-specific quantities and sales (allowing for the recovery of farm-level prices), input expenditures, capital stocks, and land allocations to crops. To examine farm exit and the distribution of farm size, I supplement this with the French Agricultural Census (1970–2010), which provides data on the universe of French farms. Due to the lack of stable farm identifiers prior to 1998, I aggregate census data into a panel of French municipalities. Appendix B details sample construction, variable definitions, and additional data sources.

2.2 Determinants of Farm-Level Chemical Use Intensity

French farm data reveal substantial heterogeneity in chemical use at the farm level. I compute ratios of chemical expenditure to revenue for grain farms in the FADN. I use national output prices and farm-level quantities to calculate the revenue component of this ratio so as to focus on cross-farm heterogeneity in chemical use. Over 1985-1990, the standard deviation

[1] See for EU land use: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Farms_and_farmland_in_the_European_Union_-_statistics.

of this ratio is 20%, for an average of 22.3%. Figure A1 shows the entire distribution. This heterogeneity is likely driven by several factors, which will interact with agricultural subsidies at equilibrium: 1) variation across farms in the price of chemicals relative to other inputs, 2) variation across crops in their chemical requirements interacting with heterogeneity across farms in crop mix composition, and 3) heterogeneity across farms in their efficiency at using chemicals. The first two channels are intensive: subsidies that change incentives for input use or crop choice will impact all existing producers. The last channel is extensive: subsidies that change the average profitability of the market will impact the exit decision of marginal incumbent farms. If farms differ in chemical use efficiency, farms' exit decisions should correlate with both their heterogeneity in chemical use efficiency, and in turn with their chemical use intensity.

I provide suggestive evidence for the last two mechanisms in Figure A1. First, cross-crop heterogeneity in nitrogen manure per ha illustrates that different crops have different chemical input needs. Second, French farms increased their pesticide-to-land and fertilizer-to-land ratios between 1980 and 2000, despite a concomitant rise in the price of these chemicals relative to that of land. As such, if land and chemicals are substitutes in production, a leading explanation for this increased reliance on costlier inputs would be rising chemical use efficiency over time (Doraszelski and Jaumandreu (2018)).

2.3 The Common Agricultural Policy

Prior to 1992, the Common Agricultural Policy (CAP) supported farmers through guaranteed price floors and export subsidies, which kept French farm-gate prices significantly above world market levels.[2] Driven internally by an EU budgetary crisis and externally by the GATT Uruguay Round negotiations, the 1992 MacSharry Reform fundamentally altered this structure by replacing price support with direct land-based payments.

The reform had three main components pertinent to farm incentives. First, it drastically reduced intervention prices for cereals and removed them for oilseeds.[3] Figure 1 illustrates the identification leverage provided by this shock: French farm-gate prices converged to lower world market levels, but the magnitude of this drop varied substantially by crop, creating differential exposure based on a farm's pre-reform crop mix. Second, the EU introduced compensatory payments per hectare, which were decoupled from current yields but condi-

[2]These price objectives were enforced through different means on the cereal and oil crop market. In the cereal market, they were enforced through three main instruments: government purchasing at floor prices on the secondary market (purchasing mostly from grain elevators), subsidies for exports to the world market, and levies on agricultural imports. Import levies were not possible on the oil market due to the 1962 GATT Dillon Round. The EU hence enforced target prices on the oil market by reimbursing processors for the difference between international market prices and target prices when they were buying European grains.

[3]To summarize this decrease in prices, I construct a land-share-weighted average output price for each farm. Figure A4 illustrates it as a sharp, permanent drop of 33% in this farm-level price index between 1991 and 1993.

tioned on department historical references.[4] Third, it implemented a mandatory set-aside requirement, obliging larger commercial farms to leave a portion of land fallow to receive subsidies. This policy shift—exchanging heterogeneous price protections for relatively uniform land subsidies—provides the variation that I exploit for identification.[5]

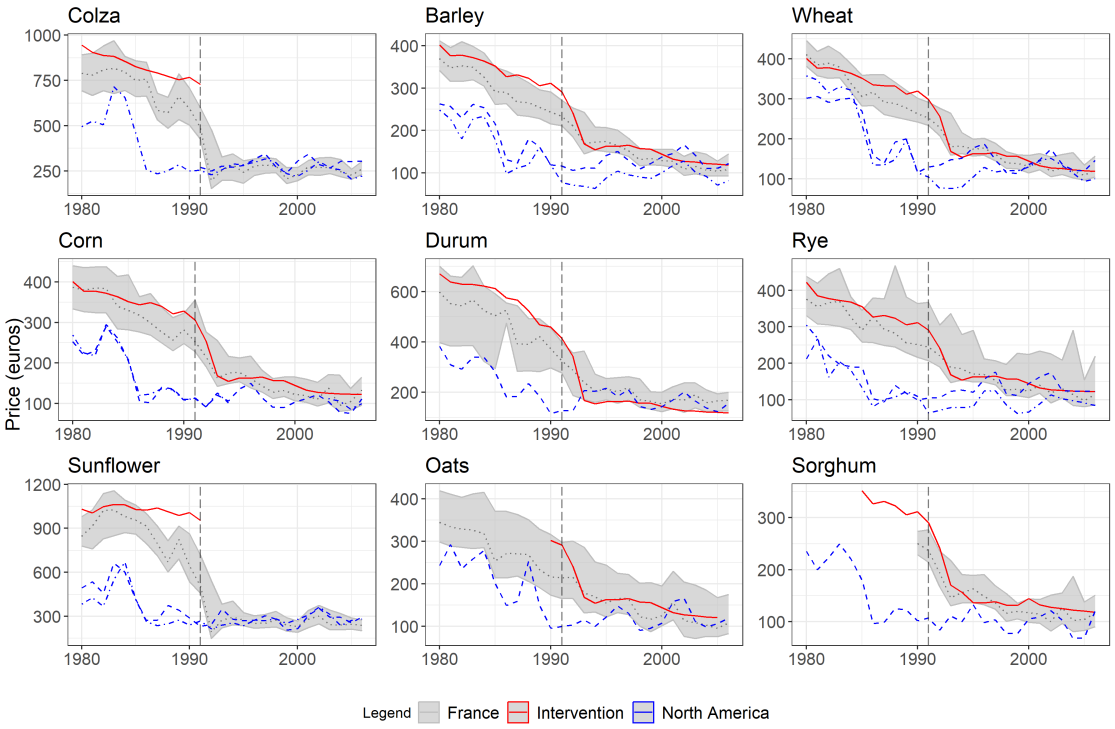


Figure 1: Convergence of French and North American Prices Following the MacSharry Reform

Notes: Evolution of French and North American farm gate prices, as well as EU intervention prices for most cereal and oil crops. French farm gate prices are shown with the grey ribbon, and taken from the FADN. The center dotted line corresponds to the weighted mean, and the edges to the bottom and top 25th and 75th centiles of the distribution of prices. Intervention prices are shown with the solid red line. The data is digitized, taken from the relevant EU directives published over time. U.S. and Canadian farm-gate prices are taken from Faostat, and shown with the dashed and dot-dashed blue lines. All prices are converted into 2020 euros using the relevant exchange rates and correcting for inflation using the Insee's consumer price index.

The MacSharry reform coincided with a significant acceleration in market consolidation among French cereal and oil crop producers. Around the time of the reform, there was both an acceleration in the growth of average farm size, and a large reallocation of land towards

[4] Figure A5 shows their evolution over time in the FADN data.

[5] Figure A3 details the evolution of French grain imports in the years before and after the MacSharry reform. Post-reform, wheat has a small increase in imports, which remains quite insignificant compared to national production levels. Imports reach 400k tons post-reform, for an average of 33 million tons of national production. Colza seed had a very precise spike in 1994, also reaching 400k tons, but immediately falling to almost zero in ensuing years. Sunflower is the only grain crop with a more marked increase between 1996 and 1999 at 294k tons, or 15% of national production. Sunflower, however, represents a small fraction of the overall grain market. These small trade effects justify the focus on a closed economy model in the second part of the paper.

larger farms. Between the 1988 and 2000 Censuses, the share of land cultivated by the top decile of farms jumped from 25% to 49%. Farm exit trends are also informative, with the 1988-2000 decline in French cereal producers being comparable to trends in prior decades, while the number of oil crop producers was increasing up to 1988, and started decreasing post-reform. These dynamics are illustrated in Figure A2. This difference in exit trends maps well with the reform-specific price decreases seen in Figure 1.

3 Shift-Share Analysis

3.1 Empirical Strategy

Approximating European Policy: The reform altered the entire system of revenue support for grain farms. I summarize the total government support for each crop c in year t as a 'revenue cushion', measured in euros per unit of output. This cushion is the sum of the pre-reform price support and the post-reform direct subsidies:

$$cushion_{ct} = \underbrace{\mathbb{1}\{Intervention_{ct}\}\left(P_{ct}^{inter} - \overline{P}_{ct}^{I}\right)}_{\text{Price Support}} + \underbrace{\left(\frac{\overline{Sub}_{ct}}{\overline{Yield}_{ct}} - setaside_{ct}\right)}_{\text{Land Subsidy}}.$$

The price support wedge captures the pre-reform system of demand-side government intervention. It is defined as the difference between the EU intervention price (P_{ct}^{inter}) and the world market price (proxied by North American farm-gate prices: $\overline{P}_{ct}^{I}$), active only when an intervention price exists.[6] The land subsidy captures the post-reform system. It consists of the new land subsidy ($\overline{Sub}_{ct}$), converted to a per-output measure through division by an average yield ($\overline{Yield}_{ct}$), minus the opportunity cost of the compulsory set-aside requirement ($setaside_{ct}$).[7][8] Figure A6 plots the evolution of these revenue cushions. The cushions for all major crops decline sharply between 1991 (the final pre-reform year) and 1995 (the final reform implementation year). Critically, the magnitude of this decline varies substantially across crops, with oil crops and durum wheat experiencing a much larger negative shock than the remaining grains. Overall, the reform enforced a convergence in subsidization across crops, as illustrated in Figure A8.

[6] The U.S. and Canada accounted for most of the export market over the 1980s to 2000s, making these prices a relevant reference point for the analysis. I use USA prices but for colza, for which I use Canadian prices—given that the USA was not a significant producer of colza at that time, but Canada was. All these prices are average farm gates prices from Faostat.

[7] I rely on the median and not average subsidies, as the subsidy data recorded in the FADN shows significant mis-measurement. The average subsidy per hectare decreases over time in the FADN, in a pattern that does not match EU policy, while the median value does.

[8] The per unit cost of the set-aside is recovered as 10% of the amount of land required to produce one ton of the given crop, times the price of land to which is removed the crop-specific land subsidy. This subsidy is removed from the price of land, as farms would still receive land subsidies for set-aside areas.

Building the Reform Exposure Variable: I measure each farm's exposure to the reform using a weighted average of these crop-specific shocks, where the weights are the farm's pre-reform land allocations.[9] I define farm j's time-varying average cushion as:

$$Cushion_{jt} = \sum_{c \in \mathbb{C}_{j,t_0}} \frac{s^l_{j,c,t_0}}{\sum_{c'} s^l_{j,c',t_0}} cushion_{ct},$$

where s^l_{j,c,t_0} is the share of farm j's land allocated to crop c in the pre-reform period, $t_0 = 1991$. These cropping decisions were made in the fall of 1990, before the reform's design was finalized. The farm-specific reform exposure variable Exp_j is then the total change in this revenue cushion over the reform period: $Exp_j = Cushion_{j,t=91} - Cushion_{j,t=95}$. This variable captures how hard each farm was hit by the reform, based entirely on its pre-reform state. Its distribution in the sample is shown in Figure A10. In Appendix C, I show results where exposure is constructed using land shares to aggregate across crops for the land subsidy and set-asides, and output shares to aggregate across intervention prices.

Identifying Assumptions: A recent literature distinguishes between shift-share designs where identification comes from a large number of plausibly random shocks (Borusyak et al. (2022)) and those where it comes from the interaction between a few dominant shocks and quasi-random exposure shares (Goldsmith-Pinkham et al. (2020)). On top of having a unique year of policy shock, decomposing the variation of the exposure variable into crop-by-crop sources following the Rotemberg weights of Goldsmith-Pinkham et al. (2020), I find that variation in Exp_j is dominated by shocks to a few key crops (sunflower, colza and durum wheat). Weights are shown in table A10, along with crop-specific losses in subsidies and the variance in pre-reform crop shares within the sample. The many endogenous shocks approach is therefore not a suitable one here.

Consequently, my design follows the identification strategy outlined in Goldsmith-Pinkham et al. (2020). This approach is conceptually similar to a continuous treatment difference-in-differences, where a farm's pre-reform crop mix determines its treatment intensity. The core identifying assumption is that the pre-reform shares, conditional on controls, are not correlated with unobserved, crop-specific trends that would have affected farm outcomes in the 1990s even in the absence of the reform. I will test and discuss the validity of this assumption in the following sections.

[9] I include the following crops to compute the farm-level cushion: wheat, barley (winter), corn, rye, oats, sunflower, rapeseed, barley (spring), durum, sorghum.

3.2 Farm-Level Analysis

Empirical Design: My primary analysis uses the FADN farm-level panel. I estimate the dynamic impact of reform exposure using the following event-study specification:

$$Y_{jt} = \alpha_j + \eta_{d(j)t} + \sum_{\tau \neq 1991} \theta_\tau Exp_j \mathbb{1}\{t = \tau\} + \sum_{\tau \neq 1991} \delta_\tau X_j \mathbb{1}\{t = \tau\} + \varepsilon_{jt}. \tag{1}$$

Here Y_{jt} is the outcome for farm j in year t, Exp_j is the farm-specific, time-invariant exposure measure, which I standardize. The coefficients of interest $\{\theta_t\}_t$ capture the year-specific effect of the reform relative to 1991. The vector X_j contains pre-reform farm characteristics, the $\eta_{d(j)t}$ are department-by-year fixed effects, and the α_j farm-specific fixed effects. Standard errors are clustered at the department level.

The causal interpretation of $\{\theta_t\}_{t \geq 1992}$ hinges on the conditional exogeneity of the exposure measure. Formally, the identifying assumption is that a farm's pre-reform crop shares are uncorrelated with unobserved factors that would have driven differential growth in the outcome variable post-1991, conditional on the controls. The pre-1992 coefficients $\{\theta_t\}_{t<1992}$ provide a direct test of this assumption, in that they must be non-statistically different from zero.

Threats to the Identification Strategy: The primary threat to identification is that 1991 crop choices were not random, but rather determined by farm characteristics that could also be correlated with future farm growth. I structure my strategy, embedded in the inclusion of a farm fixed effect, on the inclusion of controls X_j and department-specific time trends $\eta_{d(j)t}$, to neutralize two main classes of threats which relate to this.

First farm omitted state variables and technology trends. Here the central concern is that a farm's state vector—both observed and unobserved—jointly determines its crop portfolio and its future growth trajectory. An obvious example is that farms located on land particularly fertile for one crop might tend to specialize more, and earn high yields and profits. The farm fixed effects remove the time-constant influence of any time-constant unobserved farm confounder. Dynamic elements of the farm state could also bias the results.[10] To address this, X_j includes a rich set of controls for the farm's pre-reform state. To proxy for unobserved productivity, which is a key unobserved state, I include the farm's 1991 choices of flexible inputs following the logic of the production function literature (Olley and Pakes (1996)). Because farms growing different types of crops are also likely to be more different along unobserved characteristics, it is important to control for the general split of production across crop categories. Finally, it is also key to explicitly account for farm-specific technol-

[10]For example, in a multi-product setting with separate production decisions across production lines, more productive farms are likely also more diverse (Mayer et al. (2014)).

ogy adoption paths. Overall X_j contains the following elements: farm 1991 capital stock, total labor used, total land use, total profit, their chemical use, the number of crops they grew and the evenness of their land allocation across these crops, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 (the start of my period of analysis) adoption trends in chemicals measured as the evolution in their chemical use.[11]

The second threat comes from spatially correlated shocks. Local soil and climate conditions create strong spatial patterns in cropping, and I do observe systematic variation across space in farm exposure in Figure A9. If these local areas are also subject to unobserved, time-varying shocks (e.g., evolving local demand or input prices), these could induce a spurious correlation between exposure and the outcome. The inclusion of department-by-year fixed effects $\eta_{d(j)t}$ controls for any such shocks that are common to all farms within a department in a given year.

After conditioning on this rich set of controls, the identifying variation comes from sources of specialization that are plausibly orthogonal to the confounding trends of the 1990s. A primary source is stickiness in crop choice. Farms can face significant switching costs related to crop-specific physical capital (e.g., planters, harvesters), human capital, and crop cycles, as discussed by Livingston et al. (2008) and Scott (2013). This stickiness means that initial conditions at entry and past, potentially random, shocks to prices or yields can lead to persistent specialization patterns that are uncorrelated with the specific productivity trends of the study period. Similarly, historical factors, such as lumpy investments made decades prior or generational expertise in a specific crop, can create quasi-random variation in the 1991 crop portfolios that my design leverages for identification. Finally, to ensure results are driven by within-farm changes, the analysis is conducted on a balanced panel of farms present from 1985 to 2002, using FADN sampling weights in all regressions.

Testing the Design: The event-study coefficients for the pre-reform period provide the primary validation of the identification strategy. In addition, Figure A11 shows balance tests comparing trends for farms with above- and below-median exposure. Pre-reform trends are relatively similar across low and high exposure groups, while there is a sharp drop in sales and profits post 1992 for the high exposure group. Crop-specific balance tests comparing farms with an above and median crop share show comparable trends pre-treatment (Appendix C.2).

I borrow the Rotemberg weights decomposition of a shift-share variable from Goldsmith-Pinkham et al. (2020), in order to analyze the variation driving the effect $Exposure_j$ in

[11] Note that the inclusion of the share of production allocated to oil crops as a control does not prevent me from using variation coming from farms growing relatively more sunflower or colza, as shown by the Rotemberg weights in Table A10. Rather, it controls for the overall decision to grow oil crops, but then uses variation across farms which grow relatively more sunflower or colza.

Equation (1). These Rotemberg weights are shown in Table A9, and indicate which crops within the grain market are driving the variation in $Exposure_j$. After the inclusion of controls and fixed effects, this variation is driven by comparing farms that had a high or low share of land allocated to sunflower, colza and durum wheat, relative to other grain crops. These three crops are both the three crops with the highest loss in subsidization from the reform, and with a high pre-reform variance in land shares across farms. Other grain crops have both a much lower, and a much more homogeneous loss in subsidization following the reform. These weights imply that concerns about omitted variable biases should be focused on the uncontrolled-for determinants of crop choice for these three specific crops.

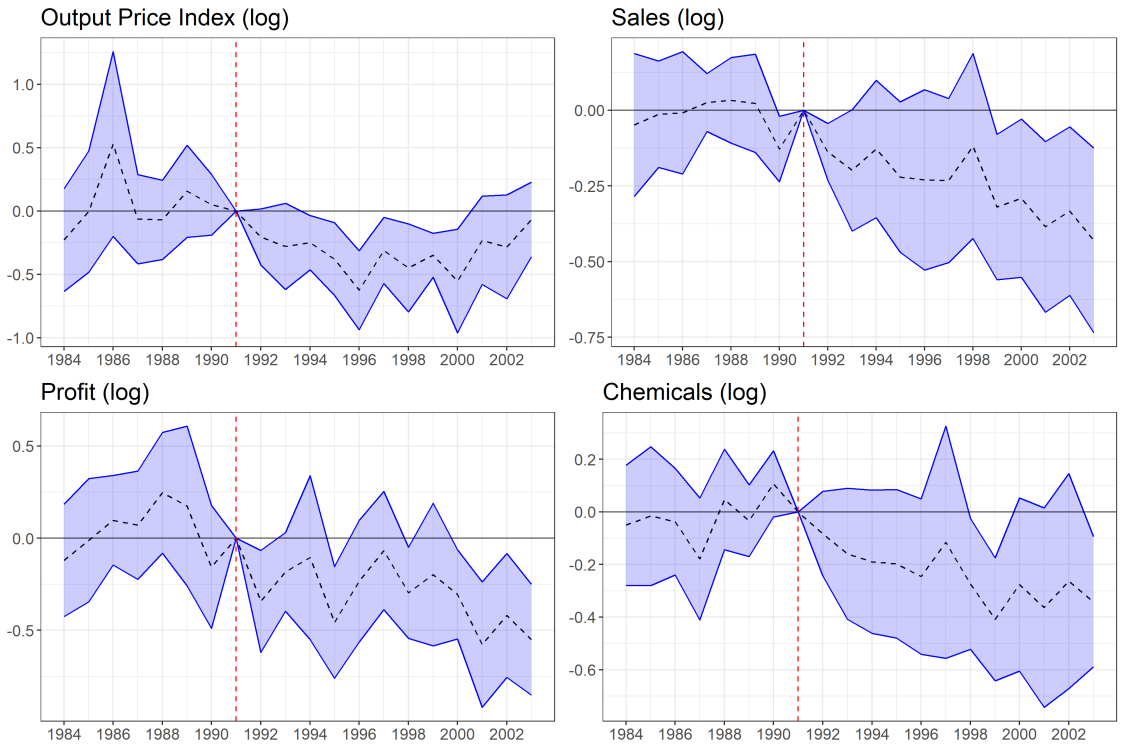


Figure 2: Exposure to the Reform—Farm-Level Event Study

Notes: All event study outcomes are measured in logs: a farm-level output price index (average price across grains), total sales, total profit, and the farms' deflated chemical bill. All coefficients give the year-specific effect of reform exposure. The vector of controls contains the following elements: farm 1991 capital stock, total labor used, total land use, total profit, their chemical use, the number of crops they grew and the evenness of their land allocation, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 adoption trends in chemicals measured as the evolution in their chemical use. I add department-year and farm fixed effects, and cluster the standard errors at the department level level.[12] The table for these results is Table A4.

I also recover negative Rotemberg weights associated with wheat, corn, spring barley and oats. Apart from wheat, all these grains have a small land share variance across farms. When I decompose exposure into crop-specific variables in Appendix C.4, I see that a high

crop share in 1991 in any of these crops is also associated with similar effects on sales, profit, and chemicals as the ones identified by the aggregate exposure in the main specification. These two features imply that these negative weights should not bias the interpretation of the aggregate effect of farm exposure (Goldsmith-Pinkham et al. (2020)).

Results: The results of the event-study analysis are presented in Figure 2. Coefficients provide the effect of a one standard deviation increase in exposure to the reform, and do not capture the aggregate effect of the reform. The effect on farm level output prices (labeled output price index) serves as a first stage confirming that an increase in the exposure variable is associated with decreased output prices.[13] The pre-reform coefficients $\{\theta_t\}_{t<1992}$ in all regressions are statistically indistinguishable from zero. This absence of pre-trends supports the core identifying assumption that farms with different ex-ante crop mixes were on parallel paths prior to the reform. From 1992 onwards, the coefficients become sharply negative and statistically significant. By the late 1990s, a one-standard-deviation increase in exposure leads to a 24% larger decline in a farm's average output price relative to the 1991 baseline. This estimated effect is consistent with the aggregate 33% price decrease shown in Figure A4.

As a complement to this first stage, the second panel of Figure A12 shows the effect of exposure on land subsidies, and indicates that more exposed farms might also be receiving less land subsidies post-reform, even though the effect is noisy. This pattern is suggestive that the variation I am using stems first from heterogeneity in ex ante exposure to subsidies (a crop mix with more protected crop prices), but also to some extent from ex post exposure heterogeneity (a crop mix with lower ex post land subsidies per unit).

The reform also had a significant negative impact on farm performance. The remaining panels of Figure 2 show that a one-standard-deviation increase in exposure led to a 24% decrease in total sales and a 31% decrease in profits by 2002. Finally, I find that a one-standard-deviation increase in reform exposure caused a 24% decrease in total chemical expenditures. Figure A12 shows that farm-level land use also decreases as a result of increased exposure, though the decrease is smaller and noisier than the one of chemicals. This joint but differentiated decrease in input use is indicative that more exposed farms are likely reducing the scale of their production, but might also be changing their input mix composition towards a less chemical-intensive one.

In Appendix C, I discuss the robustness of the results, and show the sales, profit and chemical-use regressions with crop-specific exposure variables to test the homogeneity of the reform effects across crops. I also show results with a difference-in-difference framework, and using

[13] I measure the farms' output price as the weighted average of its crop-specific prices, using relative land allocations as shares. I use the following crops to compute the average (those for which I observe a balance time series): wheat, winter barley, corn, rye, oats, sunflower, colza, spring barley, durum.

the alternative exposure variable which combines 1991 land and output shares.

3.3 Municipality-Level Analysis

To study the reform's consequences on market structure—specifically farm exit and the local distribution of land—I turn to the French Census of Agriculture, which provides me with a consistent panel at the municipality level.[14] My empirical strategy adapts the farm-level design to this repeated municipality-level panel. I measure a municipality's exposure to the reform by taking the average of the farm-level exposure Exp_j within municipality k in the 1988 Census, denoted $\overline{Exposure_k}$. This measure captures the intensity of the reform's shock for the typical farm in a locality. The formal regression specification writes as follows:

$$Y_{kt} = \alpha_k + \eta_{d(k)t} + \sum_{\tau \neq 1988} \theta_\tau \overline{Exposure_k} \mathbb{1}\{t = \tau\} + \sum_{\tau \neq 1988} \delta_\tau X_k \mathbb{1}\{t = \tau\} + \varepsilon_{kt}. \tag{2}$$

where Y_{kt} is the outcome in municipality k and year t. The coefficients of interest $\{\theta_t\}_t$ capture the dynamic effect of local reform exposure. The vector X_k contains a rich set of 1988 municipality-level controls designed to account for initial conditions that could be correlated with both local crop specialization and subsequent consolidation trends. These include moments of the farm size distribution (mean, minimum, maximum and standard deviation), the total municipality-level agricultural area and cultivated area, the minimum cultivated area across farms within the municipality, the average crop count across farms, and the crop count at the municipality level, the average crop evenness across farms, and crop evenness at the municipality level, the fraction of farms in the municipality growing grains in 1988, and the fraction of municipality land allocated to cereals (versus oil crops) in 1988. All regressions include municipality and department-by-year fixed effects, resp. α_k and $\eta_{d(k)t}$, and standard errors are clustered at the department level.

The identifying assumption here is analogous to the farm-level case: conditional on the controls and fixed effects, municipalities with a higher average exposure due to their 1988 aggregate crop patterns would not be on a differential trend for farm exit or land use. In order to test the validity of the design, I show balance tests in Appendix C.5. These indicate that municipalities with above- and below-median reform exposure exhibit parallel pre-reform trends for the market structure outcome, but there is a slight pre-reform divergence in crop diversity. I explicitly account for this by controlling for the 1988 municipality-level evenness of cropland in all specifications.

The event studies in Figure 3 show that the reform triggered significant market consolidation

[14] Note there are about 36,000 municipalities in France. I use data from metropolitan France, and results are robust to the inclusion of Corsica.

driven by the exit of farms. In 2000, a one-standard-deviation increase in a municipality's reform exposure led to a net decrease of 0.7 farms per municipality, growing to a 1.1 decrease by 2010. This same exposure led to a 0.01 unit increase in farm land use evenness (labeled 'farm crop evenness'), which corresponds to 4% of the mean value. Given that small farms tend to have a less even crop mix in the French grain market, this increase could imply that small farms are the ones exiting the market. I show results for additional outcomes in Figure A13. Of note, more exposed municipalities have a 2% increase in the fraction of land allocated to cereals (relative to oil crops). This shift in crop mix parallels the observed change in subsidization across these two categories, as well as the general trends in the number of farms growing each of the two categories shown in Figure A2. Finally, I cannot identify any significant movement in the municipality farm-size distribution (neither for the minimum, average or maximum farm size). A potential interpretation of this result is that farm exit was relatively more driven by pre-reform crop mix composition, than by farm size, and that small farms were sufficiently heterogeneous in their crop mixes.

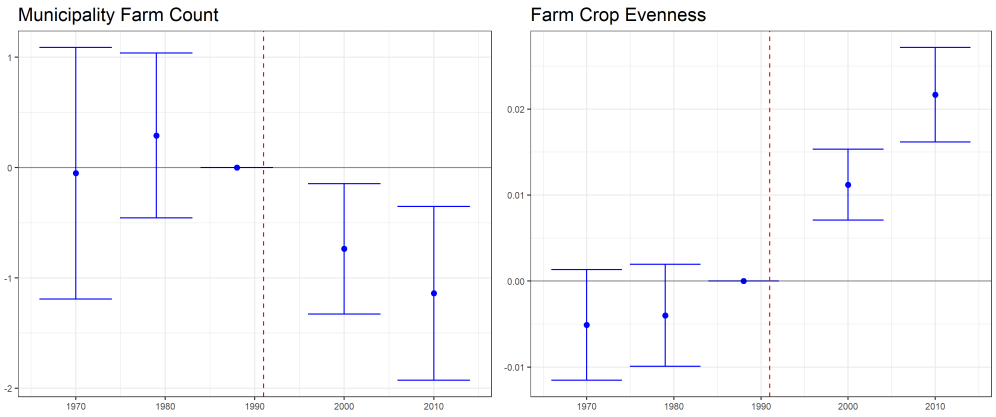


Figure 3: Exposure to the Reform—Municipality-Level Event Study

Notes: This figure gives the results for the municipality-level event study. Outcomes are specified in levels and correspond to: the number of farms operating in the grain market within the municipality and the average farm-level index for the evenness of the distribution of land across crops. The regression includes a series of controls set to their level in 1988 within the municipality, and interacted with time-varying coefficients, as well as department-by-year and municipality fixed effects. Standard errors are clustered at the department level. The table associated with these results is Table A5.

3.4 County-Level Pollution

Finally, I connect the reform directly to environmental outcomes by examining its impact on water pollution. Excessive fertilizer use is a primary cause of algal blooms in waterways, a form of nutrient pollution with significant negative consequences for ecosystems and human health (Rossi et al. (2023)). To measure this, I construct a county-year index of algal blooms using Landsat 5 satellite data, following the methodology of Taylor and Heal (2023), which is further detailed in Appendix B.5. I aggregate the analysis to the county level (approx.

2,000 in France), as municipal boundaries are often too small to reliably overlap with the geographic precision of Landsat 5.

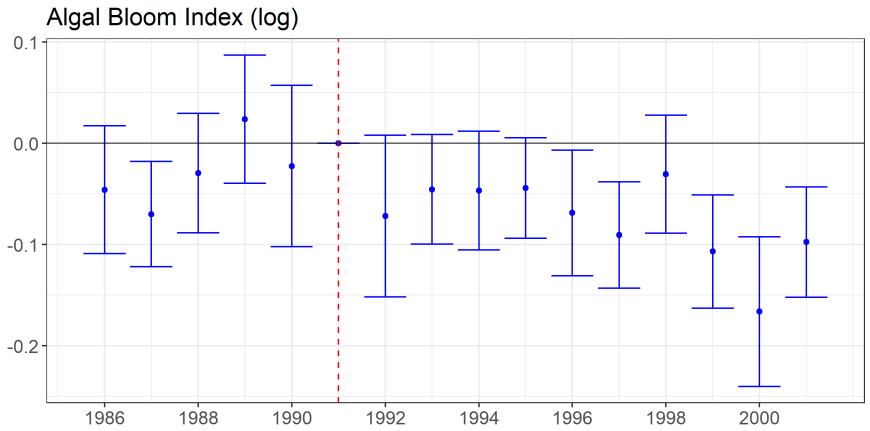


Figure 4: Exposure to the Reform—County-Level Event Study

Notes: This figure gives the results for the county-level event study. The outcome is expressed in log terms, and is a Landsat-5 based index of algal bloom intensity on the within-county water bodies. The regression includes a series of controls set to their level in 1988 within the county (last year of the Census prior to the reform), and interacted with time-varying coefficients, as well as department-by-year and county fixed effects. Standard errors are clustered at the department level. The table associated to these results is Table A6.

The empirical design mirrors the previous analysis. I estimate an event-study model at the county-year level, using the average 1988 farm exposure within the county, $\overline{Exposure_c}$, as the proxy for reform exposure. The specification is analogous to Equation (2) but for county fixed effects instead of municipality ones. I include the same set of 1988 census controls aggregated at the county level, plus controls for local climate patterns at the county level (precipitation, precipitation square, and mean temperature) as these might impact both fertilizer use and algal blooms.[15]

The results show a significant reduction in water pollution in areas more exposed to the reform. By 2001, a one-standard-deviation increase in county-level reform exposure leads to a 10% decrease in the algal bloom index relative to its 1991 level. This result provides evidence that the reform had a positive impact on environmental quality.

[15]To avoid confounding this analysis with the subsequent implementation of the EU Nitrate Directive, I restrict the sample period to end in the year 2001. Chabé-Ferret et al. (2021) analyze the effects of this separate, EU-wide regulation on fertilizer use in France from 2001 onward. The end date ensures the estimates are not contaminated by this later policy.

4 Empirical Model

4.1 Static Demand

I model the respective demands for grain crops $c \in \mathbb{C}$. Farms are price takers, and face a downward-sloping iso-elastic aggregate demand curve for each crop, where Q_{ct}^D is the total quantity demanded and P_{ct} is the market price:

$$Q_{ct}^D = exp(\gamma_c + e_{ct}) P_{ct}^{\beta_c} \ , \ \forall c \in \mathbb{C}. \tag{3}$$

The parameter $\beta_c < 0$ is the constant price elasticity of demand for c. The term $exp(\gamma_c + e_{ct})$ represents a demand shifter, composed of a crop-specific mean intercept γ_c and an idiosyncratic, mean-zero demand shock e_{ct}.

While the supply-side focus of the paper justifies a simple demand structure, Appendix D.5 discusses an alternative model that incorporates monopolistic competition through horizontal differentiation across grain crops, following the framework of Mayer et al. (2014).[16]

4.2 Supply

Farm State: Farms are characterized by an individual state vector Υ_{jt}. Farms are heterogeneous along two dimensions of productivity, both of which are exogenous to the farm's decisions and follow a persistent stochastic process. The first is ω_{jt}^h, a farm-level TFP shock, and the second is ω_{jt}^{ch}, a shock to the farm's chemical use efficiency. Farms are also characterized by their capital stock K_{jt}. Finally, farms are multi-product producers, and do not produce all crops with the same efficiency. A farm's crop mix at t is denoted $\mathbb{C}_{jt} \subseteq \mathbb{C}$. The last element of Υ_{jt} is a time-invariant 'competence ladder' σ_j, which is a farm-specific ranking of crops from most to least efficient. This ladder, drawn upon entry, interacts with the main TFP shock ω_{jt}^h to determine crop-specific productivity. The TFP for crop c is given by $\omega_{jct}^h = \omega_{jt}^h \lambda^{\sum_{n=0}^{dim(\mathbb{C})-1} n \mathbb{1}\{\sigma_j(n)=c\}}$, with n the rank of crop c on farm j's ladder, and $\lambda \in (0,1)$ is a parameter governing the strength of specialization.[17]

Market State: $\Omega_{jt} = \{\mu_t, N_t^e, \{P_{jct}\}_{\mathbb{C}}, \{P_{jt}^x\}_x, Policy_t\}$ is the market state faced by j at t, where μ_t is the measure describing the distribution of incumbents over the farm state space,

[16] This addition allows for price dispersion across producers growing the same crop: high efficiency producers selling a large quantity at a low price, and low efficiency ones selling a small quantity at a high price—in this case, the gains from high production efficiency are larger than in my perfectly competitive model, which hence offers a lower bound for the potential consequences of farm selection. In the monopolistic competition context, a single choke price ensures market clearing across all grains, while in the competitive model a vector of commodity-specific prices is needed, increasing the computational burden of solving for equilibrium. However, my aggregated data is sparse, and likely not sufficient to precisely recover an Ottaviano et al. (2020)-style demand system.

[17] As such, two farms with different TFP values will lose the same percentage of their TFP, when growing a second crop relative to the first one.

N_t^e the mass of entrants, and then vectors of output and input prices, as well as the state of agricultural subsidies $Policy_t$.[18]

Timing: For an incumbent farm, the timing within period t is as follows. First, farms observe their productivity shocks $\{\omega_{jt}^h, \omega_{jt}^{ch}\}$, and the market state Ω_{jt}. Second, they make their static production decisions, i.e. choosing variable inputs to produce quantities Q_{jct} for all crops in their mix, and profits are realized. Third, they make dynamic decisions: they decide whether to stay on the market ($\xi_{jt}^x = 0$), or exit ($\xi_{jt}^x = 1$). If a farm remains on the market, it pays a fixed cost of operation f_k, and chooses its capital stock and crop mix for the next period.

Short-lived potential entrants face an initial entry period before they can start producing. Observing the market state Ω_{jt}, potential entrants decide whether to pay a fixed entry cost f_e to enter the market ($\xi_{jt}^e = 1$). If the entry cost is paid, a farm realizes its initial productivity vector ($\{\omega_{jt}^h, \omega_{jt}^{ch}\}$) and a competence ladder (σ_j). After observing its type, a farm can confirm its entry decision, and make initial investment and crop mix decisions, or definitely leave the market.

The formal Bellman equations for the incumbent and entrant problems, as well as the law of motion for the distribution of firms and the formal definition of equilibrium, are detailed in Appendix D. Diagrams outlining timing assumptions for both entrants and incumbents are also shown in that section.

Production: I assume that profit maximization occurs independently for each crop in a farm's portfolio. This implies that all flexible inputs (land, labor, chemicals) are chosen for, and fully allocated to a specific crop's production line. Capital (K_{jt}) is treated as a shared input, fully available to all crops grown by the farm without rivalry or cost.[19][20] This framework deliberately simplifies away certain economies of scope to focus on the primary mechanisms of input substitution and productivity.[21] The production function for quantity

[18] In the numerical simulation of the model, I will assume that all farms are facing the same input and output prices, and the aggregate market state will be homogeneous across all farms as Ω_t.

[19] I assume that harvest seasons are not perfectly aligned for all crops, and that machinery and buildings can be used for these harvests without rivalry.

[20] The capital law of motion is $K_{jt+1} = (1 - \delta_k)K_{jt} + i_{jt}$, where δ_k is the rate of capital depreciation, and i_{jt} is investment.

[21] The assumption that flexible inputs are fully attributable to production lines is frequent in the literature—implied in single product settings, and formally made in multi-product ones by De Loecker et al. (2016) or Orr (2022). I later discuss how this assumption impacts potential economies of scope in the spirit of Baumol (1977). The assumption that crop-specific profits can be maximized independently from each other captures the assumption of flexible input allocations, of exogenous output prices—unaffected for example by cannibalization concerns of Nocke and Schutz (2018)—and of limited economies of scope. This also assumes away the overhead costs modeled by Foster and Rosenzweig (2022) for agricultural labor, and enforces that all input prices do not depend on the quantity purchased, and that productivity shocks are exogenous to the size and composition of the crop mix.

Q_{jct} of crop c for farm j at time t is Cobb-Douglas in its main variables:

$$Q_{jct} = e^{\omega^h_{jct}} K^{\alpha_{k,c}}_{jt} L^{\alpha_{l,c}}_{jct} \tilde{S}^{\alpha_{s,c}}_{jct} e^{\varepsilon_{jct}}, \tag{4}$$

with L_{jct} as labor, $\sum_{x\in\{k,l,s\}} \alpha_{x,c} = 1$ such that there are decreasing returns to scale for the flexible inputs or $\sum_{x\in\{l,s\}} \alpha_{x,c} < 1$, and ε_{jct} is a mean-zero, i.i.d. shock realized after input decisions are made. The key mechanisms of substitution and input-biased productivity are embedded in the composite input $\tilde{S}_{jct}$, which combines land and chemicals. This composite input is modeled as a constant elasticity of substitution (CES) function, in order to allow for flexible substitution patterns and chemical-biased productivity shocks:

$$\tilde{S}_{jct} = \left\{\delta_{s,c} S^{\rho}_{jct} + (1-\delta_{s,c})\left(e^{\omega^{ch}_{jt}} Chemicals_{jct}\right)^{\rho}\right\}^{\frac{1}{\rho}}, \tag{5}$$

The CES has constant returns to scale, and combines land (S_{jct}) with chemicals ($Chemicals_{jct}$), with an elasticity of substitution $\sigma = \frac{1}{1-\rho}$. The chemical-biased productivity shock (ω^{ch}_{jt}) acts as an augmenting shifter for chemical use: an increase in its value means the farm will be more efficient in its use of chemicals, and achieve a higher output volume with the same input levels. This structure allows the productivity shock to have a non-neutral effect on input choices (Doraszelski and Jaumandreu (2018))).

Efficiency, Chemical Use Intensity and their Correlation: In the counterfactuals, I will be measuring a farm's crop-specific production efficiency using its unit cost of production (c_{jct}), and a farm's crop-level propensity to pollute with its chemical-to-output ratio ($\frac{Chemicals_{jct}}{Q_{jct}}$). I use the unit cost as the measure of efficiency because, in my competitive market without frictions, all farms have the same equilibrium crop-level average cost, making it non-informative about individual farm heterogeneity in production efficiency.[22] It is useful to write their respective equilibrium expressions to understand what will drive the

[22] In my model, farms produce homogeneous goods sold competitively at a unique price. This implies that all farms will produce until their marginal cost equals the output price. Under my model assumptions, the average cost is a constant ratio of the marginal cost, and all farms will have the same average cost of production at equilibrium. The expression for this average cost is derived in Appendix D.1. As such, while farms differ in production efficiency, there are no misallocations across farms, and the average cost of production cannot be used to measure farms' relative production efficiency. The unit cost is then a helpful measure to recover this dispersion across producers.

crop-level correlation between chemical use intensity and production efficiency.

$$c_{jct} = \left(e^{-\omega^h_{jct}} K_{jt}^{-\alpha_{k,c}}\right)^{\frac{1}{\alpha_c}} \left(\alpha_{l,c}^{\prime -\alpha'_{l,c}} \alpha_{s,c}^{\prime -\alpha'_{s,c}}\right) (P^l_{jt})^{\alpha'_{l,c}} \left[\delta^{\sigma}_{s,c} (P^s_{jt})^{1-\sigma} + (1-\delta_{s,c})^{\sigma} (P^{ch}_{jt} e^{-\omega^{ch}_{jt}})^{1-\sigma}\right]^{\frac{\alpha'_{s,c}}{1-\sigma}}$$

$$\frac{Chemicals_{jct}}{Q_{jct}} = \frac{P_{jct}}{P^{ch}_{jt}} \frac{\alpha_{s,c}\big(1-\delta_{s,c}\big)}{\delta_{s,c}\left[\frac{P^s_{jt}(1-\delta_{s,c})e^{\omega^{ch}_{jt}}}{P^{ch}_{jt}\delta_{s,c}}\right]^{\frac{\rho}{\rho-1}} + \big(1-\delta_{s,c}\big)}$$

Where $\alpha_c = \alpha_{l,c} + \alpha_{s,c}$, $\alpha'_{l,c} = \frac{\alpha_{l,c}}{\alpha_c}$ and $\alpha'_{s,c} = \frac{\alpha_{s,c}}{\alpha_c}$. The unit cost depends negatively on the farm's state (TFP (ω^h_{jct}), capital stock (K_{jt}), chemical use efficiency (ω^{ch}_{jt})), as well as on input prices. In contrast, the chemical-to-output ratio is independent of TFP and capital stock. While the unit cost is invariant to production scale, decreasing returns to scale to flexible inputs imply that the chemical-to-output ratio rises with scale and hence with the output price. Finally the ratio also depends on the farm's chemical use efficiency (ω^{ch}_{jt}).

The impact of chemical-use efficiency (ω^{ch}_{jt}) on chemical use intensity is ambiguous. The sign of the effect depends entirely on the elasticity of substitution (σ) between land and chemicals, governed by the parameter ρ. If land and chemicals are complements ($\rho < 0$), an increase in chemical use efficiency will decrease the chemical-to-output ratio. If they are substitutes ($\rho \in (0,1)$), an increase in chemical use efficiency will increase the ratio. This directly implies that at the crop-level, farms that are more efficient will use more chemicals per unit of output if and only if chemicals and land are substitutes. The estimated intensity of this correlation then depends on the relative contribution of ω^{ch}_{jt} versus ω^h_{jct} and K_{jt} (which only affect cost, not the chemical use ratio) to overall unit cost variation.

Aggregating crop-level correlations between production efficiency and chemical use intensity to the farm level presents two main challenges. First, generating a single measure of production efficiency across crops is complex. Because crop-level average costs are equal across farms, differences in overall farm-level average costs are only reflective of variation in the farms' crop mix composition, not in production efficiency. After later imposing that all farms share the same competence ladder σ, I will use total farm profit as a proxy for efficiency, as it will be directly correlated with the farms' underlying production efficiency. Second, the farms' overall chemical-to-output ratio is a composite measure, influenced by both the crop-specific ratios and the overall crop mix composition (which is itself influenced by efficiency, as more efficient farms tend to have more homogeneous crop mixes). The farm-level correlation is thus a compound of these crop-specific and crop-mix effects. While the crop-level correlation between unit cost and chemical-to-output ratio is bound by the value of ρ, the farm-level correlation is further influenced by crop choices, and the crops' relative

reliance on chemicals.

Extensions: Appendix D.3 extends the production function to model pesticides and fertilizers as separate inputs bound in a second CES nest. The coefficients of this extended specification are close to the one found for the simpler design. Additionally, I find that pesticides and fertilizers are complements in production. This implies that their quantities will change in lock-step, and that they can be bundled into a unique chemical input without loss of generality. I also discuss and estimate an alternative joint production function system where inputs are shared publicly across products with a penalty in Appendix D.2.

Dynamics: The final components of the model govern the farm's dynamic decisions. To make the full dynamic equilibrium model computationally tractable, I introduce a few simplifying assumptions at this stage. These assumptions are used to solve the dynamic programming problem and simulating the model, but are not imposed during the more flexible estimation of the production function itself. First, I abstract from farm-level heterogeneity in crop specialization ability by assuming a homogeneous competence ladder ($\sigma_j = \sigma, \forall j$). Second, input and output prices are assumed to be common across all farms in a given period.[23] As such Ω_t is common to all farms.

To capture frictions in investment, I introduce convex costs of capital adjustment: the cost of adjusting capital from K_{jt} to K_{jt+1} via investment i_{jt} is given by:

$$C(i_{jt}, K_{jt}) = P_{jt}^K i_{jt} + C_k^Q K_{jt} \left(\frac{i_{jt}}{K_{jt}} \right)^2.$$

The Role of Dynamics: A static model would fail to properly capture the role of market exit and investment in shaping the industry's response to policy. The dynamic components of the model are essential for two primary reasons. First, dynamics are necessary to correctly model farm exit and market selection. Farms make forward-looking decisions based on their expectations of future productivity. The central variables for these decisions are future productivity which is uncertain, and capital, which is costly to adjust and serves as a buffer insulating farms from negative productivity shocks. Second, these dynamic state variables directly influence the static input choices that determine chemical use. A farm's capital stock affects its optimal scale and its crop mix. These choices, in turn, determine its demand for chemical inputs.[24]

[23] This implies that the market state Ω_t is now also common to all farms. These assumptions simplify the farm's state space considerably.

[24] Capital itself grows slowly over the life time of a French farm, as shown in (Figure A14, something which can only be accurately modeled in a dynamic model. Figure A14 shows the relation between farm size (in land use measured in ha) and their age in years of tenure. I show this relation in the 2010 census. The relationship in prior years is heavily impacted by a trend in entry size growth, which biases the overall size-tenure relation. I measure size in land

Equilibrium: I focus on stationary equilibria where the aggregate market state is constant, $\Omega^\star = \{\mu^\star, N^{e,\star}, \{P_c^\star\}_{\mathbb{C}}, \{P^x\}_x, Policy\}$. The equilibrium is such that, given input prices and the state of agricultural policy, the zero ex ante profit condition holds for entrants, the measure $\mu^\star$ is constant across periods, and the crop markets clear. Appendix D provides the formal definition.

4.3 Estimation & Results

For both estimation and counterfactuals, I aggregate crops into two categories: $\mathbb{C} = \{wheat, other\}$.[25] This aggregation is necessary to recover the production function using only data from single product farms.[26] The estimation is done using the 1980-2006 FADN and the 2000 Census of Agriculture.

4.3.1 Demand

The primary challenge in estimating demand is that market prices are endogenous, affected by unobserved demand shocks. To address this endogeneity, I employ an instrumental variable strategy, using weather shocks as instruments that shift crop supply but are plausibly uncorrelated with demand. This approach has long been used to estimate demand for agricultural commodities (e.g. Roberts and Schlenker (2013)).

My set of instruments corresponds first to national yearly growing and heating degree day realizations (GDD and HDD), for which I detail the construction in Appendix B.4. The second set of instruments corresponds to the national average of department level deviations between yearly GDD and HDD realizations, and their thirty years average. This second set focuses on local weather shocks, which should be uncorrelated to EU-wide demand shocks. Finally, I also include interactions between levels and deviations. The identifying assumption is that these weather variables are orthogonal to aggregate demand shocks for a given crop category. All specifications use national output and price data aggregated from the FADN with sampling weights. They all include a linear time trend to capture smooth changes in demand. I compute Newey-West standard errors to account for autocorrelation in the error terms. The resulting demand elasticity estimates are shown in Table A7.

As a robustness check, I run similar regressions which only include the deviations, and their interaction with level shocks, and obtain similar coefficients (Table A7). This set

use here, as this is the only available relevant variable observed in the Census.

[25] List of crops in *other* category: sunflower, colza, oats, barley, rye, triticale, corn, sorghum and durum. Finer or different aggregations dramatically reduce the number of single-crop farms available for identification and significantly increase the computational burden of solving the model.

[26] This aggregation is also a reasonable representation of the French grain market. As shown in the Census of Agriculture (Table A3), wheat is the dominant crop, and it is the one most frequently grown in monoculture. Even among the two-thirds of farms that are multi-crop, approximately 90% include wheat in their portfolio.

of instruments is less likely to be correlated to aggregate demand shocks through spatial correlation in weather realizations across EU countries.

4.3.2 Production Function

The estimation strategy is designed to address the two main challenges caused by the model specification and the structure of the data: the allocation of inputs for multiproduct farms, and the unobserved productivity shocks that are correlated with input choices. To address the second point, the parameters of the production function are estimated in a multi-step process, and the estimating equations are derived from the first-order conditions of the farm's static profit-maximization problem as in Doraszelski and Jaumandreu (2018).[27]

Input Observability: While crop-specific land use is observed in the data, all other inputs are measured at the farm level. This makes it impossible to directly estimate a crop-level production function for multiproduct firms. To address this, I follow De Loecker et al. (2016) and estimate the production function using data from single-product farms for which the input-to-output allocation is known. I then use the estimated parameters and the model's first-order conditions to recover the unobserved productivity shocks for the full sample of multiproduct farms.[28] Using only single-product farms for estimation introduces a potential selection bias, as these farms are likely systematically different from multiproduct farms.[29] To correct for this, I extend the selection-correction methodology of Olley and Pakes (1996), adapted by De Loecker et al. (2016), for my specific context. The procedure explicitly models the farm's choice to be a single-product producer and includes a correction term in the production function estimation. As argued by De Loecker et al. (2016) and Backus (2020), this bias will affect the coefficients of variables correlated with the threshold rule used for the introduction of a second product. Such variables are only present in the second step of my estimation in the form of the farms' capital stocks and chemical productivity. I therefore include the selection correction only in the second step of my estimation routine.[30]

Additional Considerations: Three final data considerations are worth noting. First, a key advantage of FADN data is that it provides crop-level output in volumes, allowing me to use

[27] Because variable input choices have no dynamic implications in the model, the static first-order conditions are equivalent to those from the full dynamic problem. Independence of crop profits, and the fact that all flexible inputs are fully allocated to product lines finally means that I will be looking at the first order conditions for crop-specific static profit maximization. As frequently done in the production function literature, I assume interior solutions and therefore restrict the estimation sample to observations with non-zero flexible input use (De Loecker et al. (2016)).

[28] A key advantage of my data is that observing crop-specific land allocations allows me to conduct this procedure, while still permitting TFP to be heterogeneous across crops within a farm, an extension to the standard approach.

[29] For instance, if higher TFP makes it easier to manage multiple crops, then the sample of single-product farms will be negatively selected on productivity (Mayer et al. (2014)).

[30] De Loecker et al. (2016) argue one should keep observations related to both always single-product firms, and sometimes single product firms when they happen to be single product. I follow their dataset construction guidelines for all my estimation.

physical quantities in production and avoid the biases associated with using revenue data to measure production (Foster et al. (2008)). Second, expenditures on chemical inputs must be deflated to recover quantities, while I directly observe labor in hours, and land in hectares. I use national-level price indices for this purpose. Store-level data suggests minimal regional price dispersion in France, and no dataset is readily available to build representative regional agricultural input prices.[31] Third, I observe aggregate input bills for both the chemical categories, and cannot address issues of varying consumption basket composition—across farms or across time. Homogeneous input quality updates at the country level might be reflected in price changes. Any upgrade in quality that comes with an increase in price not reflected in the national price index—for example a farm starting to buy better and more expensive products—will be manifested as a higher volume of chemicals used rather than an increase in quality. As shown above, a higher chemical-to-land ratio unexplained by varying prices or a change in crop mix will be classified as an increase in chemical productivity.[32]

First Stage: The first stage aims to estimate the parameters of the CES function ($\{\rho, \{\delta_{s,c}\}_{\mathbb{C}}\}$) and to recover the unobserved chemical-biased productivity shock. The strategy, following Doraszelski and Jaumandreu (2018), uses the combination of the farm's first-order conditions (FOC) for land and chemical choice to derive an expression that inverts for the unobserved productivity shock:

$$\rho\omega_{jt}^{ch} = (1-\rho)log\Big(\frac{Chemicals_{jct}}{S_{jct}}\Big) - log\Big(\frac{P_{jt}^{s}}{P_{jt}^{ch}}\Big) - log\Big(\frac{1-\delta_{s,c}}{\delta_{s,c}}\Big) \tag{6}$$

This equation expresses the unobserved shock ω_{jt}^{ch} as a function of observed data and the structural parameters. To separate the shock from the parameters, I introduce the assumption that ω_{jt}^{ch} follows a first-order Markov process. This allows me to decompose the shock into its conditional expectation and a random innovation ζ_{jt+1}^{ch}, such that: $\omega_{jt+1}^{ch} = \mathbb{E}[\omega_{jt+1}^{ch}|\omega_{jt}^{ch}] + \zeta_{jt+1}^{ch} = g_{ch}(\omega_{jt}^{ch}) + \zeta_{jt+1}^{ch}$. The innovation ζ_{jt+1}^{ch} is, by definition, mean independent of any information available at time t. By substituting Equation (6) into this expression, and approximating g_{ch} with a polynomial, I can form moment conditions. With $A_{jt+1}^{z,ch}$ the matrix of instruments, the parameters are estimated via a generalized method of moments (GMM) using moments of the form: $\mathbb{E}\Big[\zeta_{jt+1}^{ch} A_{jt+1}^{z,ch}\Big] = 0$.

[31] These indices are built using an agricultural price survey. However, an initial inquiry ran in 1995 by Insee established that the geographic dispersion in agricultural input prices in France was too small to warrant a geographic stratification of the survey. As such, the survey only achieves representativeness at the country-level. Labor and land prices are observed separately, and specified at the department-by-year level.

[32] In that sense, the non-Hicksian productivity will not discriminate between input upgrading, learning or changes in farm management that lead to increased chemical use. All these forces will count towards what I measure as increased chemical use efficiency.

Second Stage: The second stage relies on the results from the first stage to estimate the output elasticities ($\{\alpha_k^c, \alpha_l^c, \alpha_s^c\}_{\mathbb{C}}$), and to recover the distribution of TFP shocks. As in the first stage, I can rely on the FOC for labor to derive an expression for the TFP shock:

$$\omega_{jct}^h = p_{jt}^l - p_{jct} - \alpha_{k,c} k_{jt} - log(\alpha_{l,c}) + (1 - \alpha_{l,c}) l_{jct} - \alpha_{s,c} \tilde{s}_{jct}. \tag{7}$$

The standard approach would be to assume the TFP also follows a Markov process, and form moments in a similar way as I did for the first step. However, there is a potential bias caused by the restriction of the data to single-product farms. The decision to grow a second crop likely depends on the farm's state variables, particularly its capital stock and its chemical-use efficiency. Because these variables also appear in the second-stage estimating equation, the selection process induces a correlation between the regressors and the error term, likely biasing the estimates of $\{\alpha_{k,c}, \alpha_{s,c}\}_{\mathbb{C}}$.

To address this bias, I explicitly model the crop selection process and incorporate a correction term, following the logic of De Loecker et al. (2016). The true conditional expectation for TFP must account for the information that the farm possessed when deciding to remain a single-crop operation. I write as Ξ_{jct} the dummy indicating whether farm j only produces c in t. Under the assumption that the TFP shock follows a first order Markov process, I can then account for the truncated TFP process: $\omega_{jct+1}^h = \mathbb{E}[\omega_{jct+1}^h | \omega_{jct}^h, \Xi_{jct+1} = 1] + \zeta_{jct+1}^h$. Appendix D.4 details the correction procedure. With the inclusion of the correction, I can build a matrix of instruments $A_{jct+1}^{z,h}$ orthogonal to the the full residual of the production function composed of the TFP innovation and the idiosyncratic shock to production ($\zeta_{jct+1}^h + \varepsilon_{jct+1}$). I use them to form the moment conditions for the GMM:
$\mathbb{E}\left[\left(\zeta_{jct+1}^h + \varepsilon_{jct+1}\right) A_{jct+1}^{z,h}\right] = 0.$

Estimation: I estimate the parameters for both stages using GMM.[33] The validity of the estimation relies on using a set of instruments that are orthogonal to the productivity innovations and random shocks, while being correlated to flexible input choices.

Instruments for the first stage must be correlated with the choice of chemicals and land, but uncorrelated with the chemical-use productivity innovation. These will correspond to variables realized prior to the innovation, or taken to be exogenous to the farm-level innovation process such as France-wide agricultural policy. I use: a constant (innovation is assumed to be mean zero), pre-determined farm state variables (current and lagged capital stocks), lagged input choices which are correlated with current choices due to persistence but uncorrelated with the current innovation (the lag log ratio of chemicals used to land used), and

[33] Following the approach of Wooldridge (2010), the parameters of the second order polynomial functions used to approximate the conditional expectation of productivity ($\{g_{ch}(.), g_h(.)\}$) are absorbed and not estimated directly.

exogenous factor prices and policy variables (the lag log ratio of chemical and land prices, the current land subsidy, and a measure of current-farm exposure to EU agricultural policy using lagged crop land shares as weights following my shift-share instrument from the reduced-form section).

Instruments for the second stage must be correlated with the farm's input decisions at $t+1$ but uncorrelated with the full residual $(\zeta^h_{jct+1} + \varepsilon_{jct+1})$. I use: pre-determined farm state variables (current and lagged capital stock), lagged input and output choices (lagged labor, lagged composite land input, the lag log ratio of chemicals used to land used) and finally exogenous factor prices and policy variables (lag wages, lag land prices, the lag log ratio of chemical and land prices, the lagged farm output price index, current land subsidy, and the same measure of current-farm exposure to EU agricultural policy).[34]

Sensitivity of Results to Moments: To understand how specific moment conditions identify the production function parameters, I implement the sensitivity analysis of Andrews et al. (2017). This procedure quantifies the influence of each moment on each parameter estimate, highlighting which exclusion restrictions are most critical for identification. The analysis reveals that the elasticity of substitution between chemicals and land is identified both by moments related to the assumptions on the productivity process and its innovation, by moments related to the chemical-to-land lagged price and quantity ratios, and by the two subsidy instruments (variation in the land subsidy, and farm exposure to the current structure of CAP subsidies). The procedure provides some clear economic intuition: for the same lag input price ratio, a higher lag chemical-to-land quantity ratio would be associated with more substitutability between chemicals and land. Derivations and results are shown in Appendix F.

4.3.3 Dynamic Parameters

Three key parameters still need to be estimated: the fixed cost of incumbency (f_k), the convex cost of capital adjustment (C^Q_k) and the fixed cost of entry (f_e). Together, these parameters determine firm survival thresholds, investment behavior, and the equilibrium market size. I estimate these parameters using a Simulated Method of Moments (SMM) procedure (Pakes and Pollard (1989); Gouriéroux et al. (1993)).

The algorithm proceeds as follows: for a vector of potential parameters $\Theta = \{f_k, C^Q_k, f_e\}$, I solve for the post-MacSharry stationary equilibrium of the market, and obtain the farms' policy functions for investment, crop choice and exit.[35] I then simulate a panel of farms,

[34] I use Amemiya (1974) GMM weights and a Nelder-Mead minimization algorithm—searching over the space of starting parameters to ensure homogeneous convergence to these values.

[35] Specifically, I solve for an equilibrium with a 32% land subsidy.

drawing their initial type from the equilibrium distribution. I use this panel to compute a vector of moments $\Phi^s(\Theta)$. The parameter estimates are those that minimize the distance between these simulated moments and their empirical counterparts (Φ^d) calculated from the FADN and Census data, where W is the weighting matrix:[36]

$$\min_{\Theta} \left[\Phi^d - \Phi^s(\Theta)\right]' W \left[\Phi^d - \Phi^s(\Theta)\right], \tag{8}$$

Following Gouriéroux et al. (1993), identification requires that the function mapping parameters to moments ($\Phi^s(.)$) is injective (a relevance condition), and that the simulation procedure provides a consistent estimate of this function. The first condition is about the relevance of the chosen moments, shown in Table 1. I use of bootstrapping and a large simulated panel to ensure that the second condition holds.[37][38]

Table 1: Targeted and Untargeted Moments

Moments for Estimation	Moment	Observed Value	Simulated Value
Auto-Correlation Investment		0.11	0.11
Correlation Investment-TFP		0.050	-0.046
Coef. Variation in Tenure		0.64	0.82
Coef. Variation in Profits		0.75	0.80
S^2 (Measure of Fit)			0.96
Untargeted Moments	Moment	Observed Value	Simulated Value
Price (wheat)		158 €/t	164 €/t
Price (other)		184 €/t	242 €/t
Supply (wheat)		36.8 million t	41.4 million t
Supply (other)		27.5 million t	22.6 million t
Chemical Expenditure to Revenue Ratio		27.5 %	23.6 %

Notes: The coefficient of variation in tenure is computed in the 2000 Census of Agriculture. All other targeted moments are computed using the FADN survey over the 1995-2006 period. The auto-correlation in investment tracks the correlation in investment rates within farms over time, the correlation between investment and TFP the correlation between farm investment and productivity levels. The coefficients of variation in tenure and profits capture the spread in farm tenure and

[36] Following Asker et al. (2014), I use the identity matrix for W as the targeted moments are similarly scaled. I proceed to the estimation in two steps: I first run a random grid search algorithm, and then use a simulated annealing algorithm to refine the estimation, starting at the optimum located by the grid search, and setting bounds for the search also derived from the results of the grid search.

[37] To account for the influence of initial draws on estimation, in this model with endogenous exit, I bootstrap this entire simulation B=100 times, drawing new entrants from the equilibrium entrant distribution to maintain a constant sample size. In a model without exit (Cooper and Haltiwanger (2006); Asker et al. (2014)), removing initial simulation periods is sufficient to address the initial conditions problem. Here, new entrants will keep being added to the sample and might represent a significant portion of the sample. As such, the bootstrap provides a more fitting solution to the initial conditions problem. Averaging moments across bootstrap replications provides a robust alternative.

[38] I draw a panel of 1,000 farms, for T=50 periods.

profits within the market at equilibrium. [39]

While all moments are used to estimate all parameters, the relevance condition is best discussed by going through the respective use of each moment in estimating the targeted parameters.[40] The within-farm autocorrelation of investment, and the correlation between investment and farm productivity (proxied by their TFP) are primarily informative about capital adjustment cost. Without adjustment costs, investment would track productivity shocks closely, and have little serial auto-correlation. Convex costs also force firms to smooth their investment over time, creating higher autocorrelation and weakening the contemporaneous link to productivity. The cross-sectional coefficients of variation in profits and farm tenure are primarily informative about the fixed cost of operation. A higher fixed cost truncates the profit distribution from below and raises the productivity threshold for survival, leading to shorter average tenures and a less dispersed tenure distribution. Under the zero-profit condition, the fixed cost of entry equates the expected value of entry, which depends both on expected tenure and expected profit. These moments capture the intensity of selection in the market. The first three moments are measured in the FADN, while the coefficient of variation in farm tenure is taken from the Census. I compare the values of simulated and observed targeted and un-targeted moments in Table 1 as a measure of model fit.

4.3.4 Estimation Results

I present the main parameter estimates in Table 2. Table A8 provides the values for the additional model parameters that I calibrate directly from the FADN or external sources.

From $\rho = 0.32$, I recover an elasticity of $\sigma = 1.47$. As derived in Section 4.2, this elasticity signs the crop-level correlation between farm efficiency and chemical use intensity: because land and chemicals are substitutes, farms that are more productive and have lower unit costs will also have higher chemicals-per-output ratios. This correlation is retained at the farm level, where the more efficient farms, which have higher overall profits, are also using more chemicals per unit of output. This implies that the infra-marginal incumbent farms, whose incumbency decisions will not be impacted by subsidies, use on average more chemicals per output than the marginal incumbent ones. These correlations are plotted in Section 5.1. [41]

I also obtain a yearly fixed cost paid by incumbents of $6,640€$, Industry estimates of administrative costs for the average wheat farm in my sample place these around $9,400€$ (Guillermet

[39] The S^2 measure corresponds to: $S^2 = 1 - \frac{(x-\hat{x})'(x-\hat{x})}{x'x}$, with x the vector of observed moments and $\hat{x}$ the vector of predicted moments, replicating the measure of fit of Asker et al. (2014) used in a similar exercise.

[40] Their distribution is shown in Figure A16.

[41] This estimate aligns well with the literature on agricultural technical change. Land-chemical substitutability is consistent with the broad history of post-WWII agriculture, where technological progress was characterized by new, high-yield seed varieties that substituted away from land by requiring more intensive chemical use (Hayami and Ruttan (1971)). The magnitude of my estimate is comparable to other key elasticities found in agriculture, such as the elasticity of 1.71 between labor and capital estimated for U.S. agriculture by Kislev and Peterson (1982).

(2015)). The fixed cost of entry is estimated to be 58,100€. Industry estimates set the initial investment value to enter the agricultural market to be around 250,000€, but these include initial capital investments in machinery, while the farms in my model have to pay a cost of entry on top of their initial capital investment, which sum to a comparable number.[42]

Table 2: Estimated Model Parameters

Coefficient	Parameter	Estimate	Std. Error
Demand Parameters			
Demand constant (wheat)	γ_{wheat}	22.08	(1.69)
Demand constant (other)	γ_{other}	20.10	(1.97)
Demand elasticity (wheat)	β_{wheat}	-0.89	(0.27)
Demand elasticity (other)	β_{other}	-0.58	(0.32)
Production Function			
Substitution Land-Chemicals	ρ	0.32	(0.15)
Labor Share (Wheat)	$\alpha_{l,wheat}$	0.29	(0.20)
Labor Share (Others)	$\alpha_{l,other}$	0.49	(0.18)
Land-Nest Share (Wheat)	$\alpha_{s,wheat}$	0.17	(0.14)
Land-Nest Share (Others)	$\alpha_{s,other}$	0.29	(0.14)
Land Share (Wheat)	$\delta_{s,wheat}$	0.35	(0.19)
Land Share (Others)	$\delta_{s,other}$	0.40	(0.28)
Capital & Fixed Costs			
Convex Adjustment Cost	C_k^Q	2.50	(0.44)
Fixed Cost of Incumbency (€1K)	f_k	6.64	(2.91)
Fixed Cost of Entry (€1K)	f_e	58.1	(11.1)

Notes: Demand parameters come from IV regressions, and standard errors are corrected for auto-correlation using the Newey-West procedure. Production function parameters are obtained from a two-step GMM estimation. Both estimations are done on the FADN French sample restricted to 1980-2006. Prior to 1980 the FADN does not contain output price data, while years post-2006 come with significant changes in EU policies not modeled in this paper. I keep farms observed for at least three periods in a row[43], and that produce either only wheat or the other aggregated crop. I remove farms that are not observed with positive input values for the set of considered inputs (land, labor, capital, fertilizers, pesticides). Standard errors are obtained with a block bootstrap, at the farm level, using $B = 1000$. Capital and fixed costs parameters are obtained with a simulated method of moments. Because the production function estimation is run on a protected server containing the administrative data, while the numerical model is solved on a high-performance computing cluster, standard errors for the simulated method of moments do not account for estimation error in prior steps.

[42] See https://www.agri-france.com/reprendre-une-ferme/.

[43] This is done to smooth potential measurement issues, and follows from De Loecker et al. (2016).

5 Pollution-Inclusive Welfare Evaluation of Agricultural Subsidies

5.1 Which Channels Relate Subsidies to Chemical Use?

The empirical model allows for three different channels to impact chemical use on the market: 1) an intensive margin of input choice, 2) an intensive margin of crop choice, 3) an extensive margin of farm selection. For the first, subsidy-induced shifts in relative input prices will change input choices. For the second, changes in both input and output prices will change crop choices. For the third, changes in the expected profitability of incumbency will impact the marginal incumbents' exit decisions. Finally, equilibrium changes in output prices will change the farms' scale of production and their crop-level chemical use per unit of output.[44]

In this sub-section, I compare a budget equivalent series of two different farm subsidy designs, in order to highlight how subsidies can interact with the aforementioned channels. The first of these two subsidies is the land subsidy introduced post-MacSharry reform and still in place in the EU today. The second is a lump sum payment to low chemical use efficiency producers who cannot survive on the market in the no-intervention equilibrium, i.e., assuming the government can observe farm types, a subsidy scheme where the government pays a yearly lump sum to farms whose ω_{jt}^{ch} is sufficiently low. These producers have lower chemical-per-output ratios.

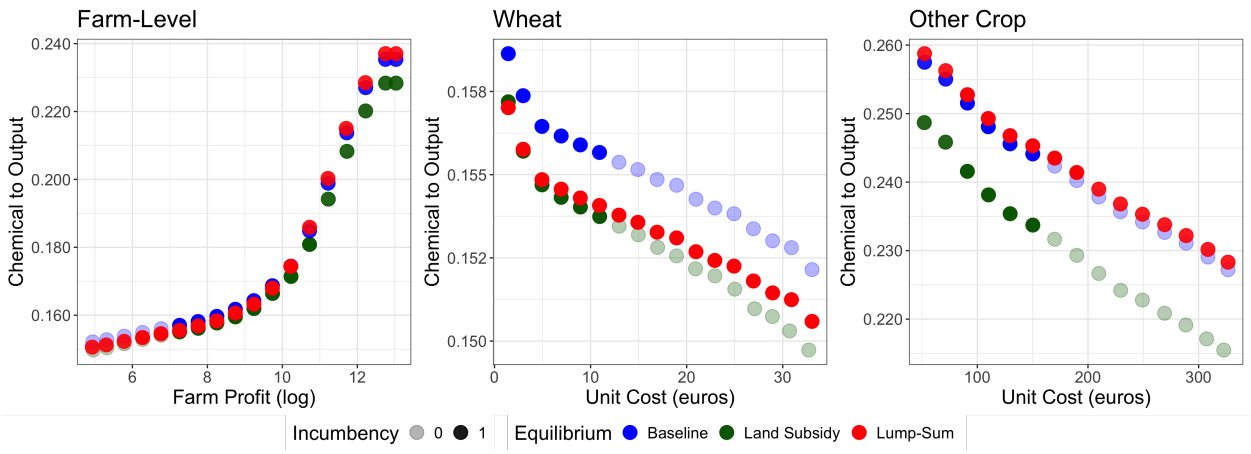


Figure 5: Intensive & Extensive Policy Effects

Notes: This figure plots equilibrium correlations between farm production efficiency and chemical use intensity, under different subsidy schemes. Production efficiency is measured at the crop-level using unit costs, and proxied at the farm level with total farm profit. Chemical use intensity is measured as the crop-specific and farm-level ratio of chemicals per quantity produced. Shaded dots indicate farms that are not incumbents under a given equilibrium.

The correlations between production efficiency and chemical use intensity are illustrated in Figure 5 for the no-intervention equilibrium, as well as for a 30% land subsidy, and a 22k€ lump sum payment. The shaded dots on the figures identify farms that are not profitable

[44] See derivations in Section 4.2.

enough to survive in a given equilibrium, and changes in shading across policies illustrate their selection effects.

Figure 5 summarizes how subsidies impact chemical use at equilibrium. The land subsidy has no impact on selection, as the shaded section of the farm profit distribution remains similar to the baseline equilibrium.[45] It however plays along the input choice margin by reducing the chemical use intensity of all farms: for the same profit or unit cost on Figure 5, the green land subsidy line is always lower than the blue baseline line.

On the contrary, the lump sum plays along the margin of farm exit: allowing for low productivity but also low pollution farms to remain on the market, and shifting more production towards these producers. This is illustrated by the extension of the set of incumbents (non-shaded dots) on Figure 5. By increasing expected profits at entry, this subsidy also significantly lowers the price of wheat, while keeping the price of the other crop relatively constant. This implies a reduction in production scale for wheat, and a decrease in the chemical-to-output ratio used by wheat producers. The chemical-to-output ratio for the other crop is left mostly unchanged. On Figure 5, for a given unit cost, the red line is lower than the blue for wheat, but relatively similar to the blue line for the other crop.

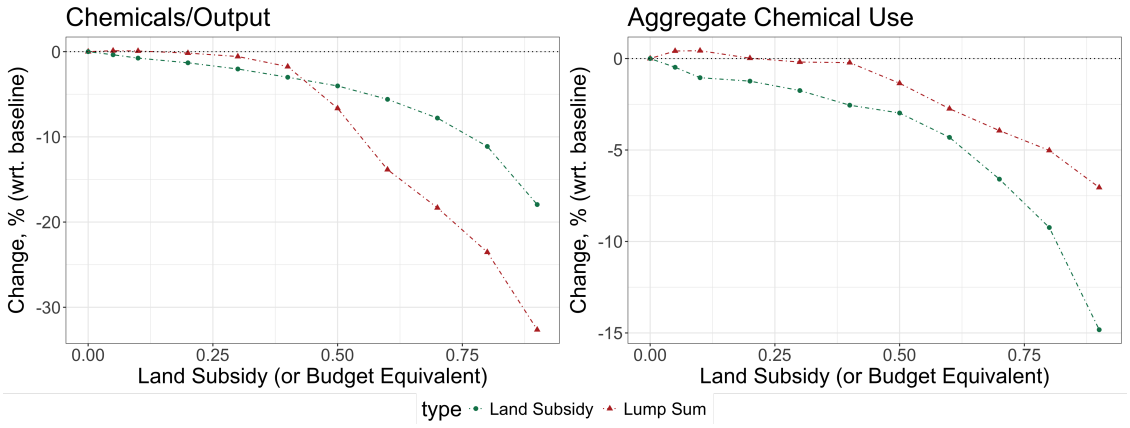


Figure 6: Chemical Use Outcomes

Notes: These figures describe equilibrium chemical use outcomes under a series of respectively increasing land subsidies, and budget-equivalent lump sum payments paid to low chemical use efficiency producers who cannot survive under the baseline no-intervention equilibrium. The first figure tracks the intensity of chemical use measured with the chemical use to output ratio. The second tracks total chemical use.

In summary, the land subsidy has a negative—homogeneous across crops and across farms—effect on the chemical-to-output ratio, but no selection effect. The lump sum payment has a negative—asymmetric across crops, and homogeneous across farms—effect on the chemical-to-output ratio, and a large selection effect.

[45] The land subsidy's cost reduction is partially passed through to consumers via lower equilibrium output prices. This price drop offsets much of the subsidy's direct benefit for marginal farms, thus keeping expected profit at entry and market selection relatively constant.

Figure 6 illustrates equilibrium chemical use under these two subsidy designs, and for different levels of government expenditure. The first figure of Figure 6 shows the subsidies' impact on the intensity of chemical use, and the second on total chemical use. Figure 1 of Figure 6 shows that for large levels of intervention, the lump sum payment achieves a larger reduction in the intensity of chemical use than the budget-equivalent land subsidy. However, figure 2 shows that land subsidies always lead to a larger decreases in total chemical use, as they have a smaller positive effect on aggregate production levels. They reach a 15% decrease wrt. no intervention for a 90% land subsidy, compared to a 7% decrease with the equivalent lump sum payment.[46]

5.2 Optimal Policy Design

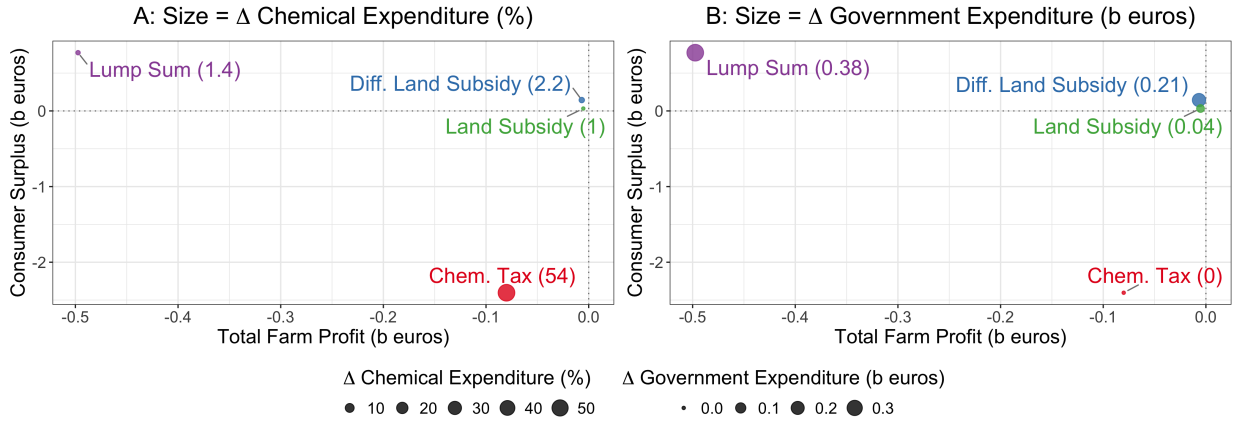


Figure 7: Comparison of Welfare Maximizing Policies

Notes: This figure plots the welfare implications of different optimal subsidy and taxes. Within each design, a grid search is performed over the range of possible intervention levels, and the welfare optimizing level is selected. The welfare criterion used for this is composed of changes in consumer surplus, total farm profit and environmental gains from reduced chemical use minus government expenditure. The marginal cost of agricultural chemical pollution is taken from CGDD (2011). The figure plots the difference in consumer surplus, total farm profit and government expenditure relative to the no-intervention equilibrium. The change in chemical expenditure is expressed in percent, also compared to the no-intervention equilibrium.

In this last counterfactual exercise, I compare the welfare implications of a series of agricultural subsidies and of a chemical tax. Conditional on a valuation for the damages caused by agricultural chemical use, I perform a grid search that identifies the optimal subsidy or tax level within each design, and then compare these optimal interventions across designs.[47] I use this comparison to ask two questions: how do the environmental implications of these

[46] In Appendix E, I use an accounting equivalence to precisely decompose these changes into channel specific changes.

[47] The valuation for chemical damages used in this exercise is the lowest of the values proposed by CGDD (2011). Figure A17 shows results taking the valuation for chemical pollution to be equal to the 2025 French marginal tax rate for glyphosate, and applying it to all chemical expenditure in the market. This yields a tax rate of 0.14 euros per euro of chemical expenditure. The main valuation chosen in the paper accounts for all types of agricultural chemicals, and comes from a French government report aimed at estimating total environmental damages from agricultural chemical use. As such, this glyphosate derived valuation serves as a counterfactual low-valuation scenario.

subsidies change one's understanding of their optimality (within and across designs), and conditional on a valuation of environmental pollution, what is the extent to which these subsidies can substitute for the direct taxation of chemicals? Figure 7 plots the welfare implications of these optimal subsidy and tax levels.

Using a valuation for the damages caused by agricultural chemical pollution from CGDD (2011) of 828 euros per ha, I find that both a positive level of land subsidy, land subsidy differentiated by crop, and of a lump payment to low chemical use efficiency producers can generate welfare gains compared to the no-intervention equilibrium. My search identifies an annual 28k€ lump sum payment to sustainable producers, a homogeneous 10%, and a differentiated (10% for wheat, 40% for others) land subsidy schemes as welfare optimal within their design. Their generated welfare gains (sum of gains to consumer surplus, total farm profit and the environment minus government expenditure), and their marginal value of public funds (MVPF) defined here as the sum of gains (to consumer surplus, total farm profit and the environment) divided by government expenditure, are resp. of 119M€ and MVPF of 1.6 for the differentiated land subsidy, 75M€ and 2.9 for the homogeneous land subsidy, and 9.6M€ and 1.03 for the lump sum payment.

Out of these three designs, the differentiated land subsidy does best overall, reaching the largest decrease in chemical use (2.2% compared to baseline), while also increasing consumer surplus slightly. Compared to the homogeneous land subsidy, it subsidizes land use for the crop whose production function land-chemical nest coefficient is largest ($\alpha_{S,other} = 0.29$ vs. $\alpha_{S,wheat} = 0.19$), and for which substitution will be easiest. As a consequence, it reaches higher decreases in chemical use at a relatively constant cost to producers. The lump sum payment does best for consumer surplus, but at a relatively high cost for producers.

In a case where environmental pollution is not valued, only a small homogeneous 5% land subsidy can generate positive, albeit very small, welfare gains. No level of lump sum payment to producers can generate positive welfare gains. Within the set of previously identified optimal policies, the optimality ranking is also shifted, with the homogeneous land subsidy becoming the dominating subsidy. In that sense, the unintended environmental consequences of these agricultural subsidies are meaningful, capable of overturning their welfare effects, and changing our understanding of which subsidies should be optimally implemented. This comparison does not account for additional goals of subsidies which could be added to the welfare criterion, for example a reduction in agricultural profit dispersion.

A second natural question is whether an optimal subsidy can be an efficient substitute to a direct tax of chemicals. To isolate the tax's distortionary effect from its large income effect on producers, I assume all tax revenues are recycled back to farms via a direct, lump-sum transfer. This revenue-neutral policy will provide a lower bound for the negative conse-

quences of the tax on farm profit. Figure 7 plots the optimal chemical tax next to the optimal subsidies. This comparison yields a straightforward take-away: a chemical tax is significantly more efficient at reducing total chemical use: the optimal tax reduces chemical expenditure by 54%, compared to 2.2% for the optimal subsidy. This reduction is achieved while generating significant tax revenue (which, in my scenario, is recycled to producers and results in zero net government expenditure). For the chosen valuation of agricultural chemical pollution, the tax reaches a 2.1 billion euros welfare gain. However, this tax achieves these gains at a significant cost, both for consumers and producers. Indeed, it significantly raises output prices and decreases produced quantities. Within this framework, subsidies are poor substitutes for the direct taxation of chemicals. Yet, taxation has a significant negative effect on economic surplus measured as consumer surplus and total farm profit, which is avoided by subsidies.

I perform one additional counterfactual analysis in Appendix G, where I simulate the effects of the policy design change at the heart of the MacSharry reform, moving from government purchasing to a land subsidy.

6 Conclusion

In this paper, I document the equilibrium effects of the EU Common Agricultural Policy subsidies on chemical use and environmental pollution within the French grain market.

To establish these results, I first use the 1992 MacSharry CAP reform as a natural experiment. A shift-share analysis demonstrates that the reform significantly decreased the profitability of French agriculture, led to the exit of smaller specialized farms, and to a reduction in both chemical use and water pollution.

To better understand the equilibrium consequences of agricultural subsidies, I then develop an empirical model of the French grain market, and show that there is a positive correlation between farm production efficiency and chemical use intensity. I use this model to analyze the welfare effects of a series of counterfactual policies. My results show that both subsidies supporting the survival of low-pollution producers, and subsidies aimed at shifting the relative price of chemicals and land, can be effective at reducing both the intensity of chemical use, and aggregate chemical use. Finally, I show that while these subsidies have meaningful environmental externalities which can transform our understanding of their welfare consequences, they remain poor substitutes for a direct tax on chemicals, which achieves far greater environmental gains, albeit at a significant cost to both producer and consumer surplus.

References

Adamopoulos, T. and D. Restuccia (2014): "The Size Distribution of Farms and International Productivity Differences," *American Economic Review*.

Amemiya, T. (1974): "The Nonlinear Two-Stage Least-Squares Estimator," *Journal of Econometrics*.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*.

Armitage, S. (2023): "Technology Transitions and the Timing of Environmental Policy: Evidence from Efficient Lighting," *Working Paper*.

Asker, J., A. Collard-Wexler, and J. De Loecker (2014): "Dynamic Inputs and Resource (Mis)Allocation," *Review of Economic Studies*.

Aspelund, K. M. and A. Russo (2024): "Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation Auctions," *Working Paper*.

Backus, M. (2020): "Why is Productivity Correlated with Competition?" *Econometrica*.

Baumol, W. J. (1977): "On the Proper Cost Tests for Natural Monopoly in a Multiproduct Industry," *American Economic Review*.

Binswanger, H. (1974): "The Measurement of Technical Change Biases with Many Factors of Production," *American Review of Economics*.

Blundell, W., G. Gowrisankaran, and A. Langer (2020): "Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations," *American Economic Review*.

Borusyak, K., P. Hull, and X. Jaravel (2022): "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*.

Bovenberg, L. A. and L. H. Goulder (1996): "Optimal Environmental Taxation in the Presence of Other Taxes: General-Equilibrium Analysis," *American Economic Review*.

Burlig, F., L. Preonas, and M. Woerman (2024): "Groundwater and Crop Choice in the Short and Long Run," *Working Paper*.

Bustos, P., B. Caprettini, and J. Ponticelli (2016): "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *American Economic Review*.

Carlton, D. W. and G. C. Loury (1980): "The Limitations of Pigouvian Taxes as a Long-Run Remedy for Externalities," *Quarterly Journal of Economics*.

CGDD (2011): "Assessing Water Pollution Costs of Farming in France," *Studies and Documents, Department of the Commissioner General for Sustainable Development*.

Chabé-Ferret, S., A. Reynaud, and E. Tene (2021): "Water Quality, Policy Diffusion Effects and Farmers' Behavior," *Working Paper*.

Clemens, M. A., E. G. Lewis, , and H. M. Postel (2018): "Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion," *American Economic Review*.

Cooper, R. W. and J. C. Haltiwanger (2006): "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*.

Davis, L. W. (2017): "The Cost of Global Fuel Subsidies," *The Energy Journal*.

De Loecker, J. and A. Collard-Wexler (2015): "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016): "Prices, Markups, and Trade Reform," *Econometrica*.

Dhyne, E., A. Petrin, V. Smeets, and F. Warzynski (2022): "Theory for Extend-

ing Single-Product Production Function Estimation to Multi-Product Settings," *NBER Working Paper*.

Dias, M., R. Rocha, and R. R. Soares (2023): "Down the River: Glyphosate Use in Agriculture and Birth Outcomes of Surrounding Populations," *Review of Economic Studies*.

Doraszelski, U. and J. Jaumandreu (2018): "Measuring the Bias of Technological Change," *Journal of Political Economy*.

Englander, G., J. Zhang, J. C. Villasenor-Derbez, Q. Jiang, M. Hu, O. Deschenes, and C. Costello (2023): "Input Subsidies and the Destruction of Natural Capital: Chinese Distant Water Fishing," *NBER Working Paper*.

European Court of Auditors (2024): "Common Agricultural Policy Plans: Greener, but not matching the EU's ambitions for the climate and the environment," *Special Report*.

FAO (2024): "Government expenditures in agriculture 2001–2022 – Global and regional trends," *FAOSTAT Analytical Briefs*.

Foster, A. D. and M. R. Rosenzweig (2022): "Are there too many Farms in the World? labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size," *Journal of Political Economy*.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*.

Frank, E. (2021): "The Economic Impacts of Ecosystem Disruptions: Costs From Substituting Biological Pest Control," *Working Paper*.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik Instruments: What, When, Why and How," *American Economic Review*.

Gollin, D., D. Lagakos, and M. E. Waugh (2013): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*.

Goulder, L. H. and I. W. Parry (2008): "Instrument Choice in Environmental Policy," *Review of Environmental Economics and Policy*.

Gouriéroux, C., A. Monfort, and E. Renault (1993): "Indirect Inference," *Journal of Applied Econometrics*.

Gowrisankaran, G., A. Langer, and M. Reguant (2024): "Energy Transitions in Regulated Markets," *Working Paper*.

Grupp, T. E., P. Mishra, M. Reynaert, and A. A. van Benthem (2024): "An Evaluation of Protected Area Policies in the European Union," *Working Paper*.

Guillermet, J. (2015): "Couts de Production en Grandes Cultures," *Technical Report*.

Hayami, Y. and V. W. Ruttan (1971): *Agricultural Development: An International Perspective*, John Hopkins University Press.

Hopenhayn, H. (1992): "Entry and Exit in Long Run Equilibria," *Econometrica*.

Hsiao, A. (2022): "Coordination and Commitment in International Climate Action: Evidence from Palm Oil," *Working Paper*.

Kislev, Y. and W. Peterson (1982): "Prices, Technology, and Farm Size," *Journal of Political Economy*.

Larsen, K. (2025): "The Environmental Value of Private Land Conservation: The Role of Conservation Easement Tax Incentives," *Working Paper*.

Livingston, M., M. Roberts, and J. Rust (2008): "Optimal Corn and Soybean Rotations," *Working Paper*.

Mayer, T., M. J. Melitz, and G. I. Ottaviano (2014): "Market Size, Competition,

and the Product Mix of Exporters," *American Economic Review*.

Missirian, A. (2020): "Yes, in your backyard: Forced technological adoption and spatial externalities," *Working Paper*.

Nocke, V. and N. Schutz (2018): "Multiproduct Firm Oligopoly: An Aggregative Games Approach," *Econometrica*.

Oates, W. (1991): "Pollution Charges as a Source of Public Revenues," *Springer-Verlag Economic Progress and Environmental Concerns*.

Oberfield, E. and D. Raval (2021): "Micro Data and Macro Technology," *Econometrica*.

Obolensky, M. (2025): "Protect or Prepare? Crop Insurance and Adaptation in a Changing Climate," *Working Paper*.

Olley, G. S. and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*.

Orr, S. (2022): "Within-Firm Productivity Dispersion: Estimates and Implications," *Journal of Political Economy*.

Ottaviano, G., T. Mayer, and M. Melitz (2020): "Product Mix and Firm Productivity Responses to Trade Competition," *Review of Ecoonomics and Statistics*.

Pakes, A. and D. Pollard (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*.

Rafey, W. (2023): "Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation," *American Economic Review*.

Roberts, M. and W. Schlenker (2013): "Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate," *American Economic Review*.

Rossi, D., G. Hoque, Y. Ji, and C. Kling (2023): "The Economics of Nutrient Pollution from Agriculture," *Annual Review of Resource Economics*.

Scott, P. (2013): "Dynamic Discrete Choice Estimation of Agricultural Land Use," *Working Paper*.

Shapiro, J. S. (2021): "The Environmental Bias of Trade Policy," *The Quarterly Journal of Economics*.

Taylor, C. and G. Heal (2023): "Fertilizer and Algal Blooms: A Satellite Approach to Assessing Water Quality." *NBER Risks in Agricultural Supply*.

Taylor, C. A. (2022): "Cicadian Rhythm: Insecticides, Infant Health, and Long-term Outcomes," *CEEP Working Paper*.

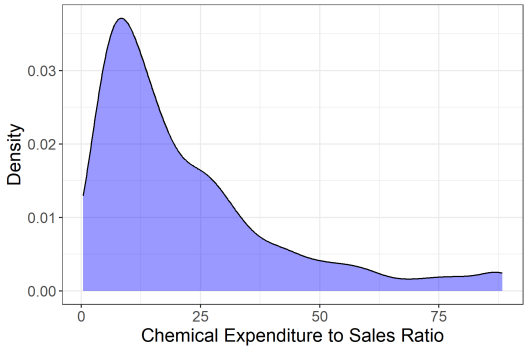
Tschofen, P., I. L. Azevedo, and N. Z. Muller (2019): "Fine particulate matter damages and value added inthe US economy," *PNAS*.

Wooldridge, J. M. (2010): "Econometric Analysis of Cross Section and Panel Data," *MIT Press, Cambridge*.

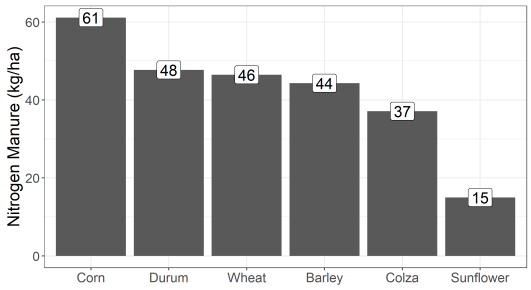
Supplemental Appendix:
An Empirical Model of Agricultural Subsidies with Environmental Externalities

A Additional Figures and Tables

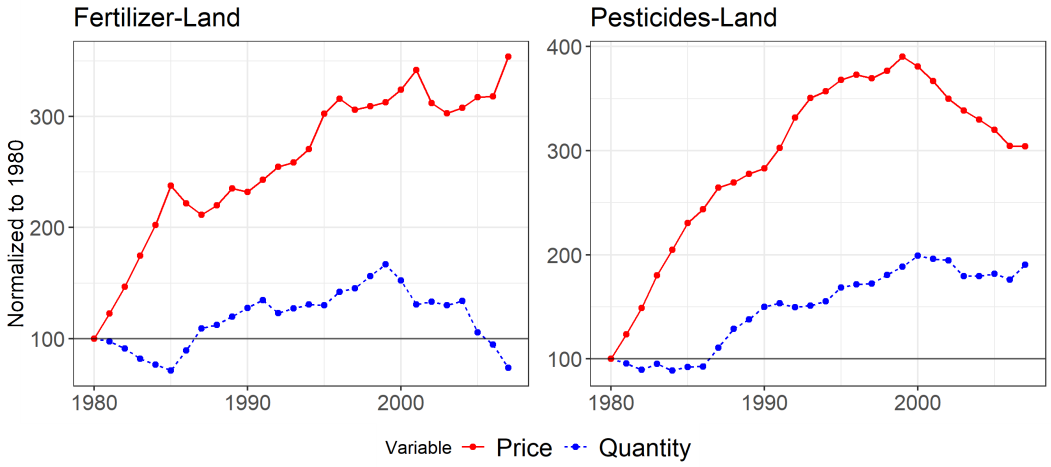
A.1 Figures



(a) Heterogeneity in Chemical Use (1985-1990)



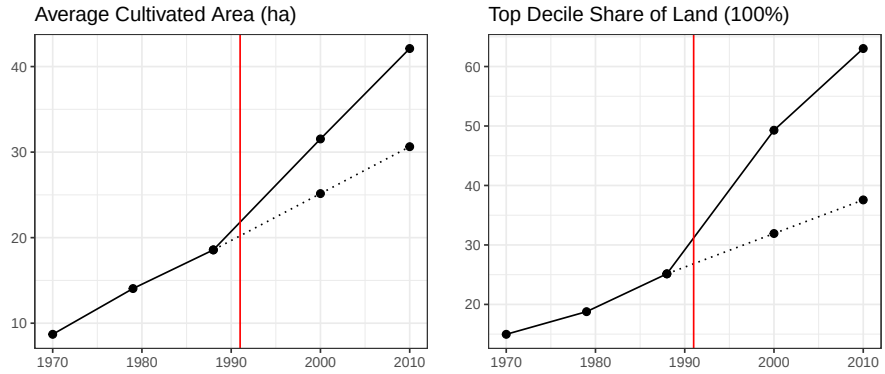
(b) Crop-Specific Nitrogen Use



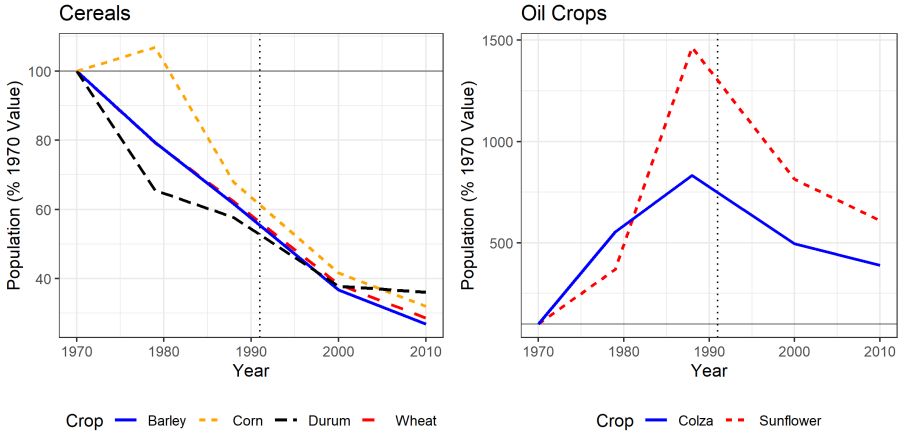
(c) Relative Input Quantities & Prices

Figure A1: Determinants of Chemical Use

Notes: Figure a) is computed using the FADN panel, and shows the 1985-1990 distribution across farms in the ratio of their chemical expenditures to their total revenue. Revenues are computed here using national average output prices, in order for the ratio to focus on the dispersion of relative chemical use. Figure b) highlights heterogeneity in nitrogen manure use (kg/ha) across crops, computed using the PKGC plot survey. The binscatters of figure c) are computed using combined data from the FADN (input quantities), the INSEE Ipampa price index series (chemical prices), and the Valeur Venale des Terres land price series. Land prices are corrected for the EU land subsidies. I compute input-to-input price and quantity ratios. Inputs are measured as deflated bills for pesticides and fertilizers, and in hectares for land. Prices for chemicals are Laspeyres price indices, while land prices correspond to an average land price in France. All indices are normalized to 100 in 1980.



(a) Farm Growth and Reallocations



(b) Differentiated Farm Exit across Markets

Figure A2: Farm Growth, Reallocations and Exit around the MacSharry Reform

Notes: The data comes from the full count Census of Agriculture for years 1970, 1979, 1988, 2000 and 2010. Average farm size is measured in hectares, and the share of land allocated to the top decile of farms (in size of land) is expressed in percent. These two variables are computed for the set of farms active on the grain market. The evolution of the respective farm populations is expressed wrt. to the 1970 baseline value, in order for the rates to be comparable across crops. The vertical line is for 1991, the last year prior to the reform. The dotted line for the first three figures show the average trend over 1970-1988.

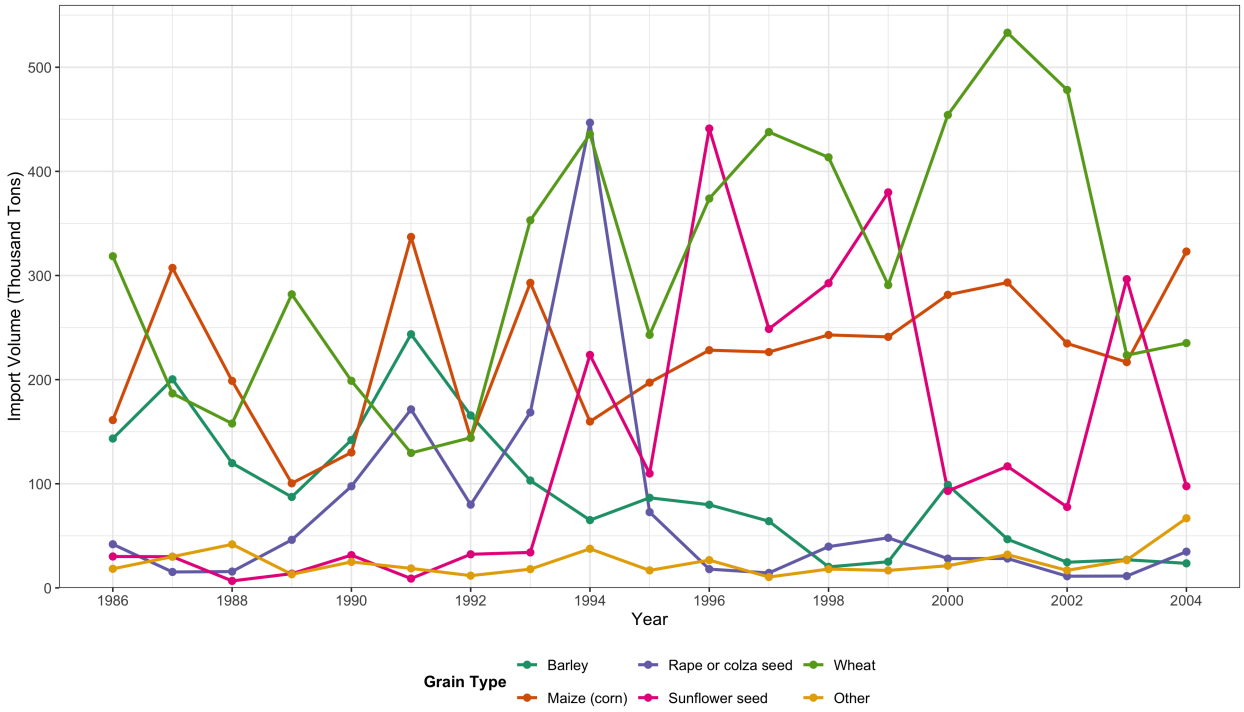


Figure A3: French Grain Imports

Notes: This figure shows the evolution of French grain imports expressed in thousands of tons around the period of the MacSharry reform. The data comes from FAOSTAT.

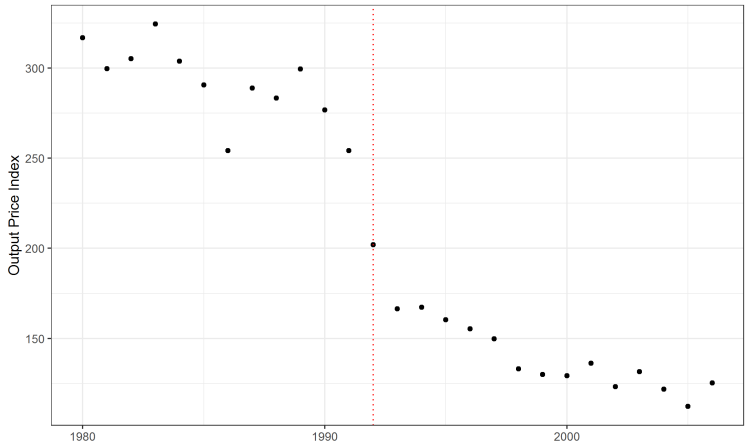


Figure A4: Average Farm Price

Notes: Binscatter for the average output price across FADN, computed over oil and cereal crops, using relative land shares as weights.

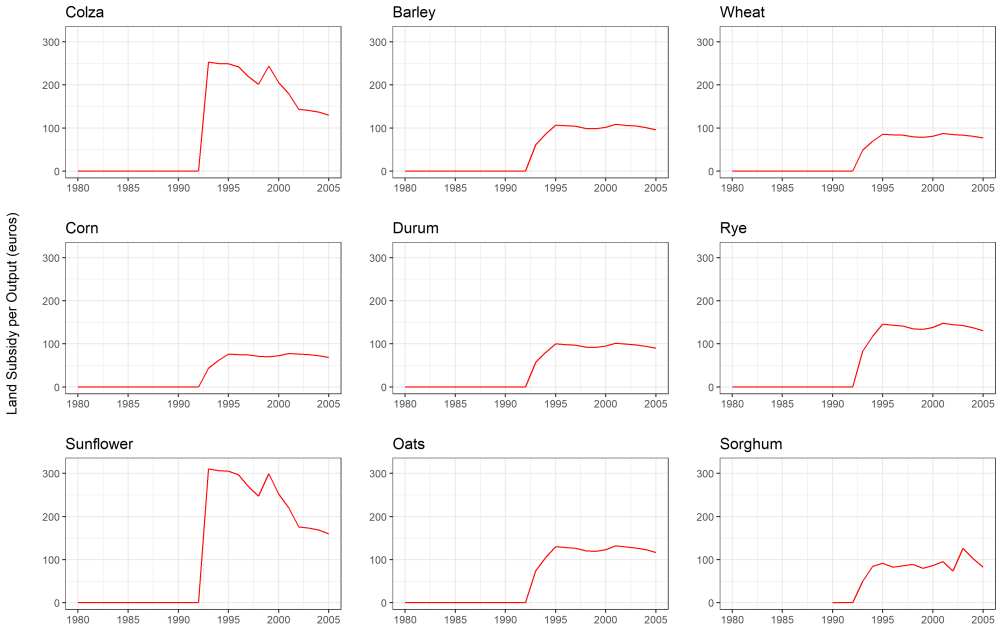


Figure A5: EU Land Subsidies per Crop under the CAP

Notes: The figure shows the evolution of the average land subsidy per unit of output over time, as measured in the FADN. I measure the average subsidy per hectare for cereals and oil crops, dividing the farms' total subsidy for each crop category, by their land allocated to each category. I then divide this measure by the average crop-specific yield observed in France in that year, in order to recover a subsidy per unit of output. Subsidies per land are relatively stable across years, while yields increase, making it so that subsidies slightly decrease over time absent any policy adjustment.

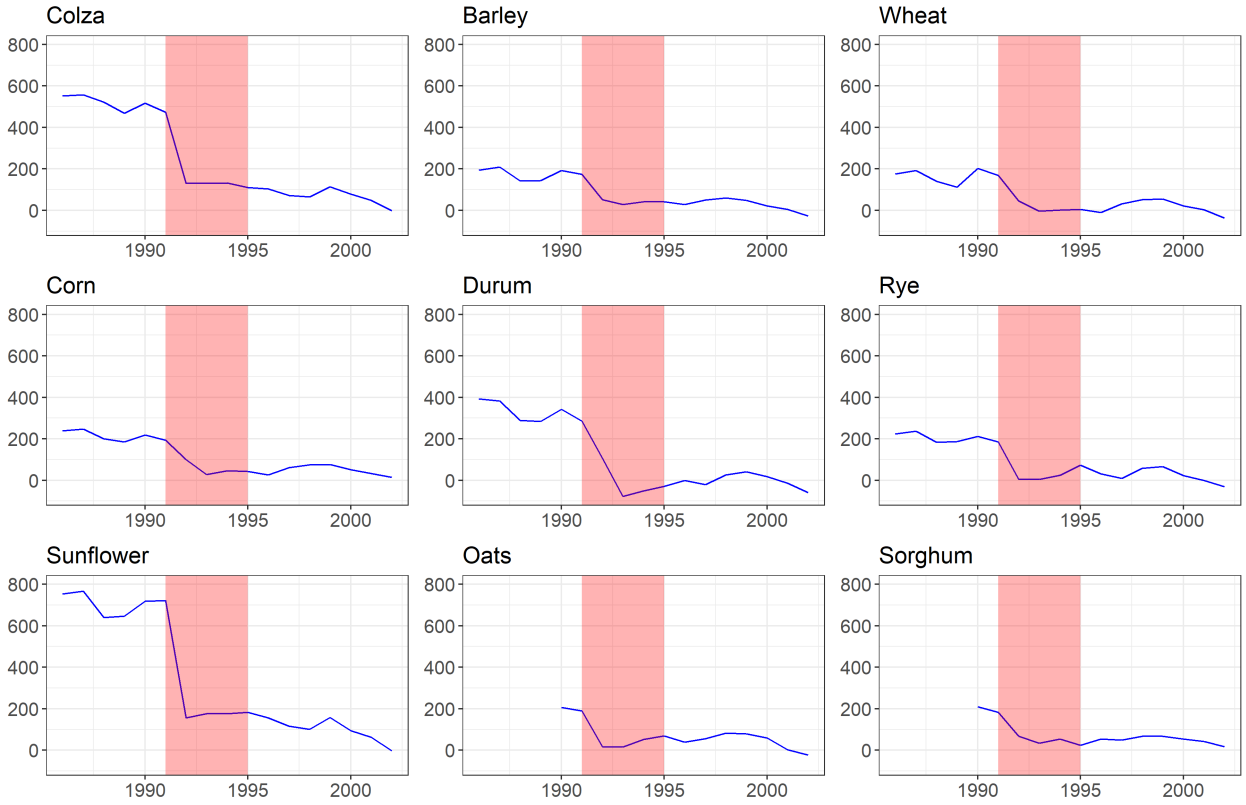


Figure A6: EU Subsidization per Unit of Output—Heterogeneity in Losses around the Reform

Notes: The figure plots crop-specific EU subsidies per unit of output over time, using data from the FADN, and digitized time series of EU intervention prices over time, as well as U.S. and Canadian farm gate prices from Faostat. The shaded area is for 1991-1995, going from the last year pre-reform to the last year of the reform. Data on oats and sorghum is only introduced in 1990 in the FADN.

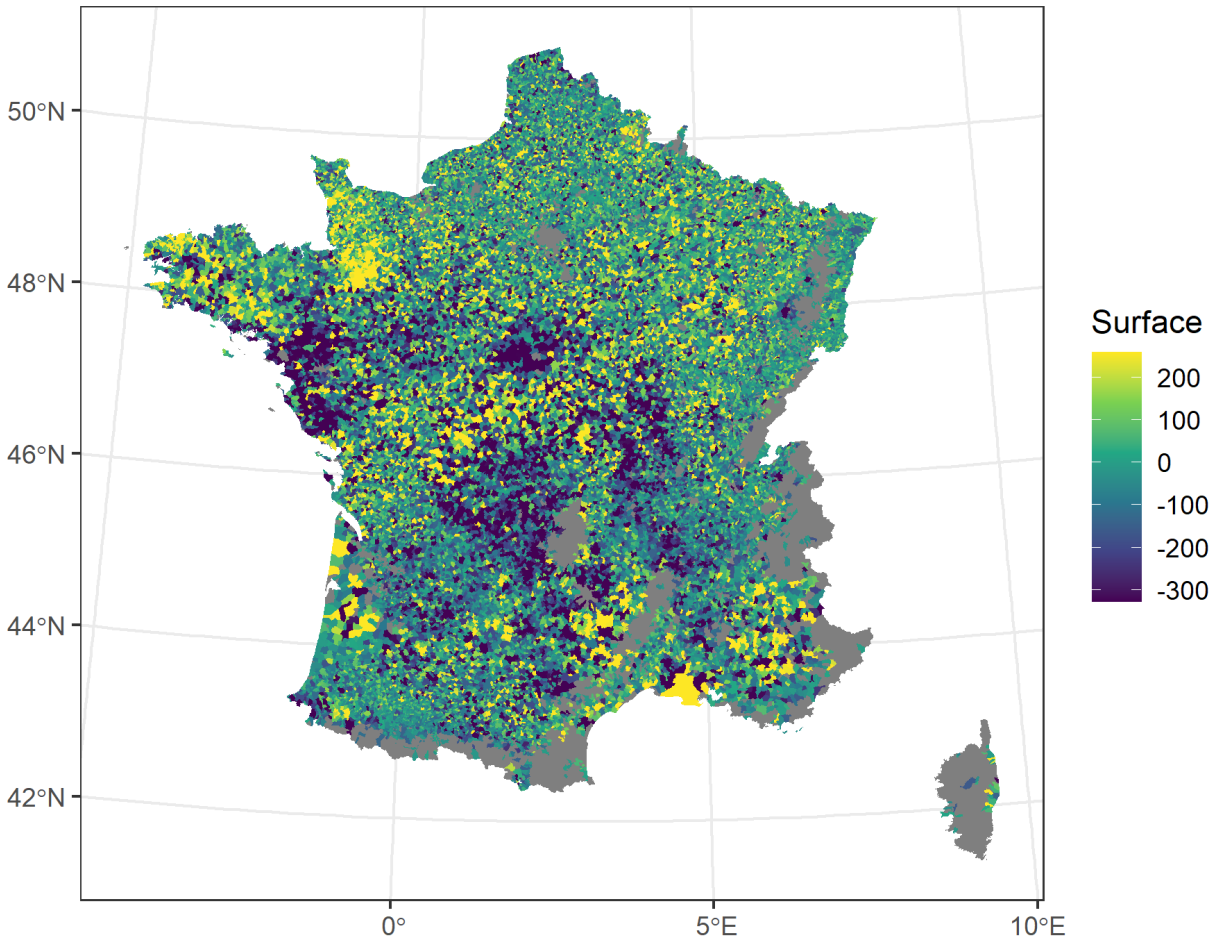


Figure A7: Evolution of Village-level Total Agricultural Area

Notes: This map shows the evolution of within-village total agricultural area between 1988 and 2000. I use the data from the relevant waves of the French Agricultural Census. Decreases in agricultural area mainly happen in the center of France around the Massif Central, as well as in Vendée and possibly lower Brittany.

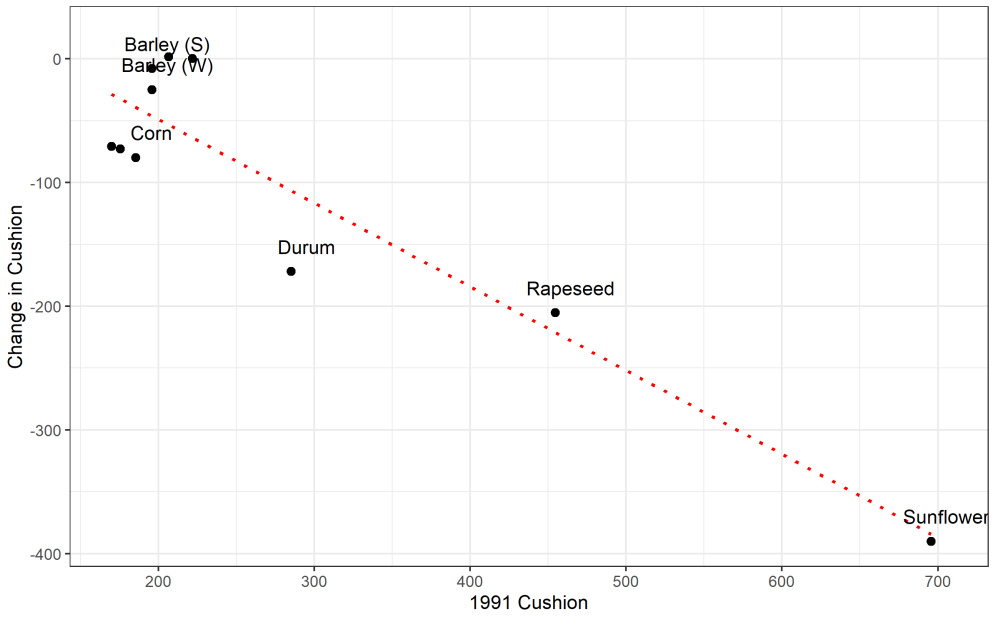


Figure A8: Convergence of Crop Cushions Post Reform

Notes: The figure illustrates the convergence of crop cushions post reform. Here the cushion is the French average of the crop-specific cushions in the relevant year, computed according to the formula outlined in Section 3.1. The change in cushion is the change observed between 1991 and 1995, corresponding to the period of the MacSharry reform.

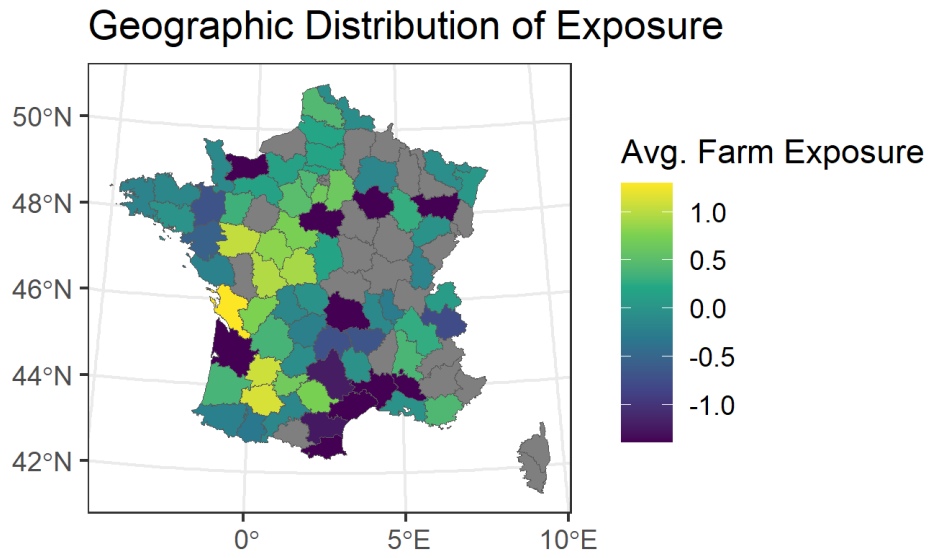


Figure A9: Department Weighted Averages of Farm-Level Exposure

Notes: The figure shows the geographic distribution of the standardized measure of farm-level exposure. The exposure is computed using data from the FADN, and department-level averages are computed using the extrapolation weights provided there. I plot the geographic variation of the standardized exposure here.

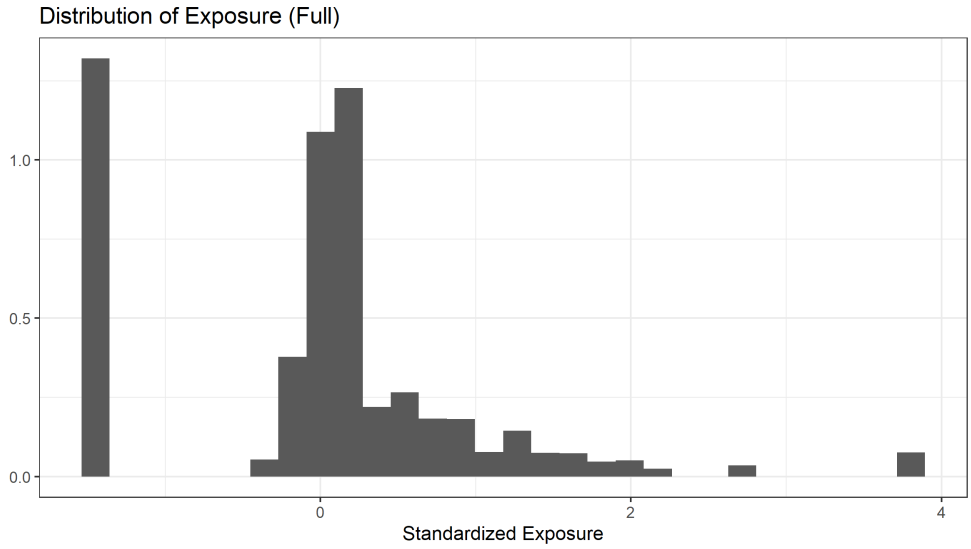


Figure A10: Distribution of Farm-Level Exposure

Notes: The figure shows the distribution of the standardized measure of farm-level exposure. The exposure is computed using data from the FADN.

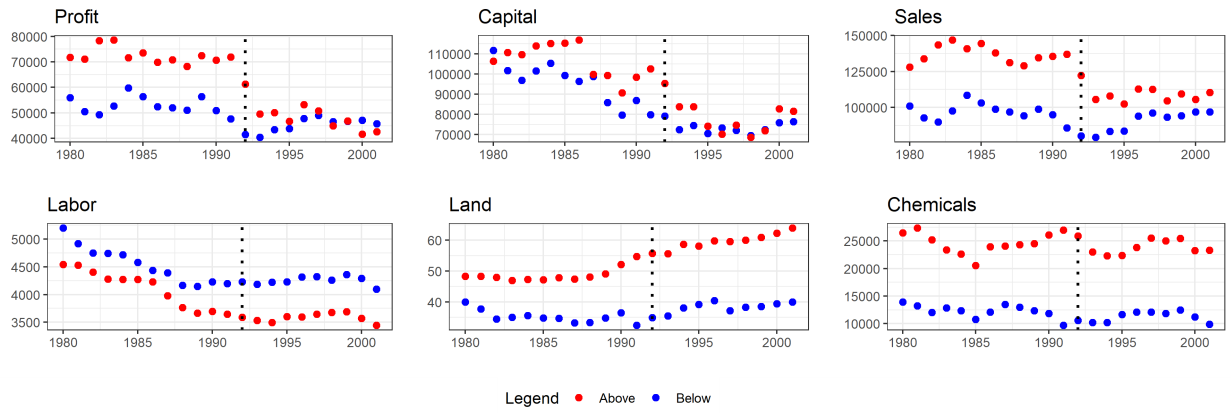


Figure A11: Balance in Levels and Trends

Notes: These binscatters are computed using data from the FADN, and use sampling weights to compute bin-specific averages within each category. Profit corresponds to total farm profit, capital is the sum of total farm value in buildings and machinery, sales is total sales. All of these values are computed in 2020 euros. Labor is total labor in hours per year, and land total utilized agricultural area in hectares.

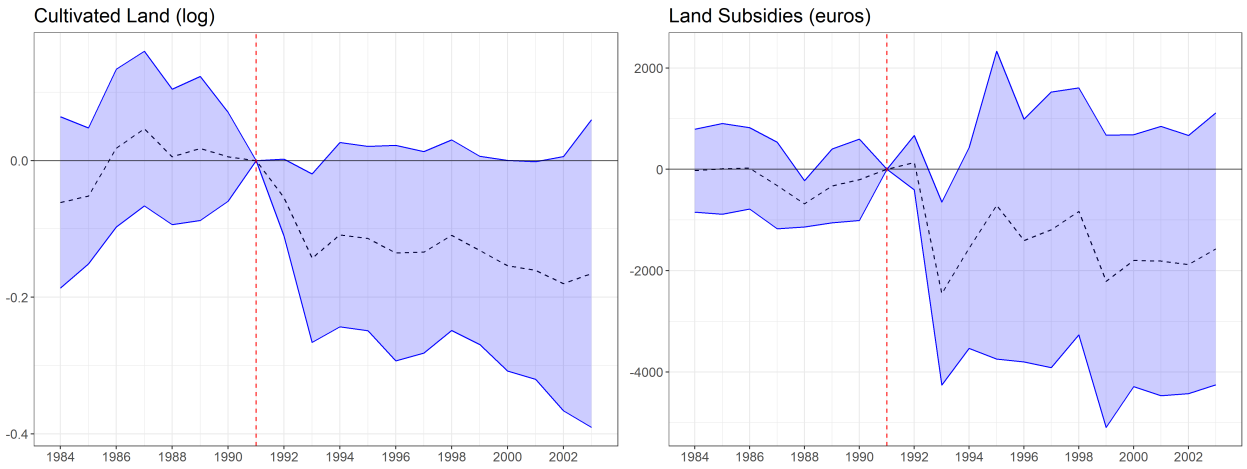


Figure A12: Exposure to the Reform—Farm-Level Event Study (Additional results)

Notes: The outcomes of the event studies are the following: the amount of land cultivated within the farm in log terms, and the amount of subsidies received from the land subsidy post-MacSharry reform in euros (this measure is 0 for all farms up to 1992, which is why it is measured in levels). All coefficients correspond to the year-specific coefficient associated to farm exposure, and give the effect of exposure in that given year. The vector of controls contains the following elements: farm 1991 capital stock, total labor used, total land use, their chemical use, the number of crops they grew and the evenness of their land allocation, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 adoption trends in chemicals measured as the evolution in their chemical use. I add department-year and farm fixed effects, and cluster the standard errors one-way at the department level.

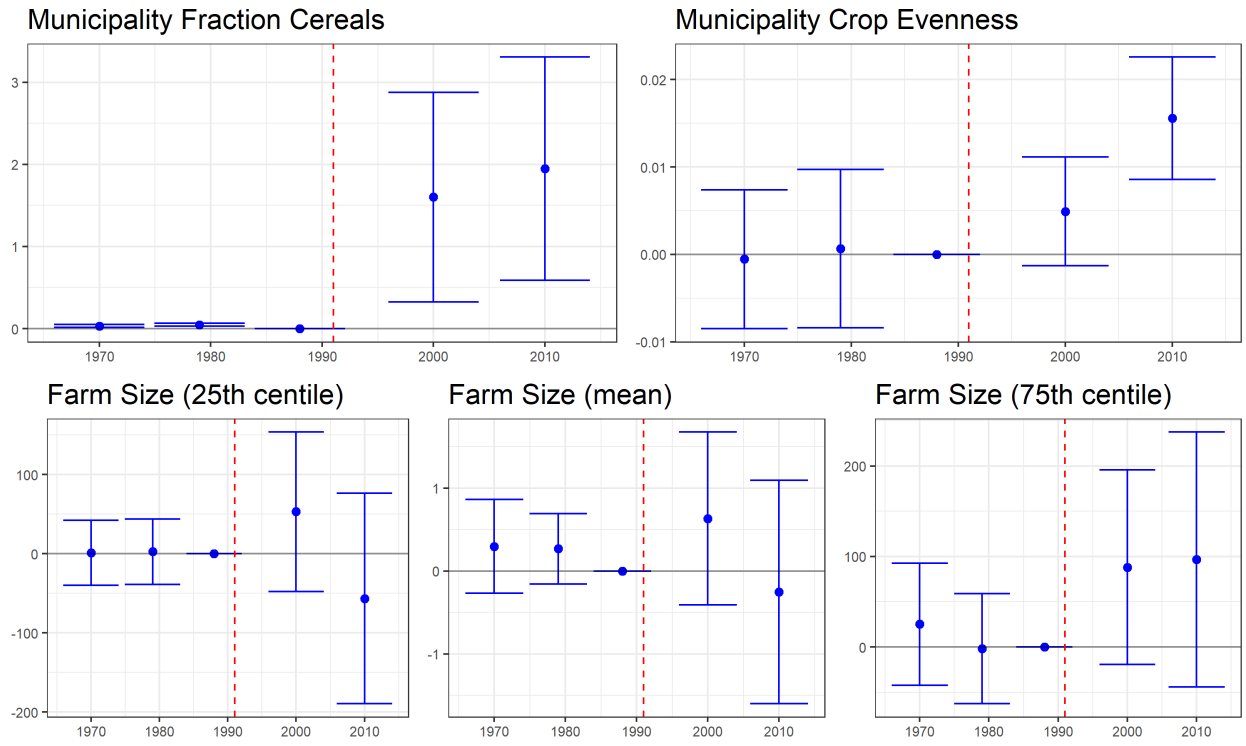


Figure A13: Additional Municipality Level Results

Notes: This figure gives the results for the municipality-level event study. Outcomes are the following: the fraction of land in the municipality allocated to cereals, the evenness of land allocation across crops at the municipality level, the 25th, mean and 75th centile of the farm size distribution, measured in ha cultivated land. The regression includes a series of controls set to their level in 1988 within the municipality, and interacted with time-varying coefficients, as well as municipality and department-by-year fixed effects. Standard errors are clustered at the department level.

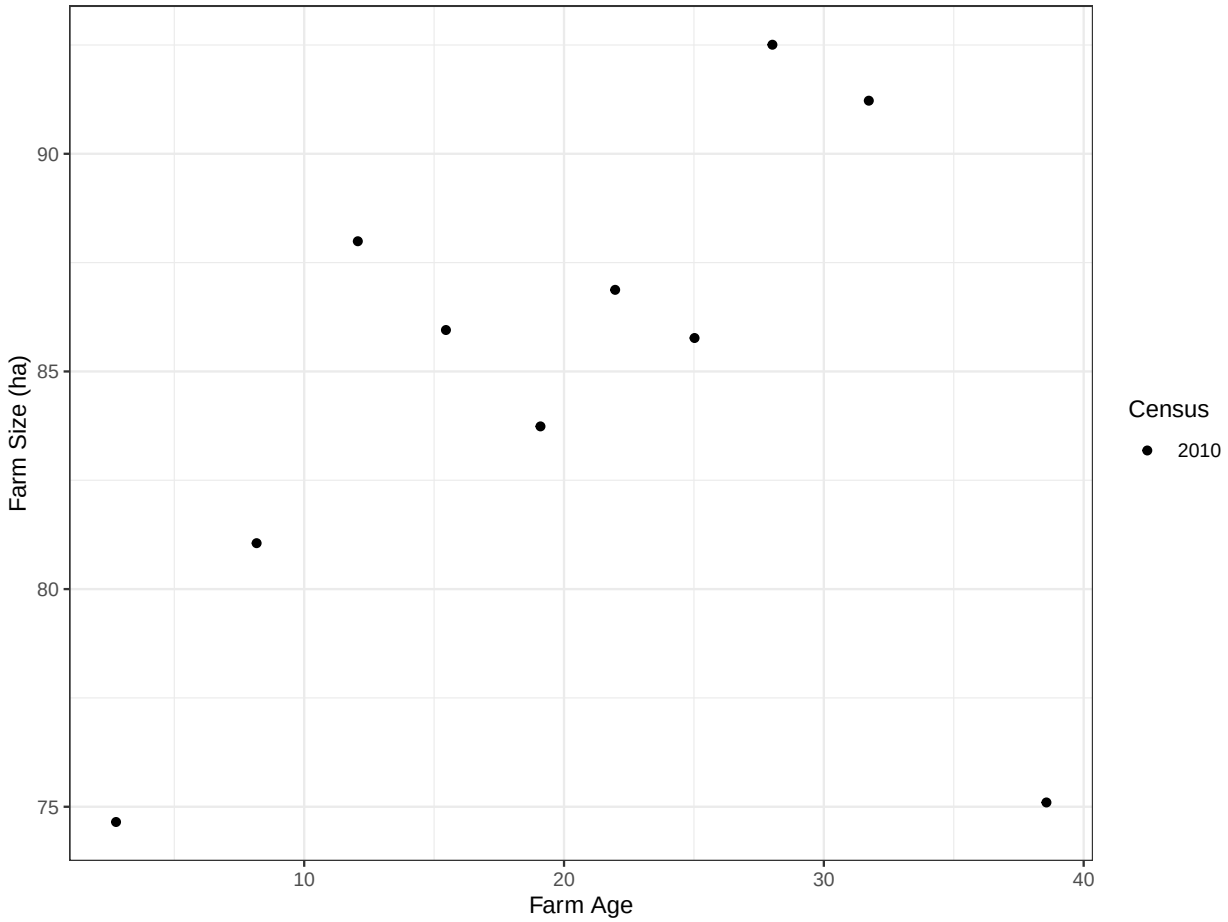
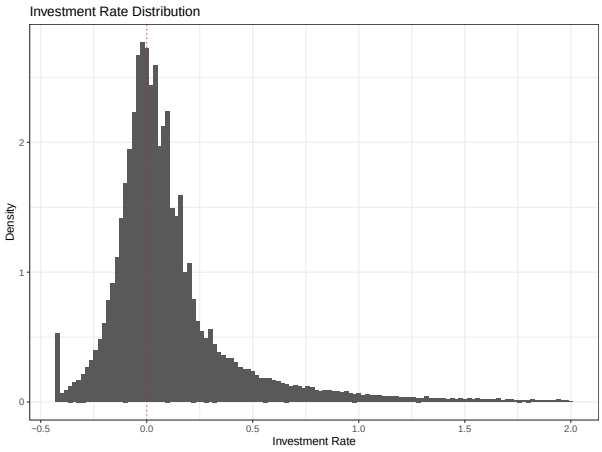
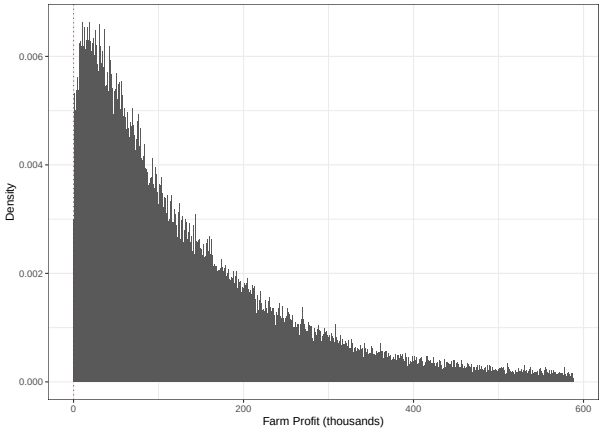


Figure A14: Relation between Farm Size and Farm Age in the 2010 Census

Notes: The figure shows a binscatter showing the relation between farm size and farm tenure, as observed in the 2010 French agricultural census. I use the 2010 Census to capture this relation, as the total number of farms ceases to decrease as quickly between 2000 and 2010, and one can interpret the market as relatively more stable than in the previous decades. In previous waves of the Census, older farms had likely entered the market under very different conditions, and were more likely to be small non-commercial farms with different management styles.

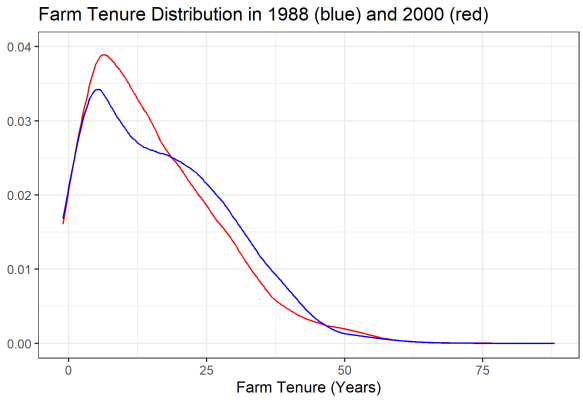


(a) Distribution of Investment Rates (FADN)



(b) Distribution of Profits (FADN)

Figure A15



(a) Distribution of Tenure (Census)

Figure A16

Notes: I plot the distributions from which I obtain the moments used for the simulated method of moments estimation routine. The investment rate is recovered as the difference between the current capital level and the lagged depreciated capital level, normalized by the lagged depreciated capital level. Profits are directly observed in the data, as well as farm tenure. The first two graphs are obtained from FADN data, while the last one is built with 1988 and 2000 Census data.

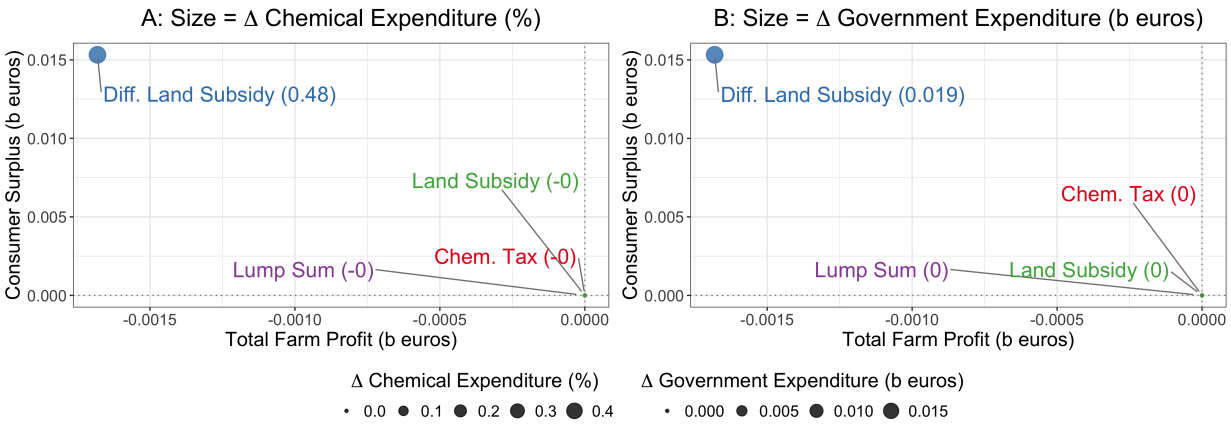


Figure A17: Comparison of Welfare Maximizing Policies (alternative valuation)

Notes: This figure plots the welfare implications of different optimal subsidy and taxes. Within each design, a grid search is performed over the range of possible intervention levels, and the welfare optimizing level is selected. The welfare criterion used for this is composed of changes in consumer surplus, total farm profit and environmental gains from reduced chemical use minus government expenditure. The marginal cost of agricultural chemical pollution is here taken as the 2025 marginal tax rate for glyphosate in France, of roughly 0.14 euros per euro of expenditure. In this exercise, I consider that this tax corresponds to the chosen social valuation for agricultural chemical pollution, and apply it to all of the chemical expenditure on the market. This exercise serves as a robustness check of the main analysis, where the value accounts for al types of agricultural chemicals and is taken from CGDD (2011). The figure plots the difference in consumer surplus, total farm profit and government expenditure relative to the no-intervention equilibrium. The change in chemical expenditure is expressed in percent, also compared to the no-intervention equilibrium.

A.2 Tables

Table A1: Descriptive Statistics - Farm-Level Dataset

Statistic	N	Mean	St. Dev.
Output Volume—Sum (t)	308,984	179	303
Output Volume—Wheat (t)	191,357	149	204
Output Volume—Corn (t)	95,925	116	205
Output Volume—Sunflower (t)	36,263	31.0	33.9
Output Price Index (€/t)	169,219	213.3	94.8
Farm Surface (ha)	308,982	71.0	67.6
Capital (€)	308,984	128,136	167,984
Total Labor (hours)	308,982	4,406	3,739
Profit (€)	308,984	161,577	264,342
Phytosanitary (expenses in €)	268,893	8,071	11,066
Fertilizer (expenses in €)	268,893	16,040	17,643

Table A2: Agricultural Market Trends

Panel A	All Farms (0/1)							
	Year	Farms (K)	Total Farm Surface (K ha)	Total Cultivated Land (K ha)	Avg. Farm Surface	Avg. Cultivated Land	Gini (all land use)	Gini (cultivated)	Land Share of Top 10^{th} Decile
	1970	1583	29905		18.8		0.58		0.23
	1979	1257	29497		23.4		0.58		0.29
	1988	1006	28596		28.4		0.5816		0.37
	2000	653	27856		42.6		0.61		0.62
	2010	519	27833		53.6		0.62		0.73
Panel B	Row Crops Farms (1/1)								
	Year	Farms (K)	Total Farm Surface (K ha)	Total Cultivated Land (K ha)	Mean Farm Surface	Mean Cultivated Land	Gini (all land)	Gini (cultivated)	Land Share of Top 10^{th} Decile
	1970	1035	25984	9006	25.1	8.7	0.49	0.66	0.15
	1979	823	25882	11568	31.4	14.0	0.48	0.66	0.19
	1988	602	23828	11189	39.5	18.6	0.48	0.65	0.25
	2000	352	22885	11106	64.9	31.5	0.48	0.63	0.49
	2010	261	21828	10980	83.7	42.1	0.46	0.60	0.63

Notes. The data is computed using the full count agricultural census for years 1970, 1979, 1988, 2000 and 2010. The average farm size and total agricultural land are both computed in hectares, and the total agricultural area for France corresponds to the sum of all total used agricultural areas at the farm level. Cultivated land for row crop farms is computed for the largest set of row crops that I can track across the years of the Census: wheat, durum, barley, oats, rye, corn, sorghum, rice, beetroot, rapeseed, sunflower and soy. The Gini-Coefficient is computed at the country level, for the distribution of farm size.

Table A3: Crop Mixes within French Farms

Year	Only Wheat	Non-Wheat Only	Wheat +	Other +	Only Barley	Only Corn	Only Oats	Only Rye	Only Sunflower	Only Colza
1979	0.041	0.19	0.64	0.14	0.048	0.11	0.011	0.0071	0.00020	0.00040
1988	0.094	0.18	0.64	0.092	0.054	0.085	0.014	0.011	0.0058	0.0023
2000	0.13	0.19	0.62	0.066	0.045	0.096	0.013	0.0051	0.0090	0.0022
2010	0.14	0.17	0.63	0.063	0.046	0.079	0.0078	0.0057	0.0078	0.0042

Notes. The data is computed using the full count agricultural census for years 1970, 1979, 1988, 2000 and 2010.

Table A4: Farm-Level Shift-Share Results

Dependent Variables:	Output Price (log)	Sales (log)	Profit (log)	Chemical Expenditure (log)
Model:	(1)	(2)	(3)	(4)
Variables				
$Exposure_j \times 1984$	-0.2278	-0.0485	-0.1200	-0.0504
	(0.2068)	(0.1203)	(0.1554)	(0.1166)
$Exposure_j \times 1985$	-0.0038	-0.0133	-0.0111	-0.0158
	(0.2447)	(0.0897)	(0.1703)	(0.1347)
$Exposure_j \times 1986$	0.5301	-0.0081	0.0971	-0.0374
	(0.3720)	(0.1031)	(0.1238)	(0.1035)
$Exposure_j \times 1987$	-0.0639	0.0255	0.0704	-0.1787
	(0.1798)	(0.0488)	(0.1497)	(0.1182)
$Exposure_j \times 1988$	-0.0692	0.0330	0.2465	0.0476
	(0.1596)	(0.0719)	(0.1669)	(0.0975)
$Exposure_j \times 1989$	0.1564	0.0228	0.1745	-0.0332
	(0.1854)	(0.0828)	(0.2214)	(0.0697)
$Exposure_j \times 1990$	0.0513	-0.1275**	-0.1554	0.1067
	(0.1226)	(0.0552)	(0.1702)	(0.0642)
$Exposure_j \times 1992$	-0.2034*	-0.1370***	-0.3424**	-0.0828
	(0.1131)	(0.0478)	(0.1411)	(0.0820)
$Exposure_j \times 1993$	-0.2790	-0.1985*	-0.1827*	-0.1593
	(0.1736)	(0.1023)	(0.1087)	(0.1269)
$Exposure_j \times 1994$	-0.2486**	-0.1279	-0.1053	-0.1900
	(0.1097)	(0.1156)	(0.2266)	(0.1389)
$Exposure_j \times 1995$	-0.3789**	-0.2209*	-0.4566***	-0.1974
	(0.1462)	(0.1269)	(0.1545)	(0.1442)
$Exposure_j \times 1996$	-0.6239***	-0.2295	-0.2339	-0.2456
	(0.1588)	(0.1520)	(0.1684)	(0.1507)
$Exposure_j \times 1997$	-0.3100**	-0.2322*	-0.0670	-0.1150
	(0.1331)	(0.1385)	(0.1638)	(0.2252)
$Exposure_j \times 1998$	-0.4480**	-0.1182	-0.2958**	-0.2740**
	(0.1771)	(0.1559)	(0.1262)	(0.1268)
$Exposure_j \times 1999$	-0.3481***	-0.3201**	-0.1974	-0.4082***
	(0.0885)	(0.1226)	(0.1973)	(0.1192)
$Exposure_j \times 2000$	-0.5517**	-0.2905**	-0.3045**	-0.2764
	(0.2088)	(0.1335)	(0.1233)	(0.1679)
$Exposure_j \times 2001$	-0.2311	-0.3850***	-0.5775***	-0.3637*
	(0.1778)	(0.1440)	(0.1736)	(0.1936)
$Exposure_j \times 2002$	-0.2828	-0.3332**	-0.4189**	-0.2631
	(0.2089)	(0.1420)	(0.1718)	(0.2085)
$Exposure_j \times 2003$	-0.0658	-0.4303***	-0.5509***	-0.3416***
	(0.1506)	(0.1556)	(0.1531)	(0.1262)
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Department-Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	1,910	2,978	2,945	2,978
R^2	0.96634	0.95419	0.89874	0.94800

Clustered (Department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. The event study outcomes are the following: a farm-level output price index (average price across row crops, using relative areas as weights), total sales, total farm profit and the farms' deflated chemical bill. All coefficients correspond to the year-specific coefficient associated to farm exposure, and give the effect of exposure in that given year relative to the effect in 1991. The vector of controls contains the following elements: farm 1991 capital stock, total labor used, total land use, their chemical use, the number of crops they grew and the evenness of their land allocation, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 adoption trends in chemicals measured as the evolution in their chemical use. I add farm and department-year fixed effects, and cluster the standard errors one-way at the department level. The figure for these results is Figure 2.

Table A5: Municipality-Level Shift-Share Results

Dependent Variables:	Municipality Farm Count	Farm Crop Evenness
Model:	(1)	(2)
Variables		
$Exposure_k \times 1970$	-0.0501	-0.0051
	(0.5820)	(0.0033)
$Exposure_k \times 1979$	0.2912	-0.0040
	(0.3815)	(0.0030)
$Exposure_k \times 2000$	-0.7364**	0.0112***
	(0.3016)	(0.0021)
$Exposure_k \times 2010$	-1.139***	0.0217***
	(0.4017)	(0.0028)
Fixed-effects		
Municipality	Yes	Yes
Department-Year	Yes	Yes
Fit statistics		
Observations	147,157	147,157
R^2	0.95	0.84

Clustered (Department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. This table gives the results for the municipality-level event study. Outcomes are specified in levels and correspond to: the number of farms operating in the grain market within the municipality and the average farm-level index for the evenness of the distribution of land across crops. The regression includes a series of controls set to their level in 1988 within the municipality, and interacted with time-varying coefficients, as well as department-by-year and municipality fixed effects. Standard errors are clustered at the department level. The figure for these results is Figure 3.

Table A6: County-Level Shift-Share Results: Algal Blooms

Dependent Variable:	Algal Bloom (log)
Model:	(1)
Variables	
$Exposure_c \times 1986$	-0.0458
	(0.0322)
$Exposure_c \times 1987$	-0.0700**
	(0.0265)
$Exposure_c \times 1988$	-0.0293
	(0.0301)
$Exposure_c \times 1989$	0.0239
	(0.0323)
$Exposure_c \times 1990$	-0.0224
	(0.0407)
$Exposure_c \times 1992$	-0.0717*
	(0.0408)
$Exposure_c \times 1993$	-0.0454
	(0.0277)
$Exposure_c \times 1994$	-0.0467
	(0.0299)
$Exposure_c \times 1995$	-0.0440*
	(0.0253)
$Exposure_c \times 1996$	-0.0687**
	(0.0317)
$Exposure_c \times 1997$	-0.0905***
	(0.0268)
$Exposure_c \times 1998$	-0.0305
	(0.0297)
$Exposure_c \times 1999$	-0.1068***
	(0.0286)
$Exposure_c \times 2000$	-0.1662***
	(0.0377)
$Exposure_c \times 2001$	-0.0975***
	(0.0278)
Fixed-effects	
County	Yes
Department-Year	Yes
Fit statistics	
Observations	25,764
R^2	0.88

Clustered (Department) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. This table gives the results for the county-level event study. The outcome is a Landsat-5 based index of algal bloom intensity on the within-county water bodies. Algal blooms are generally caused by the over-fertilization of agricultural land. The regression includes a series of controls set to their level in 1988 within the county (last year of the Census prior to the reform), and interacted with time-varying coefficients, as well as county and department-by-year fixed effects. Standard errors are clustered at the department level. The figure for these results is Figure 4.

Table A7: Demand Results

Dependent Variables:	Log Quantity (Wheat)	Log Quantity (Other)	Log Quantity (Wheat)	Log Quantity (Other)
Model:	(1)	(2)	(3)	(4)
Variables				
Constant	24.4***	22.4***	21.3***	22.5***
	(1.69)	(1.97)	(1.21)	(2.18)
Price (log)	-0.89***	-0.58*	-0.77***	-0.96*
	(0.28)	(0.32)	(0.19)	(0.35)
Instruments	All	All	Only Deviations & Interactions	Only Deviations & Interactions
Fit statistics				
Observations	28	28	28	28
R^2	0.79	0.63	0.78	0.45

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. I obtain estimates of total quantity sold and average prices for France for the two crop categories using the FADN and its sampling weights. I then run regressions of quantities on prices, instrumenting the prices with weather shocks. These are used as supply shifters orthogonal to demand. The full set of instruments includes: yearly GDD and HDD realizations at the national level, the national yearly average of GDD and HDD department level deviations from the thirty-year department average, and the interaction between levels and deviations. The second set of regressions does not include national realizations in levels. I include a linear time trend in all regressions, and use the Newey-West correction for serial auto-correlation.

Table A8: Model Parameters

Coefficient	Parameter	Value
Other Estimated Parameters		
Competence Ladder	λ	0.70
Mean TFP	μ	3.66
Auto-Correlation TFP	ρ_h	0.79
Shape TFP Pareto Distribution	σ_h	0.32
Location TFP Pareto Distribution	δ_h	0.07
Auto-Correlation Non-Hicksian shock	ρ_{ch}	0.80
Shape Non-Hicksian Pareto Distribution	σ_{ch}	0.23
Location Non-Hicksian Pareto Distribution	δ_{ch}	0.00
Calibrated Parameters		
Time Preference	$\frac{1}{1-r}$	0.9615
Capital Depreciation	δ_k	0.08
Price Capital (k €)	p_k	1,000
Price Land	p_s	1,034
Price Chemicals	p_c	347
Price Labor	p_l	842

Notes: I use the recovered productivity shocks to estimate these other estimated parameters. Specifically, I demean TFP shocks, and use the average as an estimate of μ. Demeaned TFP shocks are used to compute the following parameters. The competence ladder corresponds to the average ratio of the highest crop TFP to lowest crop TFP within a farm, across farms and years. The Pareto distribution parameters are recovered across farms and years, and the auto-correlation in both shocks are computed using within-farm variation. For the additional calibrated parameters, the rate of capital depreciation is set to 8%. I do not observe chemical use quantities in the FADN, but rather the overall chemical expenditures. As such, I express the price of chemicals as the chemical expenditure for one ton of wheat observed in the FADN. For consistency, the price of other inputs is similarly expressed as the input-specific expenditure needed to produce one ton of wheat.

B Data Description

B.0.1 Farm Accountancy Data Network (FADN)

My primary source of farm-level data is the French subset of the European Farm Accountancy Data Network (FADN), a rich panel survey designed to be representative of commercial agriculture in the European Union.[48] The FADN is an ideal dataset for the analysis due to its detail and panel structure. Crucially, while the FADN is a sample (about 8,000 farms per year in France), it is designed to capture the vast majority of economic activity. In 2000, the commercial farms it represents accounted for 95% of France's total agricultural production.[49] The dataset's richness stems from its foundation in farm accounting records. For each farm, it provides annual data on crop-specific sales and produced quantities, which allows me to recover farm-level output prices. It also contains detailed information on input expenditures (e.g., fertilizer, pesticides), capital stocks, and crop-specific land allocations. Finally, farms are geocoded at the department level for the period of analysis, enabling me to match them with local input prices and weather data.

B.0.2 French Agricultural Census

To analyze market-level dynamics, I leverage the farm-level French Agricultural Census, a unique resource that provides data on the universe of French Farms. The census is fielded approximately every ten years (I use the 1970, 1979, 1988, 2000, and 2010 waves) and covers virtually every agricultural establishment in France.[50]

While the Census lacks the detailed financial data of the FADN, its comprehensive scope makes it ideal for measuring market structure. Its primary variables for my purposes are farm municipality-level location and crop-specific acreage. A key limitation is the absence of a stable farm identifier prior to 1998, which prevents me from tracking individual farms across census waves. I therefore use the repeated cross-sections of the Census to construct a municipality-level panel. Other sources of data used in this paper are described in Appendix B.

[48] My analysis focuses on France, as many of the datasets I match the FADN to only exist for France.

[49] The FADN's definition of a "commercial farm" has evolved but consistently targets farms that are professionally managed and market-oriented. Prior to 2010, it included requirements on the manager's work hours and a minimum land area of 5ha for non-specialized farms. Since 2010, the threshold is based on a minimum standard production capacity of €25,000.

[50] Farms are defined as production units with independent management that meet a minimum size of .2ha, with exceptions for specific agricultural sub-markets. The original list of establishments filtered with these requirements is built out of a combination of land registries and data from the farmers' family allowances mutual savings societies.

B.1 Intervention Prices

For the purpose of this paper, and to the best of my knowledge, I created the first database of agricultural commodities' EU intervention prices covering the end of the 1960s to the early 2000s when intervention was completely removed. Intervention prices denominated in ECUs, euros or francs were gathered from the yearly and sometimes commodity-specific directives published by the European Union over that period of time. When prices were denominated in ECUs, I used the CAP-specific exchange rate which was then used to translate ECUs to each member state's currency. I similarly gathered data on commodity-specific land subsidies, which I later combine with farm-specific land subsidies as recorded in the FADN.

B.2 Input Prices: Indices and Price Surveys

Input prices are drawn from the IPAMPA agricultural input price indices, and from the Enquête sur l'observation des Prix des Consommations Intermédiaires nécessaires aux Agriculteurs (EPCIA), the price survey used to develop the IPAMPA. I use the IPAMPA indices for pesticides and fertilizers interpolated by INSEE back to 1980.

The EPCIA has been run since 1996, and provides price data for pesticides and fertilizers. Its sample is based on a 1995 survey, which helped identify a sample of products representing 50% of the total sales within their product category. The sampling of establishments selling these products was done based on the establishments' respective market share within that product category. Finally, within a product and establishment, the series chosen corresponded to the product's most frequent sales conditions. As such, series are good, company and mode of sale specific. When series disappear, they are replaced with their closest equivalent. The EPCIA is made of 4165 series, sold by 250 companies. The EPCIA is aggregated into national and regional price indices using 1990 plutocratic weights, and the following index of category-specific month-to-month price evolution:

$$i_t = i_{t-1}\left(\prod_i^n \frac{P_{i,t}}{P_{i,t-1}}\right)^{\frac{1}{n}} \tag{9}$$

I also recover agricultural hourly wages from the continuous labor survey "Enquête Emploi", which is a survey fielded for a sampled set of households, each household being drawn once and then followed for the next six trimesters. I consider the hourly wage data provided by workers whose occupation is classified as that of an agricultural worker or farmer.

Finally, I use land prices taken from the Land Market Value survey (Valeur Venale des

Terres), a yearly and department-level survey, which is fielded every year by the statistical services of the French departmental administration for agriculture and forestry. These are based on data provided by the public company in charge of land management (SAFER), which authorizes agricultural land purchases and consolidations when transactions surpass a given threshold. This data is then complemented by data provided by local notaries, and several local administrations. I use this data for 1994-2015. The data was digitized from scanned data catalogues for the first years of the series.

B.3 Plot Level Data

I use the "Pratiques Cultures sur les Grandes Cultures" or PKGC (Agricultural Practices for Field Crops) survey, a plot-level dataset, to highlight the heterogeneitya cross crops in their use of nitrogen.[51] It is important to note that this dataset is a repeated cross-section, not a panel.

B.4 Weather Data: Realizations and Forecasts

Our realized climate data comes from the European Centre for Medium-Range Weather Forecasts (ECMWF)'s ERA5 reanalysis product. ERA5 gives hourly estimates of climate data, out of which I use precipitation and temperature (temperature 2m above the surface of the Earth). ERA5 combines observational weather data with model-based data into a 0.25*0.25 gridded dataset.

I extract that data at the French department level, cropping the grid with department shapes, and averaging the data using simple area weights. Using the time separability assumption common in the literature, I aggregate the hourly data into growing-season observations [52]: growing degree days (GDD) and heat degree days (HDD) for temperature[53], and total precipitation for rainfall. In order to match the forecast dataset, I only use 4 daily measurements of temperature to compute the GDDs and HDDs, specifically at midnight, 6h, 12h and 18h.

The forecast data is taken from ECMWF's SEAS5 seasonal forecasting system. Forecast are

[51] Plots surveyed are selected among the farms that benefit from the European Union's Common Agricultural Policy. The survey focuses on land plots defined as the set of contiguous land for which the same crop is cultivated, with homogeneous agricultural practices (fertilizer and pesticide use for example). For each crop, the survey selects the minimum number of regions covering at least 95% of that crop's production, and within each region the minimum set of departments accounting for at least 90% of the region's production. Within departments, the survey selects farms with at least .1 hectare cultivated, and less than 200ha. A unique plot is selected within each farm. For the waves that we study, around 20,000 plots are sampled each time.

[52] I use an extensive definition of the agricultural growing season for France, running from October of the previous year, to July of the current one.

[53] Growing degree days are computed over the $[4°, 30°]$ degree interval, and heat degree days sum the realized temperature above 30°C.

produced on the first of each month for the following 5,160 hours [54]. For temperature, they are produced at a 6h interval, and give an instantaneous prediction of temperature, while for rainfall, they give the accumulation of rainfall every 24h. As such, the rainfall forecasts for the second day of the month will be the following: a 24h ahead forecast, a 30 to 32 days ahead forecast produced on the first day of the previous month, and so on until the lead value exceeds 5,160 hours. Temperature forecasts work in a similar way, but are just produced with more granular time steps.

The main issue for my purpose is that given that forecasts are produced every first of the month, different days within a month will not be provided a forecast with the same lead. I would ideally like to build the forecast-equivalents of the growing-season aggregates for realized weather, for different lead times. For example, the rainfall forecast for the growing season, with a constant one-month lead throughout the season (or the equivalent of the farmer's knowledge about rainfall one month in advance throughout the growing season). I approximate this by bundling together forecasts produced one calendar month ago (the one month lead for the rest of the paper), produced two calendar months ago, up to five months ago. As such, the forecasts that I aggregate into growing season observations are not homogeneous in terms of lead value, but are the closest equivalent of it that I can get.

Similarly to the weather realization data, I extract the gridded data into department-level observations, using area weights.

B.5 Landsat 5 Remote Sensing Data

Both the remote sensing indices are based on Landsat 5 imagery. Landsat 5 was a low orbit satellite jointly managed by the U.S. Geological Survey and NASA, and ran between 1984 and 2013. Its unusual longevity making it a good source of data to build long time series. It had a repeat cycle of about 16 days, and was equipped among else with a Thematic Mapper (TM) and a Multi-Spectral Scanner.

My algal bloom index follows the methodology of Taylor and Heal (2023). For every year between 1985 and 2001, I filter Landsat data for images taken between June and August, which I crop for surface water and treat for clouds and cloud shadows. I then run the following function on each pixel:

$$Bloom = NIR - 1.03 * SWIR \tag{10}$$

NIR corresponds to the near infrared band of Landsat 5 (.77 to .90μm), and SWIR to the shortwave infrared band (1.55 to 1.75μm). In both cases, these come from atmospherically corrected surface reflectance data produced by the Landsat TM series. The index is then

[54] ECMWF provides an ensemble of 25 forecasts, which I average.

averaged across the selected months of the year, and averaged over surface water area at the county level.[55]

C Robustness and Additional Results for the Reduced Form Analysis

I start this robustness section by showing the Rotemberg weights for the shift-share variable, following the formula provided by Goldsmith-Pinkham et al. (2020). The weights are crop-specific, and can be compared to both the crop-specific losses in subsidization from the reform, and the pre-reform variance in land shares across farms. These comparisons are useful to understand the source of the variation identifying the reform effects. I then compare them to the weights obtained without controls. This comparison highlights the role of controls in shifting the source of the variation used to identify the effect of farm reform exposure. I provide a more in-depth analysis of the conditional parallel-trends assumption driving the design. I show trends in the same set of variables as shown in Figure A11, but for farms with high or low shares for each crops used to build the shift-share. Finally, I provide the complete set of results in a difference-in-difference format, and give the shock-by-shock regression results, in order to decompose potential heterogeneity in treatment across crop-mix composition.

C.1 Rotemberg Weights

I give the Rotemberg weights below—with and without the inclusion of the controls used in the farm-level regressions. This has the advantage of both showing what cross-crop variation is used for identification, and how this variation is influenced by the controls. Table A9 shows the Rotemberg weights obtained when including all the controls used in the main FADN event study specification. The first thing I note is that sunflower, colza and durum are the three crops with high positive weights, while other grains have small or negative weights. These three crops are also the high exposure crops (with per-unit losses in subsidies of resp. 540, 315 and 362), while other grains have lower and very homogeneous losses. Sunflower, colza and durum are also among the set of crops with highest land share variance. Only wheat has a comparable share variance. As such, comparing farms with high and low shares of sunflower, colza and durum wheat is a good proxy for the overall cross-farm heterogeneity in crop mix.

When using the exposure variable, I am making comparisons across farms that had a high share of their land allocated to sunflower/colza/durum, versus others who were more special-

[55] There are 2,054 counties in France, making them much smaller geographic entities than U.S. counties, and the smallest geographic entity above villages (communes).

ized in the remaining cereals. As such, worries about omitted variable biases should focus on the uncontrolled determinants of crop choice for these three specific crops. The balance tests for these three crops shown later are then of particular relevance to build trust in the design. The correlation matrix between weight values, shocks, and the variance of the shares in Table A9 shows that the exposure variable can be taken as a reasonable summary of the differential exposure to the reform across farms, coming from the heterogeneous effects of the reform across crops, and from heterogeneity across farms in crop mix composition.

Table A9: Rotemberg Weights: With Controls

Panel A	Rotemberg Weights		
Crop	Weight	Shock	Share Variance
Sunflower	1.330	540	0.117
Colza	0.205	315	0.031
Durum	0.345	362	0.082
Winter Barley	0.038	133	0.015
Sorghum	-0.008	160	0.024
Rye	-0.031	113	0.006
Oats	-0.063	121	0.010
Spring Barley	-0.135	128	0.0001
Wheat	-0.294	165	0.034
Corn	-0.386	151	0.012
Panel B	Correlation Matrix		
Weights	1	0.90	0.83
Shocks	0.90	1	0.95
Variance Shares	0.83	0.95	1

Notes. The weights are computed after residualizing exposure on the fixed effects and controls used in the farm-level regressions. These Rotemberg weights are obtained when using only land shares to measure farm exposure. The shock column indicates the value of the crop-specific shock, and the share variance corresponds to the variance observed in the FADN sample of the crop-specific share.

The second thing to note about the weights is that oats, spring barley, wheat and corn have negative weights. These negative weights could raise issues in the case of heterogeneous treatment effects. I investigate this possibility Appendix C.4, by decomposing the shift-share variable into a series of crop-specific measures made of the farms' 1991 crop share interacted

with the crop-specific shock. These results highlight that the effects associated with a high crop share for oats, spring barley, wheat and corn are of the same sign as the ones shown in the main event study. The only exception is the effect of a higher oats crop share on chemical use, which is non-significant and very close to zero.

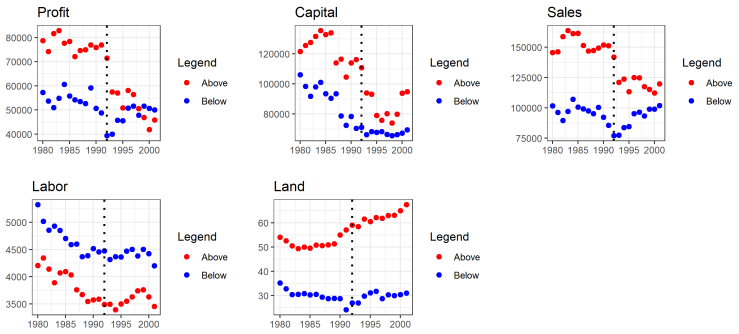
Finally, I note a change in the ordering of weights when comparing weights obtained with and without controls: wheat plays a much larger role in instrument variation when no controls are included, and colza and durum play a much smaller role in the variation. The addition of controls shifts the variation in the instrument from being driven by the high-shock crop sunflower, and then wheat and corn which play a large role in the French grain market, to comparing farms that were growing high versus low shock crops.

Table A10: Rotemberg Weights: Without Controls

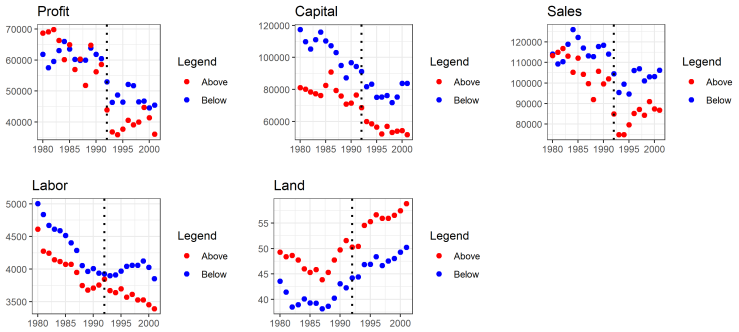
Crop	Weight	Shock	Variance Land Share
Sunflower	0.539	540	1.17-01
Wheat	0.18	165	1.05e-02
Durum	0.088	315	3.07e-02
Corn	0.076	151	3.37e-02
Colza	0.065	362	8.29e-02
Spring Barley	0.056	128	1.51e-02
Winter Barley	0.0047	133	9.58e-05
Sorghum	0.0016	160	2.37e-02
Oats	0.0015	121	5.74e-03
Rye	-0.0061	113	1.05e-02

Notes. The weights are computed without residualizing exposure on the fixed effects and controls used in the farm-level regressions. These Rotemberg weights are obtained when using only land shares to measure farm exposure. The shock column indicates the value of the crop-specific shock, and the share variance corresponds to the variance observed in the FADN sample of the crop-specific share.

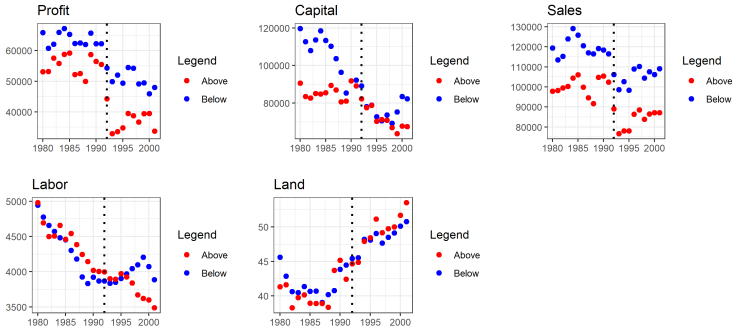
C.2 Balance Tests: Across Crop Shares



(a) Balance Tests - Wheat



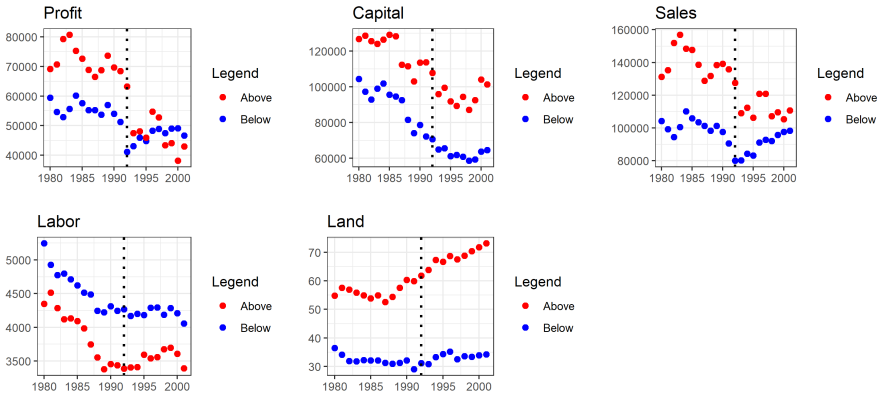
(b) Balance Tests - Sunflower



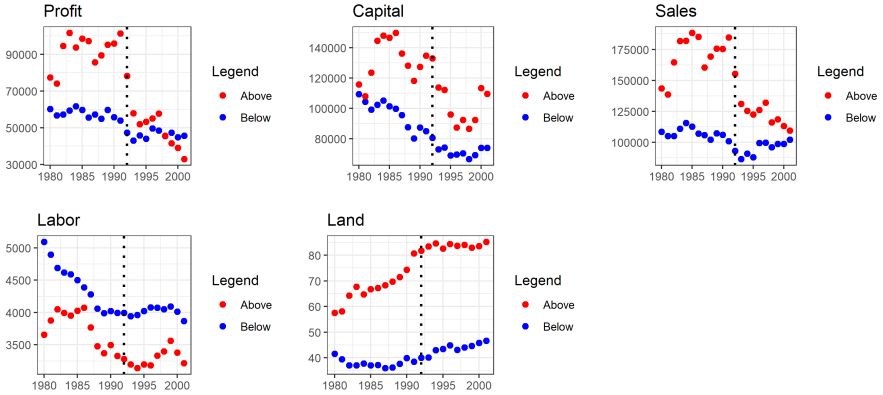
(c) Balance Tests - Corn

Figure A18: Trends per Category

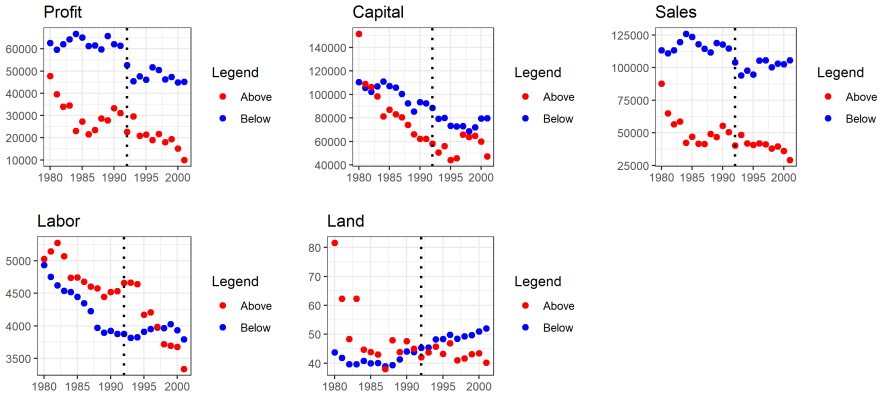
Notes: The figures show the trends of relevant farm characteristics for the bottom and top third of farms in terms of their share of land allocated to the relevant crop in 1991.



(a) Balance Tests - Barley



(b) Balance Tests - Colza



(c) Balance Tests - Rye

Figure A19: Trends per Category (continued)

Notes: The figures show the trends of relevant farm characteristics for the bottom and top third of farms in terms of their share of land allocated to the relevant crop in 1991.

Next, I show the balance tests across farms with an above and below median share of wheat, corn, winter barley, rye, sunflower and colza. The main threat to identification lies with the presence of unobservables that are linked with heterogeneity in exposure, and which affect differential farm-level growth. The following elements stand out: across all crops, capital is the one variable that differs the most in trends across groups prior to the reform—making it a necessary control in the design. Profit, sales and labor show relatively comparable trends prior. Trends in land use also show some notable differences in the cases of rye and colza. While the differences for rye are quite significant, I also know from the previous Rotemberg weights that rye plays an almost absent role in the variation of the shift-share variable, and hence should not drive results too much. Differences in colza are potentially more important, again justifying the control for the share of production coming from oil crops pre-reform.

C.3 Additional Farm-Level Results

It is useful to recast event study design in a pure difference-in-difference design in order to recover a unique coefficient summarizing the effect of the reform. I can do so using a modified estimating regression of the following form:

$$Y_{jt} = \alpha_j + \eta_{d(j)t} + \beta Exp_j Post_t + X_j Post_t \Gamma + \varepsilon_{jt}. \tag{11}$$

I also run this design with an alternative construction of exposure, where price intervention shocks are averaged across crops using output-based weights, and land subsidy ones using land weights. In both cases, results match the ones obtained using the event study design.

Table A11: Difference-in-Difference Results

Dependent Variables:	Output Price (log)	Sales (log)	Profit (log)	Chemical Expenditure (log)
Model:	(1)	(2)	(3)	(4)
Variables				
$Exposure_j \times Post_t$	-0.2380*	-0.2376**	-0.3117***	-0.2473
	(0.1334)	(0.1041)	(0.0996)	(0.1482)
Mean Outcome in Level	224	109k	50k	17k
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Department-Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	1,929	2,978	2,945	2,978
R^2	0.94	0.94	0.88	0.93

Clustered (Department) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Outcomes are the following: a farm-level output price index (average price across row crops, using relative areas as weights), total sales, profit, and the farms' deflated chemical bill. All coefficients correspond to the period-specific coefficient associated to farm exposure, and give the effect of exposure in that given year. The vector of controls contains the following elements: farm 1991 capital stock, total labor used, total land use, profit, their chemical use, the number of crops they grew and the evenness of their land allocation, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 adoption trends in chemicals measured as the evolution in their chemical use. I add farm and department-year fixed effects, and cluster the standard errors at the department level level.

Table A12: Difference-in-Difference Results (Alternative Exposure)

Dependent Variables:	Output Price (log)	Sales (log)	Profit (log)	Chemical Expenditure (log)
Model:	(1)	(2)	(3)	(4)
Variables				
$Exposure_j^{alt} \times$ post	-0.1928*	-0.2322**	-0.3228***	-0.2357
	(0.1042)	(0.1042)	(0.1032)	(0.1422)
Mean Outcome in Level	224	109k	50k	17k
Fixed-effects				
Farm	Yes	Yes	Yes	Yes
Department-Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	1,929	2,978	2,945	2,978
R^2	0.94	0.94	0.88	0.94

Clustered (Department) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Outcomes are the following: a farm-level output price index (average price across row crops, using relative areas as weights), total sales, profit, and the farms' deflated chemical bill. All coefficients correspond to the period-specific coefficient associated to farm exposure, and give the effect of exposure in that given year. The vector of controls contains the following elements: farm 1991 capital stock, total labor used, total land use, profit, their chemical use, the number of crops they grew and the evenness of their land allocation, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 adoption trends in chemicals measured as the evolution in their chemical use. I add farm and department-year fixed effects, and cluster the standard errors at the department. Exposure is computed using land shares for land subsidies, and output shares for price intervention.

C.4 Independent Exposures

Table A13: Heterogeneity in Crop-Specific Effects

Dependent Variables:	Sales (log)	Profit (log)	Chemical Expenditure (log)
Model:	(1)	(2)	(3)
Variables			
$Wheat_j \times Post_t$	-0.0024	-0.0030*	-0.0032
	(0.0016)	(0.0015)	(0.0019)
$WinterBarley_j \times Post_t$	0.0002	-0.0003	-0.0010
	(0.0025)	(0.0025)	(0.0030)
$Corn_j \times Post_t$	-0.0024**	-0.0020	-0.0024
	(0.0012)	(0.0015)	(0.0019)
$Rye_j \times Post_t$	-0.0053***	0.0005	-0.0023
	(0.0018)	(0.0021)	(0.0020)
$Sunflower_j \times Post_t$	-0.0014	-0.0050***	-0.0010
	(0.0020)	(0.0013)	(0.0027)
$Colza_j \times Post_t$	-0.0030	-0.0075***	-0.0022
	(0.0019)	(0.0019)	(0.0027)
$Durum_j \times Post_t$	-0.0093***	-0.0077***	-0.0069***
	(0.0025)	(0.0018)	(0.0026)
$Sorghum_j \times Post_t$	0.0161	0.0746***	0.0060
	(0.0147)	(0.0128)	(0.0219)
$SpringBarley_j \times Post_t$	-0.0006	-0.0015	-0.0010
	(0.0008)	(0.0014)	(0.0024)
$Oats_j \times Post_t$	-0.0009	-0.0046***	0.0015
	(0.0021)	(0.0016)	(0.0018)
Fixed-effects			
Farm	Yes	Yes	Yes
Department-Year	Yes	Yes	Yes
Fit statistics			
Observations	2,978	2,945	2,978
R^2	0.94	0.88	0.94

Clustered (Department) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. I show the heterogeneity of the difference-in-difference results across crop-specific shocks. Exposure is decomposed in multiple crop-specific measures. I include the crop shares in levels, but do not show the coefficients in the table. The vector of controls contains the following elements: farm 1991 capital stock, total labor used, total land use, profit, their chemical use, the number of crops they grew and the evenness of their land allocation, their fertilizer-to-land and pesticides-to-land ratios, the share of their production which corresponds to oil crops (colza and sunflower), and the farms' 1983-1984 adoption trends in chemicals measured as the evolution in their chemical use. I add farm and department-year fixed effects, and cluster the standard errors at the department level.

I decompose the shift-share variable into crop-specific variables. Each crop-specific exposure measure corresponds to the crop-specific change in subsidization between 1991 and 1995,

interacted by the land share allocated to that crop by the farm in 1991. The regression is otherwise similar to the difference-in-difference regressions presented previously. The coefficients for the crop exposures interacted with the post-1991 dummy isolate the effect of crop-specific exposure to the reform after the reform happened. All crop-specific exposures but sorghum and winter barley have negative effect on sales, and these two are non statistically significant. Sorghum is the one crop with a a positive effect for profit. This means that the shift-share exposure hides some degree of heterogeneity in terms of the effects of exposure. I note however, that sorghum production in France corresponds to about 1.5% of the overall production of the grain market, and hence does not correspond to a very significant share of the market. The crop-by-crop results otherwise describe a high level of homogeneity in the consequences of exposure to the reform.

C.5 Municipality-Level Results

In this section, I first present balance tests comparing municipality with relatively more and relatively less exposure to the reform, and then the distribution of the measure of exposure used for the municipality-level design. I end by showing the shift-share results using different aggregate measures to go from farm to municipality.

Starting with balance tests, I show the trends (averages for each wave of the census) for six municipality characteristics, splitting them nationally between municipalities with a median exposure above or below the French median in 1988. Trends are overall similar, apart from the evenness variable. In the estimation design, I control for this evenness, both at the municipality and farm level.

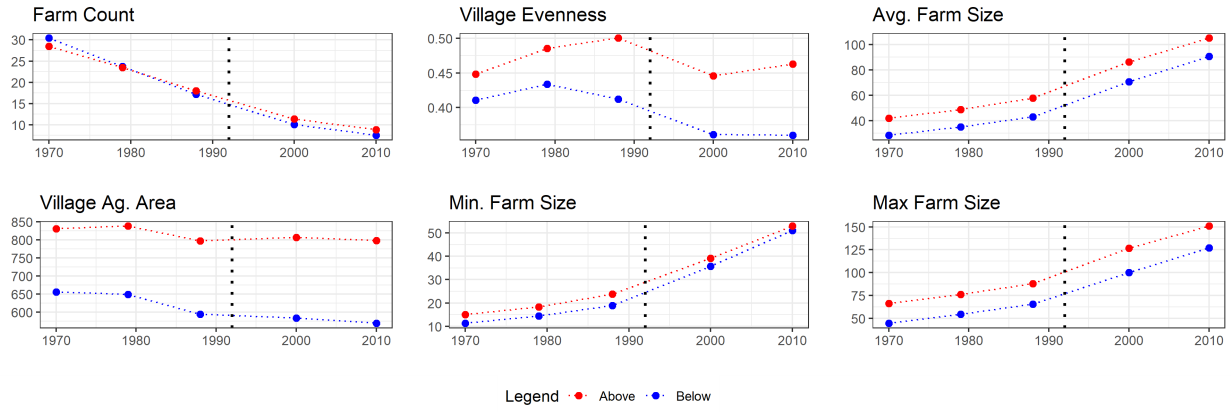


Figure A20: Municipality-Level Trends

Notes: This figure gives the evolution of municipality level outcomes over time, for the group of municipalities for which median exposure is above or below the median exposure in France as a whole.

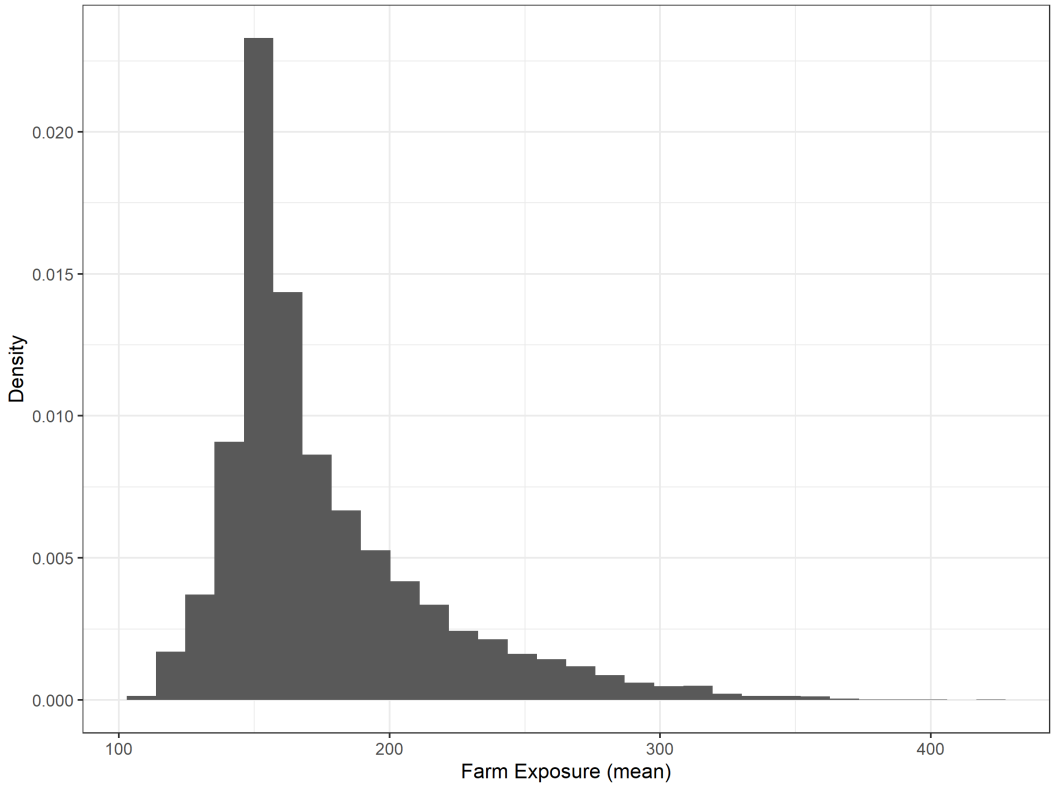


Figure A21: Distribution of Median Exposure

Notes: This figure gives the distribution of the municipality-level exposure to the MacSharry reform.

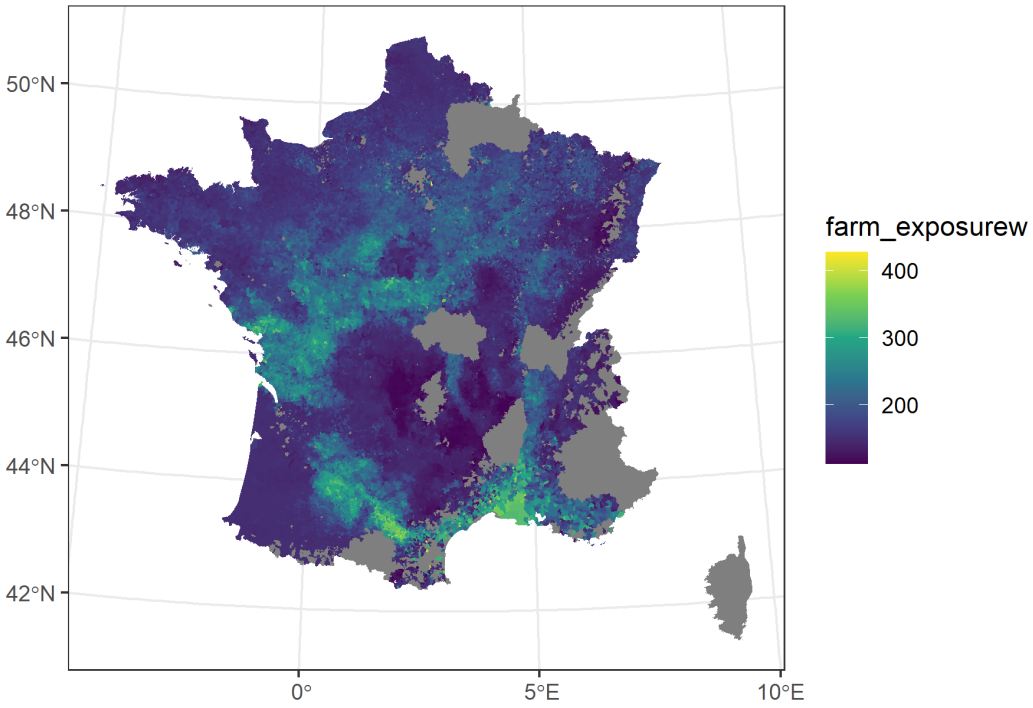


Figure A22: Geographic distribution of Mean Exposure

Notes: This figure gives the distribution of the municipality-level exposure to the MacSharry reform.

Next, I show the distribution of median municipality exposure. As mentioned previously, all the measures of municipality-level exposure are computed after winsorizing the farm-level measure of exposure in the Census for the bottom and top 1ppt values. Exposure is mostly distributed between 100 and 200€ per output unit, with a large right tail which likely corresponds to municipalities which have some of their land allocated either to colza, durum and sunflower, which are the crops with a shock value higher than 100 € per unit. These are mostly located in two areas of France, as shown in Figure A22, along the Mediterranean coast, as well as the main grain region of France around the Beauce region below Paris.

D Model Details and Extensions

I start with the Bellman decisions describing the farms' decisions—both incumbent and entrant farms—which correspond to the general outline of the model from Section 4.2.

The incumbent's decision can be characterized by:

$$V(\Upsilon_{jt};\Omega_{jt}) = \max_{\{X_{jct}\}_{\mathbb{C}_{jt}},\xi^x_{jt},K_{jt+1},\mathbb{C}_{jt+1}} \underbrace{\Pi_{jt}}_{\text{Static Profit}} - \underbrace{C(K_{jt},K_{jt+1})}_{\text{Capital Adjustment}}$$
$$+ \xi^x_{jt}\underbrace{\left[(1-\delta_k)K_{jt}P^K_{jt}\right]}_{\text{Scrap Value if Exit}} + (1-\xi^x_{jt})\underbrace{\left(\beta\mathbb{E}\left[V(\Upsilon_{jt+1};\Omega_{jt+1})|\Upsilon_{jt},\Omega_{jt}\right] - f_k\right)}_{\text{Expected Value of Continuation}}$$

And I can further write static profit as:

$$\Pi_{jt} = \sum_{c\in\mathbb{C}_{jt}} \underbrace{\left[P_{jct}Q_{jct} - P^L_{jt}L_{jct} - (1-\tau^c_s)P^S_{jt}S_{jct} - P^F_{jt}F_{jct} - P^P_{jt}Pest_{jct}\right]}_{\text{Crop Specific Profit: } \pi_{jct}}$$

The entrant's problem takes the form:

$$V_e(\Omega_{jt}) = \max_{\xi^e_{jt},\xi^x_{jt},K_{jt+1},\mathbb{C}_{jt+1}} \mathbb{E}\Bigg[\xi^e_{jt}\Bigg\{ - f_e - \underbrace{C(0,K_{jt+1})}_{\text{Capital Adjustment}}$$
$$+ (1-\xi^x_{jt})\underbrace{\beta\mathbb{E}\left[V(\Upsilon_{jt+1};\Omega_{jt+1})|\Upsilon_{jt},\Omega_{jt}\right]}_{\text{Expected Value of Continuation}}\Bigg\}|\Omega_{jt}\Bigg]$$

The timing assumptions which relate to these Bellman equations are represented in Figure A23 and Figure A24 below.

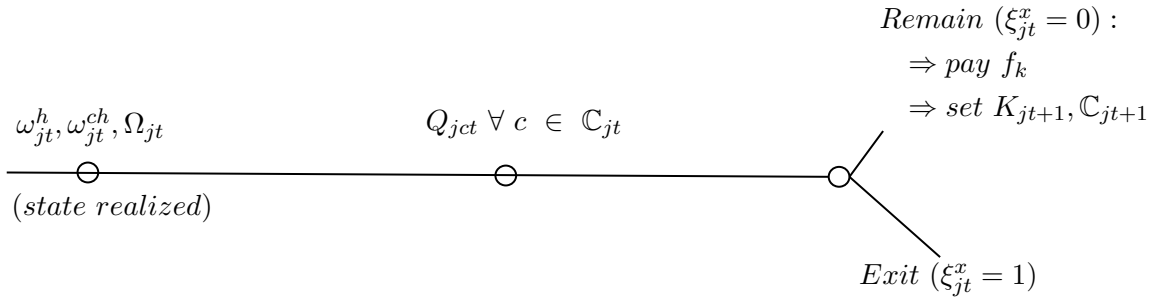


Figure A23: Timing for Incumbent Farms

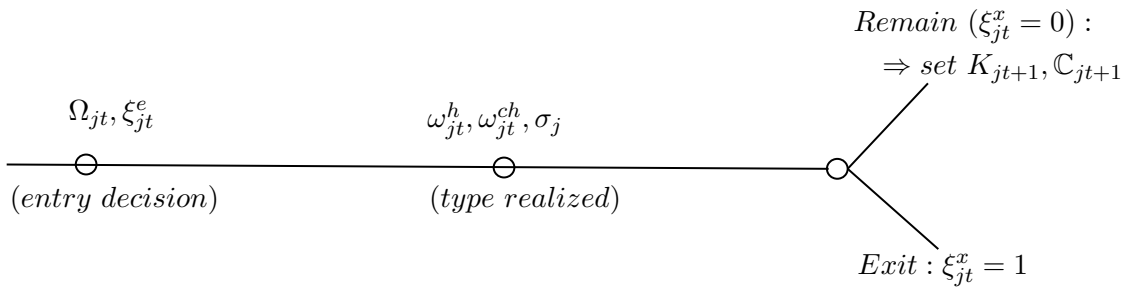


Figure A24: Timing for Entrants

After the simplifying assumptions regarding the homogeneity of σ_j and output prices, I can further write the transition operators which describe the evolution of entrant and incumbents farms as: $\phi_t(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}, K, \mathbb{C})$ and $\phi_{e,t}(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch})$. They require the additional introduction of the transition operator $J^h(.|.)$ and $J^{ch}(.|.)$ for both shocks, and of the policy functions κ^k for incumbent capital investment, $\kappa^{\mathbb{C}}$ for crop choice, and κ^k_e for entrant capital investment. These transition operators now write as:

$$\begin{aligned}\phi_t(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}, K, \mathbb{C}) =& J^h(\omega'|\omega) J^{ch}(\omega^{ch'}|\omega^{ch}) \mathbb{1}\{K' = \kappa^k(\omega, \omega^{ch}, K|\Omega_t)\} * \\ & \mathbb{1}\{\mathbb{C}' = \kappa^{\mathbb{C}}(\omega, \omega^{ch}, K|\Omega_t)\} \mathbb{1}\{\xi^x(\omega, \omega^{ch}, K|\Omega_t) = 0\}\end{aligned}$$

I note that crop choice is a dynamic decision simply because of the timing assumption—which allows me to address the selection bias induced by focusing on single crop farms for estimation—but crop choice has no impact on a farm's transition. The current state of a farm's crop mix in a given period has no bearing on the dynamic decisions made by the farm regarding the next period, including its crop mix decision. $\phi_t(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}, K, \mathbb{C})$ could as such be re-written as $\phi_t(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}, K)$, though I will keep the inclusion of $\mathbb{C}$ in the conditioning for clarity.

The analogue transition operator for entrants writes $\phi_{e,t}(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch})$:

$$\begin{aligned}\phi_{e,t}(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}) =& J(\omega'|\omega) J^{ch}(\omega^{ch'}|\omega^{ch}) \mathbb{1}\{K' = \kappa^k_e(\omega, \omega^{ch}, 0|\Omega_t)\} * \\ & \mathbb{1}\{\mathbb{C}' = \kappa^{\mathbb{C}}(\omega, \omega^{ch}, 0|\Omega_t)\} \mathbb{1}\{\xi^x(\omega, \omega^{ch}, 0|\Omega_t) = 0\}\end{aligned}$$

Once I have defined these, I can then define the law of motion for the measure μ_t. $G_1(.)$ and $G_2(.)$ are the distributions from which entrants draw resp. their original TFP and chemical productivity. With N^e_t the mass of entrants, the allocation of farms over the market space is then fully described by the measure μ_t, whose transition operation $T(.)$ can be written as:

$$\mu_{t+1}(\omega', \omega^{ch'}, K', \mathbb{C}') = \int \phi_t(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}, K, \mathbb{C}) d\mu_t(\omega, \omega^{ch}, K, \mathbb{C}) + N_t^e \int \int \phi_{e,t}(\omega', \omega^{ch'}, K', \mathbb{C}'|\omega, \omega^{ch}) dG_1(\omega) dG_2(\omega^{ch})$$

With all these elements, I can finally define the stationary equilibria on which I will focus as follows. I look at stationary competitive equilibria composed of the tuple of $\Omega^\star = \{\mu^\star, N^{e,\star}, \{P_c^\star\}_{\mathbb{C}}, \{P^x\}_x, Policy\}$, such that—for a given policy state ($Policy$) and given input prices ($\{P^x\}_x$):

- **Zero Ex-Ante Profit**: $V^e(\Omega^\star) \leq 0$, and $V^e(\Omega^\star) = 0$ if $N^{e,\star} > 0$.
- **Stationary Market**: $\mu^\star = T(\Omega^\star)$.
- **Output Market Clearing**: $P^\star$ s.t. $Q_c^S(\Omega^\star) = Q_c^D(P_c^\star) \ \ \forall c \in \mathbb{C}$.

D.1 Measuring Farm Efficiency

It is useful to derive the expressions for the average cost of production, marginal cost of production and unit cost of production in order to understand how farm efficiency will be measured in the counterfactuals. For simplicity, I will remove the subscripts $\{j, c, t\}$ in the following derivations, and take $\varepsilon = 0$.

In a first step, I will recover the input price for the composite input $\tilde{S}$, which is defined as the total expenditure required to produce one unit of $\tilde{S}$. This input price can hence be recovered by solving the unit cost problem, taking advantage of the fact that the expression for $\tilde{S}$ has constant returns to scale:

$$\min_{S,Chemicals} Cost_S = P_s S + P_{ch} Chemicals \ s.t. \ 1 = \left[\delta_s S^\rho + (1-\delta_s)\{e^{\omega^{ch}} Chemicals\}^\rho\right]^{\frac{1}{\rho}}$$
$$\Leftrightarrow \min_{S,Chemicals} Cost_S = P_s S + P_{ch} Chemicals \ s.t. \ 1 = \delta_s S^\rho + (1-\delta_s)\{e^{\omega^{ch}} Chemicals\}^\rho$$

The Lagrangian can hence be written as:

$$\mathcal{L} = P_s S + P_{ch} Chemicals - \lambda\left(\left[\delta_s S^\rho + (1-\delta_s)\{e^{\omega^{ch}} Chemicals\}^\rho\right] - 1\right).$$

Combining the first order conditions of the problem, I can write the price of the composite input $P_{\tilde{S}}$ as:

$$P_{\tilde{S}} = P_S S + P_{ch} Chemicals = \lambda\rho.$$

Introducing the elasticity $\sigma = \frac{1}{1-\rho}$, I then solve for the expression of λ:

$$\lambda = \rho^{-1}\left[\delta_s^{\sigma} P_s^{1-\sigma} + (1-\delta_s)^{\sigma}(P_{ch}e^{-\omega^{ch}})^{1-\sigma}\right]^{\frac{1}{1-\sigma}}.$$

This implies finally:

$$P_{\tilde{S}} = \left[\delta_s^{\sigma} P_s^{1-\sigma} + (1-\delta_s)^{\sigma}(P_{ch}e^{-\omega^{ch}})^{1-\sigma}\right]^{\frac{1}{1-\sigma}}.$$

Now, we can consider the profit maximization problem faced by the farm in a given production period. Within the period, the farm takes its capital stock K as given. This implies that the production function has decreasing returns in the flexible inputs (labor and the composite input). The cost minimization problem writes:

$$\underset{L,\tilde{S}}{min}\ C = P_l L + P_{\tilde{S}}\tilde{S} \quad s.t. \quad Q = e^{\omega^h} K^{\alpha_k} L^{\alpha_l} \tilde{S}^{\alpha_s}$$

It is useful to write $\alpha = \alpha_l + \alpha_s$, as well as: $\alpha_l' = \frac{\alpha_l}{\alpha}$, and $\alpha_{s'} = \frac{\alpha_s}{\alpha}$. The normalized production function with constant returns to scale $Y = Q^{\frac{1}{\alpha}}$ has a constant marginal cost of production:

$$c_y = \left(e^{-\omega^h} K^{-\alpha_k}\right)^{\frac{1}{\alpha}} \left(\alpha_l'^{-\alpha_l'} \alpha_s'^{-\alpha_s'}\right) P_l^{\alpha_l'} P_{\tilde{s}}^{\alpha_s'}.$$

Given the transformation, the same set of inputs which produces Y with the normalized function will also produce Q with the main function. With $C_Q(.)$ the cost function of function Q, and $C_Y(.)$, the cost function of function Y, we always have that:

$$C_Q(Q) = C_Y(Y) = c_y Y = c_y Q^{\frac{1}{\alpha}}$$

This implies first that the average cost is $AC(Q) = \frac{C_Q(Q)}{Q} = c_y Q^{\frac{1-\alpha}{\alpha}}$ which is upward sloping in the quantity produced, and that the marginal cost is $MC(Q) = \frac{c_y}{\alpha} Q^{\frac{1-\alpha}{\alpha}}$. The average cost is then always a fraction of the marginal cost such that $AC(Q) = \alpha MC(Q)$.

The model has no distortions, the market is competitive and all farms face a unique price for a given crop. This implies that at equilibrium, all farms will produce such that their marginal cost is equal to the market price P. This implies that the average cost of production will be equal across all farms, and will correspond to a fraction of the market price:

$$AC(Q^{\star}) = \alpha P = \left(\alpha_l + \alpha_s\right) P.$$

There are no misallocations in the model, in the sense that inputs are allocated across incumbents in a way that minimizes overall costs of production. This equality also implies that the average cost of production cannot be used as a measure of the farms' production efficiency.

In light of this, I use the farms' unit cost of production as a measure of their heterogeneity in production efficiency. This cost is defined as the cost to produce one unit of a given crop. Denoting this unit cost c, it writes:

$$c = C(Q = 1) = c_y$$
$$c = \left(e^{-\omega^h} K^{-\alpha_k}\right)^{\frac{1}{\alpha}} \left(\alpha_l'^{-\alpha_l'} \alpha_s'^{-\alpha_s'}\right) P_l^{\alpha_l'} \left[\delta_s^\sigma P_s^{1-\sigma} + (1-\delta_s)^\sigma (P_{ch} e^{-\omega^{ch}})^{1-\sigma}\right]^{\frac{\alpha_s'}{1-\sigma}}.$$

As expected, this unit cost is decreasing in a farm's TFP (ω^h), in its capital stock (K), and in its chemical use efficiency (ω^{ch}).

D.2 Joint Production Estimation Framework

Here I describe an alternative joint production framework, for which I also provide estimates. This approach combines the work of Dhyne et al. (2022) and of Doraszelski and Jaumandreu (2018). From the first, I use the idea that inputs can be shared across product lines, and that the parametrization of a firm-level transformation function will define the degree of penalization that accompanies this sharing of inputs. I then draw from Doraszelski and Jaumandreu (2018) to model and estimate non-hicksian shocks, within that multi-product production framework. Because the recovery of non-hicksian shocks relies heavily on the first-order conditions taken from the firm-level optimization problem, I move away slightly from the exact specification used by Dhyne et al. (2022).

In effect, consider a firm j which produces the set of crops $\mathbb{C}_{jt}$ within a production period t. The firm is endowed with product-specific input-neutral shocks $\tilde{\omega}_{jct}$. I write the transformation function as:

$$f_{jct} = exp(\tilde{\omega}_{jct}) g_{jct}^{\gamma_{cc}} \prod_{c' \in \mathbb{C}_{jt}, c' \neq c} g_{jc't}^{\gamma_{cc'}}$$

Here f_{jct} is the quantity of crop c produced. That quantity depends on a specific TFP shock, and on product-specific functions g_{jct} which are evaluated at the level of inputs used for each crop grown in the season. The parameters γ_{cc} parametrize the rivalry or complementarity in production from jointly producing c and c', and potentially sharing some inputs across these two product lines. I keep the functions g_{jct} fairly general for the moment, and parametrize

them further later. One additional set of assumptions I make is that:

$$g_{jct} = g_c(K_{jt}, L_{jt}, Fert_{jt}, Pest_{jt}, S_{jct}, \omega^p_{jt}, \omega^f_{jt})$$

This means that each function g_{jct} depends on the firm-level chosen amounts of capital K_{jt}, labor L_{jt}, fertilizer $Fert_{jt}$, pesticides $Pest_{jt}$, on a crop-specific chosen amount of land S_{jct}, and on firm-level pesticide and fertilizer productivity shocks that I denote ω^p_{jt} and ω^f_{jt}. This implies that all inputs but land are selected at the firm-level, and shared across product-lines according to the γ matrix. I also consider that the g_c function only varies across crops, but not periods or farms, and that these two non-hicksian shocks fully account for the heterogeneity in these functions across farms. This assumption of land being a crop-specific input, while others are shared with the same penalty γ allows me to derive an estimating equation to recover the non-hicksian shocks.

With this specification, the firm-level profit is as follows:

$$\Pi_{jt} = \left(\sum_{c \in \mathbb{C}_{jt}} P_{jct}(f_{jct}, D_{jct}) f_{jct} - P^S_{jt} S_{jct} \right) - P^L_{jt} L_{jt} - P^F_{jt} Fert_{jt} - P^p_{jt} Pest_{jt}$$

Taking $\{L_{jt}, Fert_{jt}, Pest_{jt}, S_{jct}\}$ as flexible inputs, their setting will only impact within period profit. I write $P_{jct}(f_{jct}, D_{jct})$ for the inverse demand function, known to the farmers, which depends on the volume produced and sold f_{jct}, and a demand shock D_{jct}. Writing X_{jt} for any of the public inputs $\{L_{jt}, Fert_{jt}, Pest_{jt}\}$, I get the following expression:

$$\frac{\partial \Pi_{jt}}{\partial X_{jt}} = 0 \Leftrightarrow P^X_{jt} = \sum_{c \in \mathbb{C}_{jt}} \frac{\partial g_{jct}}{\partial X_{jt}} g^{-1}_{jct} \sum_{c' \in \mathbb{C}_{jt}} \gamma_{cc'} f_{jc't} P_{jc't} \left[1 - \frac{1}{\eta_{jc't}} \right]$$

I denote by η_{jct} the absolute value of the elasticity of demand. I then look at land, the only private input in this set-up:

$$\frac{\partial \Pi_{jt}}{\partial S_{jct}} = 0 \Leftrightarrow P^S_{jt} = \frac{\partial g_{jct}}{\partial S_{jct}} g^{-1}_{jct} \sum_{c' \in \mathbb{C}_{jt}} \gamma_{c'c} f_{jc't} P_{jc't} \left[1 - \frac{1}{\eta_{jc't}} \right]$$

Combining these two equations, and under symmetry of the γ matrix, I get:

$$\frac{P^X_{jt}}{P^S_{jt}} = \sum_{c \in \mathbb{C}_{jt}} \frac{\frac{\partial g_{jct}}{\partial X_{jt}}}{\frac{\partial g_{jct}}{\partial S_{jct}}}$$

This relation is quite intuitive, the relative allocation of inputs has to be such that the ratio of their marginal impacts on production across production lines equals the ratio of their

prices.

To go further, I assume a specific shape for the g_{jct} function[56]. I take:

$$g_{jct} = K_{jt}^{\alpha_k} L_{jt}^{\alpha_L} \left\{ \delta_s^c S_{jct}^{\rho} + \delta_p^c \left[e^{\omega_{jt}^p} Pest_{jt} \right]^{\rho} + \delta_f^c \left[e^{\omega_{jt}^f} Fert_{jt} \right]^{\rho} \right\}^{\frac{\alpha_s}{\rho}}$$

Without relying on a translog specification, this is the simplest specification which allows me to model non-hicksian productivity shocks. I note that the heterogeneity in the g_{jct} across crops only relies on the $\{\delta_s^c, \delta_f^c, \delta_p^c\}_c$ parameters. Because I am mostly worried about variations in crop-composition impacting the chemicals-to-land ratios at the farm-level, this is the main form of heterogeneity I am interested in. A more expansive heterogeneity could specify crop-specific elasticities ρ_c, which I assume away here.

With this parametrization, I obtain the following estimating equation, where small cap letters represent logs, with $x \in \{pest, fert\}$:

$$p_{jt}^x - p_{jt}^s = \rho \omega_{jt}^x + (\rho - 1) x_{jt} + log\Big(\sum_{c \in \mathbb{C}_{jt}} \frac{\delta_x^c}{\delta_s^c} S_{jct}^{1-\rho} \Big)$$

This equation mirrors the first stage estimating equation from Doraszelski and Jaumandreu (2018), adapted to the parametrized multi-product setting. Here, the input ratio can either be affected by variations in the input price ratio, by changes in the establishment's production mix, or by changes in the establishment's proficiency in using input x in the production process. Similarly to them, when the inputs of interest are substitutes, and holding everything else constant, an increase in ω_{jt}^x will increase the ratio of x_{jt} to land.

I propose to recover the $\{\delta_s^c, \delta_f^c, \delta_p^c\}_c$, and ρ in a first estimation step, and to recover the remaining parameters in a second step. To simplify the estimation process, I assume that $\gamma_{cc} = 1$, and $\gamma_{cc'} = \gamma$, and verify post-estimation that $\gamma_{cc'} \in]-\frac{1}{dim(\mathbb{C})}, 0[$ to guarantee that the second order conditions from the optimization problem are met.

First Step: I specify a structure for the two input-biased productivity shocks. For each of them, I assume an AR(1) structure such that, with ζ_{jt+1}^x an exogenous i.i.d. innovation shock:

$$\omega_{jt+1}^x = \mathbb{E}[\omega_{jt+1}^x | \omega_{jt}^x] + \zeta_{jt+1}^x = g^x(\omega_{jt}^x) + \zeta_{jt+1}^x$$

I take $g^x(.)$ to be a third-order polynomial. I use the following moments, with A_{jt}^1 the matrix of instruments:

$$\mathbb{E}[(\zeta_{jt}^p + \zeta_{jt}^f) A_{jt}^1] = 0$$

[56]Note that this framework can also accomodate a full CES specification. Given the difficulty of estimating the parameters within a CES nest, I keep the CES structure to a minimum - the land nest being itself estimated within an equation which is linear in part of the parameters.

My instruments match the production function literature, and correspond to lagged firm-level observables presumably uncorrelated to the productivity innovations, as well as the current values of farm-level land prices and hourly agricultural wages.

Second Step: For the second step, I take the log of the production function, and re-order the elements to get:

$$\tilde{\omega}_{jct} = log(f_{jct}) - \left(\sum_{c' \in \mathbb{C}_{jt}} \gamma_{cc'} \right) \left[\alpha_K k_{jt} + \alpha_L l_{jt} \right] - \sum_{c' \in \mathbb{C}_{jt}} \gamma_{cc'} \tilde{s}_{jc't}$$

Where $\tilde{s}_{jc't}$ is the log of the land-fertilizer-pesticide nest, which I can compute using the estimates from the first step. Now, I use crop-specific policy cushions, as well as land prices and agricultural wages as instruments denoted A^2_{jt} in a GMM estimation, with moments:

$$\mathbb{E}\left[\tilde{\omega}_{jct} A^2_{jt} \right] = 0$$

I obtain the following results:

Table A14: Joint Production - Parameters

	Coefficient	se
α_k	0.2415	0.0006
α_l	0.4812	0.0015
α_s	0.0931	$2.2e^{-5}$
γ	-0.0904	0.00013
ρ	0.3886	0.00073
δ_f - Cereals	0.2428	0.0027
δ_f - Oil/Protein	0.2916	0.0034
δ_f - Industrial	0.0282	0.0002
δ_p - Cereals	0.1789	0.0025
δ_p - Oil/Protein	0.2925	0.004
δ_p - Industrial	0.1024	0.002

Standard errors from the second step are corrected for the two-step procedure, following Doraszelski and Jaumandreu (2018), and the coefficients associated to the input-biased productivity processes are concentrated out, and not estimated. I note a negative value for γ indicating that the more you share the public inputs across production lines, the less a given crop will benefit from that input (either because of an increase in scope, or an increase in the scale of the other product's production levels). I also note that $\rho \in]0, 1[$, indicating that land, pesticides and fertilizers are substitutes in production, which was expected.

D.3 Extended Production Function

The production function for quantity Q_{jct} of crop c for farm j at time t remains Cobb-Douglas in its main inputs:

$$Q_{jct} \;=\; e^{\omega^h_{jct}}\; K_{jt}^{\alpha_k^c}\; L_{jct}^{\alpha_l^c}\; \tilde{S}_{jct}^{\alpha_s^c}\; e^{\varepsilon_{jct}}, \tag{12}$$

with L_{jct} as labor, $\sum_{k\in\{k,l,s\}} \alpha_k^c = 1$, and ε_{jct} is a mean-zero, i.i.d. shock is realized after input decisions are made. The key mechanisms of substitution and biased productivity are embedded in the composite input $\tilde{S}_{jct}$, which combines land and chemicals. This composite input is now modeled as two nested constant elasticity of substitution functions, in order to allow for flexible substitution patterns and chemical-biased productivity shocks:

$$\tilde{S}_{jct} = \left\{\delta_s^c S_{jct}^{\rho} + (1-\delta_s^c)\left(e^{\omega^{ch}_{jt}}\left\{\delta_p^c Pest_{jct}^{\rho_2} + (1-\delta_p^c) Fert_{jct}^{\rho_2}\right\}^{\frac{1}{\rho_2}}\right)^{\rho}\right\}^{\frac{1}{\rho}}, \tag{13}$$

This function has two layers. The inner nest combines pesticides ($Pest_{jct}$) and fertilizers ($Fert_{jct}$) into a chemical aggregate, with an elasticity of substitution governed by ρ_2. The outer nest then combines physical land (S_{jct}) with this chemical aggregate, with an elasticity of substitution governed by ρ. The chemical-biased productivity shock (ω^{ch}_{jt}) still acts as an augmenting shifter on the entire chemical aggregate.

The second step of the estimation is similar to the one used in the main estimation procedure described in the paper. I obtain the following results:

Table A15: Alternative Production Function (Additional Nest)

Coefficient	Parameter	Estimate	Std. Error
Substitution Land-Chemicals	ρ	0.3425	(0.2033)
Substitution Fertilizers-Pesticides	ρ_2	-1.8104	(1.0575)
Land Share (Wheat)	δ_s^{wheat}	0.3952	(0.2457)
Land Share (Others)	δ_s^{other}	0.5001	(0.3285)
Pesticide Share (Wheat)	δ_p^{wheat}	0.1189	(0.0924)
Pesticide Share (Others)	δ_p^{other}	0.0438	(0.1007)
Labor Share (Wheat)	α_L^{wheat}	0.1124	(0.2148)
Labor Share (Others)	α_L^{other}	0.3835	(0.2723)
Land-Nest Share (Wheat)	α_S^{wheat}	0.7785	(0.0832)
Land-Nest Share (Others)	α_S^{other}	0.3240	(0.0792)
N	7625		
Unique Farms	2671		
N - Group 1	2511		
N - Group 2	5114		

Notes: The parameters are obtained from a two-step estimation, each step performed by GMM. The estimation is run on the

FADN French sample restricted to 1980-2006. Observations prior to 1980 are removed, as they do not contain output price data. Observations post-2006 are removed, to focus on a period with significant variation in EU agricultural subsidies, both across time and across crops. I keep farms observed for at least three periods in a row[57]*, and that produce either only crops in the wheat group, or in the other crops group. I remove farms that are not observed with positive input values for the set of considered inputs (land, labor, capital, fertilizers, pesticides). standard errors are obtained using a block bootstrap procedure, where I draw all the observations related to a farm at a time, using* $B = 1000$.

Recovering Input Allocations: As discussed in Section 4.3, I estimate the production parameters on single crop-group firms, following De Loecker et al. (2016). Because I observe farm-crop level land allocations, knowing the parameters of the production specification, I can then recover both input allocations and hicksian and non-hicksian productivity shocks for multi-product firms. The production function has the following shape:

$$Q_{jct} = e^{\tilde{\omega}^h_{jct}} K_{jt}^{\alpha^c_K} L_{jct}^{\alpha^c_L} \left\{ \delta^c_s S^{\rho}_{jct} + (1-\delta^c_s) \left(e^{\omega^{ch}_{jt}} \left[(1-\delta^c_p) Fert^{\rho_2}_{jct} + \delta^c_p Pest^{\rho_2}_{jct} \right]^{\frac{1}{\rho_2}} \right)^{\rho} \right\}^{\frac{\alpha^c_S}{\rho}}$$

$$\text{With: } \tilde{\omega}^h_{jct} = \omega^h_{jct} + \varepsilon_{jct}$$

Drawing from the first-order conditions, I can write the following relation between pesticides and fertilizers:

$$Fert^{1-\rho_2}_{jct} = \frac{P^P_{jt}}{P^F_{jt}} \frac{(1-\delta^c_p)}{\delta^c_p} Pest^{1-\rho_2}_{jct} \tag{14}$$

I can use the relation to recover a relation between land and pesticide allocations:

$$\frac{P^s_{jt} S^{1-\rho}_{jct}}{\delta^c_S} = \frac{P^P_{jt}}{\delta^c_p (1-\delta^c_s)} \left(e^{-\omega^{ch}_{jt}} \right)^{\rho} Pest^{1-\rho}_{jct} \left[(1-\delta^c_p) \left(\frac{P^P_{jt}(1-\delta^c_p)}{P^F_{jt}\delta^c_p} \right)^{\frac{\rho_2}{1-\rho_2}} + \delta^c_p \right]^{1-\frac{\rho}{\rho_2}}$$

For farms producing the two crop groups, I can then take the ratio of this expression for their two crop categories c and c':

$$\frac{Pest^{1-\rho}_{jct}}{Pest^{1-\rho}_{jc't}} = \frac{S^{1-\rho}_{jct}}{S^{1-\rho}_{jc't}} \frac{\delta^c_p (1-\delta^c_s) \delta^{c'}_s}{\delta^{c'}_p (1-\delta^{c'}_s) \delta^c_s} \frac{\left[(1-\delta^c_p) \left(\frac{P^P_{jt}(1-\delta^c_p)}{P^F_{jt}\delta^c_p} \right)^{\frac{\rho_2}{1-\rho_2}} + \delta^c_p \right]^{\frac{\rho}{\rho_2}-1}}{\left[(1-\delta^{c'}_p) \left(\frac{P^P_{jt}(1-\delta^{c'}_p)}{P^F_{jt}\delta^{c'}_p} \right)^{\frac{\rho_2}{1-\rho_2}} + \delta^{c'}_p \right]^{\frac{\rho}{\rho_2}-1}}$$

Once I have estimated the parameters of the production function, the right hand side of this expression is fully known. As I know total farm pesticide use, I can then recover the volumes of pesticide used for each of the crops that they grow. I can then use Equation (14)

[57] This is done to smooth potential measurement issues, and follows from De Loecker et al. (2016) and Doraszelski and Jaumandreu (2018).

to recover the fertilizer allocations as well. With these, I can construct the crop-specific land nest values, and use the outer Cobb-Douglas shape of the production function to recover crop-specific labor allocations, from these crop-specific land nests.

D.4 Selection Correction for the Production Function Estimation

I need to account for selection while building moments for the second step of the estimation procedure. Using $\Xi_{jct} = 1$ a dummy for the fact that j only produces crop c in period t, I want to build moments based on the following corrected process:

$$\omega^h_{jct+1} = \mathbb{E}[\omega^h_{jct+1}|\omega^h_{jct}, \Xi_{jct+1} = 1] + \xi^h_{jct+1} \tag{15}$$

With I_{jt} the farm's information set at the end of period t, I want to base the moments on the following derivation—with $\overline{\omega}^h_{jcc't+1}$ the threshold used at the end of period t to decide on the inclusion of crop c' on top of c. This $\overline{\omega}^h_{jcc't+1}$ corresponds to the minimum value the farm's TFP must take (ω^h_{jt}), for c' to be added in the crop mix as a second crop, when c is the first crop included in the mix. Because the crop mix is set at the end of the period for the next, ω^h_{jt} rather than ω^h_{jt+1} is the TFP value compared to the threshold $\overline{\omega}^h_{jcc't+1}$. This threshold will be a function of both input and output prices, agricultural subsidies, the farm's capital stock and chemical-use efficiency.

Here c is the ex-ante most profitable crop to grow in period t, and c' the second (out of two crops). The farms' competence ladder is both farm-specific and time-constant, which means that there will be a time-constant relation between the farm's TFP between c and c'. I also note that because demand and input prices play a role in crop choice—the core competence of a farm might not always match the identity of the unique crop it grows. I can write, using the log of the production function and the process specified in Equation (15):

$$\begin{aligned}
\xi^h_{jct+1} + \varepsilon_{jct+1} &= q_{jct+1} - \alpha^c_k k_{jt+1} - \alpha^c_l l_{jct+1} - \alpha^c_s \tilde{s}_{jct+1} - \mathbb{E}[\omega^h_{jct+1}|I_{jt}, \Xi_{jct+1} = 1] \\
&= q_{jct+1} - \alpha^c_k k_{jt+1} - \alpha^c_l l_{jct+1} - \alpha^c_s \tilde{s}_{jct+1} - \mathbb{E}[\omega^h_{jct+1}|\omega^h_{jct}, \overline{\omega}^h_{jcc't+1}] \\
&= q_{jct+1} - \alpha^c_k k_{jt+1} - \alpha^c_l l_{jct+1} - \alpha^c_s \tilde{s}_{jct+1} - g_h(\omega^h_{jct}, \overline{\omega}^h_{jcc't+1}).
\end{aligned}$$

Where moving from line 1 to 2 relies on the first-order Markov process assumption, and the definition of the threshold $\overline{\omega}^h_{jcc't+1}$.[58] I can then specify $g_h(.)$ as a polynomial of its two terms. While Equation (7) provides an expression for ω^h_{jct}, I also need one for the threshold $\overline{\omega}^h_{jcc't+1}$. For this I derive an expression for the conditional probability of producing only crop c as a function of farm productivity and its state, focusing on the crop-addition threshold.[59]

[58] Because of the assumed competence ladder structure, firms face a unique TFP process that propagates across their crops. The process g_h is similar across crops as a direct consequence of this.

[59] Using a balanced panel of farms helps to address the selection bias induced by farm exit, when no grain production

The competence ladder structure means that the threshold for the production of the second crop c' can be expressed using the farm's observed TFP shock for crop c combined with its competence ladder σ_j—specifically by its rank on the farm's ladder ($\sigma_j^{-1}(c)$), as well as additional variables that influence the value of the threshold—input and output prices, capital stock and chemical productivity:

$$\begin{aligned}\mathbb{P}\Big(\Xi_{jct+1}=1\Big) &= \mathbb{P}\Big(\omega^h_{jct} \leq \overline{\omega}^h_{jcc't+1}(\omega^{ch}_{jct}, K_{jt+1}, \sigma_j^{-1}(c), \Omega_{jt}) | \overline{\omega}^h_{jcc't+1}(.), \omega^h_{jct}\Big) \\ &= h(\overline{\omega}^h_{jcc't+1}(.), \omega^h_{jct}) \\ &= h(\omega^h_{jct}, \omega^{ch}_{jct}, K_{jt+1}, \sigma_j^{-1}(c), \Omega_{jt}) \\ &= h(\{X_{jct}\}, \omega^{ch}_{jct}, K_{jt+1}, \sigma_j^{-1}(c), \Omega_{jt}) \\ &= SP_{jt}\end{aligned}$$

Equation (7) is used to move from the third to the fourth line. While I do not observe $\sigma_j^{-1}(c)$, I approximate it by the interaction between the farm's location and a dummy for the unique crop it produces—assuming competence ladders are determined by local soil and climatic factors. SP_{jt} is recovered non-parametrically, and as in Olley and Pakes (1996) and De Loecker et al. (2016), under some regularity conditions on the density of ω^h_{jct}, I can invert SP_{jt} to obtain a proxy for the threshold $\overline{\omega}^h_{jcc't+1}$. I can then use it in the polynomial approximation of $g_h(.)$, which depends both on the threshold and on the farm TFP.

D.5 Monopolistic Competition

An alternative to perfect competition is to adopt a monopolistic competition framework—which preserves the single agent setting—but allows for price dispersion and some horizontal differentiation across goods. While agricultural products are commodities usually thought of as homogeneous, exogenous local conditions can lead to similar crops having different moisture or nutrient content, making them more or less suitable for different uses. In that sense, a model with differentiated products might be preferable. Monopolistic competition can also allow for mark-ups that vary across producers, and along with changes in market structure.

I adopt for this section the quadratic demand of Mayer et al. (2014). The demand structure goes as follows: in a closed economy, a representative consumer has the following utility at period t:

$$U_t = q_{0,t} + \alpha \int_{i\in\Omega} q^c_{i,t} di - \frac{1}{2}\gamma \int_{i\in\Omega} (q^c_{i,t})^2 di - \frac{1}{2}\eta\Big(\int_{i\in\Omega} q^c_{i,t} di\Big)^2 \tag{16}$$

$q_{0,t}$ corresponds to the quantity of outside good consumed. Agricultural varieties $i \in \Omega$

is observed in the FADN.

are differentiated. α and η control the substitution between differentiated varieties and the outside good, while γ accounts for the degree of differentiation across varieties. Here, I note that two varieties can be two types of corn grown by two different farms, or wheat and corn grown by the same farm. I discuss next a clustering of the agricultural market, where each market correspond to a crop type, and firm produce different varieties within each crop market.

I take the outside good to be the numeraire. This preference structure implies the following demand for variety i - with $\overline{p}_t$ the average price of agricultural varieties on the market, and M_t the mass of consumed varieties:

$$q_{i,t} = \frac{\alpha}{\eta M_t + \gamma} - \frac{1}{\gamma} p_{i,t} + \frac{\eta M_t}{\eta M_t + \gamma} \frac{1}{\gamma} \overline{p}_t$$

I see that M_t and $\overline{p}_t$—market-level statistics on which atomistic farms have no individual impact—determine the characteristics of the demand curve faced by farms. From this relation, I can also derive a choke price p_t^{max} above which farms will not face any positive demand.

Estimating Demand: I use the following expression to recover the parameters of demand. From now on I denote by m the crop grown, to match the competence ladder structure that I introduce next. Because crops are differentiated in the model, each of a farm's variety is uniquely produced by that farm.

$$q_{mt} = \frac{1}{\gamma} \alpha - \frac{1}{\gamma} p_{mt} - \frac{\eta}{\gamma} Q_t$$

I take for Q_t total agricultural production sold in France, approximated using the weighted sum of sold output in the FADN, for the set of crops that I consider. q_{jmt} and p_{jmt} are respectively the quantity sold, and associated price, for each variety produced by a farm in a given period. I estimate the following regression using a two-stage least squares strategy, with exogenous shocks to firm supply as instruments. Specifically, I use local realized and forecasted weather shocks as supply shocks.

$$q_{jmt} = \beta_0 + \beta_1 p_{jmt} + \beta_2 Q_t + \varepsilon_{jmt}$$

I recover the following parameters. I only use data post-2003 for the estimation, in order to obtain demand parameters that are not impacted by intervention pricing.

Table A16: Demand - Parameters

Coefficient	Parameter	Estimate	Std. Error
Substitution with Numeraire	η	1.32e-07	6.92e-08
Substitution within Varieties	γ	3.19e-02	1.01e-02
Substitution with Numeraire	α	203.5	71.42

Notes: I recover the parameters from a 2SLS regression, using FADN data from 2003-2020, and using realized and forecasted weather data constructed at the department-year level as supply shocks. The realized weather data is constructed from the ECMWF ERA-5 hourly weather data series, and the forecasts come from ECMWF SEAS-5. Both are aggregated into growing season observations.

The standard errors for the transformed parameters are computed using the delta method. This yields a decreasing choke price over time, which drops significantly around the time of the reform.

E Decomposition

I propose a slight modification of the dynamic decomposition of Olley and Pakes (1996) and De Loecker and Collard-Wexler (2015) to highlight the role of different channels in driving the evolution of total chemical use on the market. I write evolution of market-level total chemical use between t and $t+1$, with $\mathcal{A}$ the set of farm types for which the equilibrium mass is positive across both equilibria, $\mathcal{B}$ the set of exiters, and $\mathcal{C}$ the set of entrants:

$$\begin{aligned}
\Delta Chemicals = &\sum_{j\in\mathcal{A}} Q_{jt}\left(s_{jc_1t}\Delta\frac{C}{Q}_{jc_1} + s_{jc_2t}\Delta\frac{C}{Q}_{jc_2}\right)\\
&+\sum_{j\in\mathcal{A}} Q_{jt}\left(\Delta s_{jc_1}\frac{C}{Q}_{jc_1t} + \Delta s_{jc_1}\Delta\frac{C}{Q}_{jc_1} + \Delta s_{jc_2}\frac{C}{Q}_{jc_2t} + \Delta s_{jc_2}\Delta\frac{C}{Q}_{jc_2}\right)\\
&+\sum_{j\in\mathcal{A}}\left(\Delta Q_j\frac{C}{Q}_{jt} + \Delta Q_j\Delta\frac{C}{Q}_{j}\right)\\
&-\sum_{j\in\mathcal{B}} Q_{jt}\frac{C}{Q}_{jt}\\
&+\sum_{j\in\mathcal{C}} Q_{jt+1}\frac{C}{Q}_{jt+1}
\end{aligned}$$

I write $\frac{C}{Q}_{jtc_2}$ for the chemical-to-output ratio of farm j in period t for crop c_2, s_{jtc_1} denotes the output share of crop c_1 in farm j at t, and Q_{jt} is for total output. The first line of the decomposition is for within-farm within-crop changes in chemical intensity associated with

changes in input price ratios. The second line is for changes in chemical-use intensity coming from reallocations of production within the farm. The last line is for cross-farm reallocations, as well as changes in production volumes.

Changes in the chemical expenditure to sales ratio can be obtained with a comparable decomposition, where farm revenue market share replaces production volume, and where the chemical-per-output ratio is replaced by the farm-crop specific of farm specific chemical expenditure ratio.

I use this accounting equivalence to decompose the changes in chemical use intensity, and in total chemical use generated by the series of increasing land subsidies and lump sum payments discussed in Section 5.1.

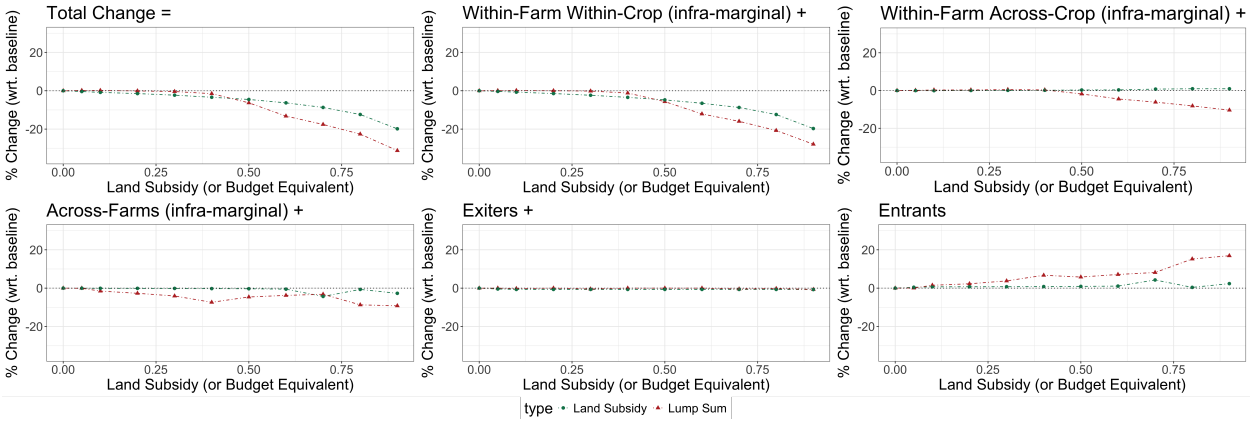


Figure A25: Decomposition of Chemical Use Intensity for Land Subsidy vs. Lump Sum Payments to Low Chemical Use Efficiency Farms

Notes: The first figure shows the evolution of the chemical expenditure to sales ratio, which is then decomposed in the remaining figures, using the accounting formula described in Appendix E. The decomposition of the change in total chemical use is shown in Figure A26.

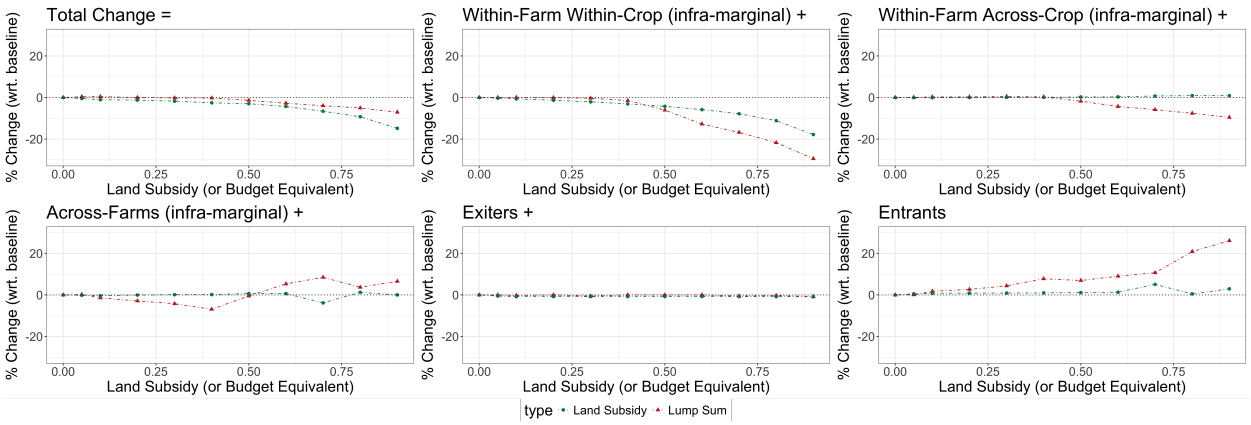


Figure A26: Decomposition Total Chemical Use for Land Subsidy vs. Lump Sum Payments to Low Chemical Use Efficiency Farms

Notes: These figures describe the evolution of total chemical use under a series of respectively increasing land subsidies, and budget-equivalent lump sum payments paid to low chemical use efficiency producers who cannot survive under the baseline no-intervention equilibrium. The first figure shows the evolution of total chemical expenditures, which is then decomposed using the accounting formula described in Appendix E.

F Sensitivity Analysis for the Production Function Estimation

I rely on the work of Andrews et al. (2017) to discuss the role of the identifying assumptions regarding instrument exogeneity in driving the estimation of the production function parameters. The purpose of the method developed by Andrews et al. (2017) is to recover the sensitivity of model parameters to the moments used in structural estimation, and in this extending the omitted variable bias formula to structural cases. They propose to recover a matrix they call the "sensitivity" matrix, which maps the relation between the parameters of the model, and the moments used for their estimation. When combined with alternative assumptions about instrument validity—recasted as the impact of alternative assumptions on moment values—this sensitivity matrix allows to predict the impact of alternative assumptions on the estimated parameters. In other words, this analysis aims at shedding light on the role of identifying assumptions in driving estimation.

I denote by $\hat{\theta}$ the parameter vector minimizing the criterion function:

$$\hat{g}(\theta)'\hat{W}\hat{g}(\theta)$$

In this case, $\hat{\theta}$ can be the estimate for the first stage and second stage of the production function estimation. For any a in the set A of alternative assumptions, they define a local perturbation of the model in the direction of a such that the estimate $\hat{\theta}$ has first-order

asymptotic bias—where Λ is the sensitivity matrix:

$$\mathbb{E}[\tilde{\theta}(a)] = \Lambda \mathbb{E}[\hat{g}(a)]$$

An alternative assumption a should be interpreted as an alternative assumption regarding the relation between the chosen instruments and the structural errors for resp. the first and second step of estimation. The sensitivity matrix can be written as:

$$\Lambda = -\left(G'WG\right)^{-1} G'W$$

Where W is the probability limit of $\hat{W}$, the weight matrix used in the GMM criterion, and G is the Jacobian of the probability limit of $\hat{g}(\theta)$ evaluated at the true parameter vector θ_0. Λ serves as a local approximation to the mapping from moments to estimated parameters.

I provide an estimate of ΛK for each of the estimation step for the production function, where K is a weighting matrix based on the standard deviation of the relevant moments. As Andrews et al. (2017) discuss, the units of Λ are contingent on the units of the different moments (and implicitly on the range of variation of the instruments). The weighting ensures that ΛK elements can be read as the effect of a one standard deviation violation of the given moment condition on the asymptotic bias in $\hat{\theta}$.

Table A17: Sensitivity for First Step

	constant	Lag Chemicals-Land Ratio	Capital	Lag Capital	Lag Chemical-Land Price Ratio	Land Subsidy	Farm Exposure
ρ	-5.6975	0.2596	-0.0330	-0.0084	-0.0233	0.2271	0.0886
δ_s (group 1)	-11.5509	0.4113	-0.0767	-0.0165	-0.1357	0.5239	0.2000
δ_s (group 2)	-3.3408	0.2259	-0.0109	-0.0049	0.0516	0.0783	0.0338

Table A18: Sensitivity for Second Step

	Constant	Capital	Lag Capital	Lag Labor	Lag Wage	Lag Land Price
α_l (group 1)	-0.0515	-0.0002	0.0002	0.0001	0.0000	0.0000
α_s (group 1)	0.0435	0.0001	-0.0002	-0.0000	-0.0000	-0.0000
α_l (group 2)	0.0223	0.0001	-0.0001	0.0000	-0.0000	-0.0000
α_s (group 2)	-0.0202	-0.0001	0.0001	0.0000	0.0000	0.0000
	Lag Chemicals-Land Ratio	Lag Land Composite	Lag Chemicals-Land Price Ratio	Lag Output Price Index	Land Subsidy	Farm Exposure
α_l (group 1)	0.0001	0.0000	-0.0005	-0.0001	0.0000	-0.0002
α_s (group 1)	-0.0002	0.0001	0.0003	0.0001	-0.0000	0.0002
α_l (group 2)	-0.0000	-0.0002	0.0002	0.0000	0.0000	0.0001
α_s (group 2)	0.0001	0.0001	-0.0001	-0.0000	0.0000	-0.0001

G A Simplified MacSharry: Budget Equivalent Government Purchasing vs. Land Subsidies

In this counterfactual exercise, I model the conceptual change in subsidization at the heart of the MacSharry reform: moving from demand support via direct government purchasing of commodities, to a land subsidy which reduces the farmers' costs of production. To capture the consequences of the change in subsidy design, I model a government expenditure-constant change from one policy to the other. Using a 32% land subsidy value for 1995, this implies a government purchase of 0.9 million tons of grain on the market, where the expenditure corresponds to the total cost of the purchase, which I then split equally between the two commodities. Table A19 shows the consequences of the policy change.

Reform Evaluation: The first effect of the policy change is to decrease output prices. Because government purchasing is modeled as symmetric across the two markets, it ends up playing a larger role on the market with the highest crop price, which is here the second crop. The price decreases generated by the move to a land subsidy are concentrated on that second crop, while the price of the first crop remains at the same level. Because farms are multi-product producers, and because of decreasing returns to scale and the absence of crop-specific fixed costs, all farms produce some amount of the two crops. This implies that the new zero-profit condition can be reached with an asymmetric price decrease, where only one of the two crops is impacted.

Table A19: Government Purchasing vs. Land Subsidization (Budget Equivalent)

Panel A	Market Outcomes							
	P (Wheat)	P (Other)	Q (Wheat)	Q (Other)	Avg. Profit	Avg. Cost	Chemical/Q (Wheat)	Chemical/Q (Other)
Government Purchasing	164 €/t	246 €/t	41.9 Mt	22.3 Mt	52.2 k€	145.8 €	0.159	0.255
Land Subsidy	164 €/t	242 €/t	41.4 Mt	22.1 Mt	53.1 k€	147.5 €	0.157	0.245
Panel B	Welfare Implications							
	Consumer Surplus		Total Profit		Chemical Expenditure		Pollution Damage	
Δ Level (%)	+ 103 M€		-223 M€ (- 8.3%)		-137 M€ (-3.2%)		[-809 M€,-274 M€] (-3.3 %)	
Panel C	Decomposing the Change in Chemical Expenditure							
	Change (%) ≡		Crop-Level (Within) +		Crop Mix (Within) +		Reallocation (Across)	
	-3.3%		-2.29%		+0.073%		-1%	

Notes: I recover the equilibrium under a 32% land subsidy, and then perform a bisection to find the level of government purchasing for which government expenditure is equivalent. The equilibrium under government purchasing has a government expenditure of 191.4M€, while the equilibrium under the land subsidy has an expenditure of 191.6M€. This table compares these two equilibria.

Second, the reform leads to a decrease in production volume. Here the decrease is of 1% on both markets. This decrease in production is driven by market exit, and a decrease in the mass of producers on the market. The farms that stay on the market increase their volumes of production slightly, and consequently reach slightly larger profits. They do so

because they benefit from the decrease in costs of production which is not fully translated into decreases in output prices.

Overall, this change of subsidization has a negative effect on market-level production efficiency (measured by the average cost of production), but a positive one on environmental outcomes. Government purchasing props up output prices without altering relative factor prices, and this has a relatively small effect on overall production efficiency. On the contrary, the land subsidy distorts input choices away from the optimal mix, which generates efficiency losses (the average cost per unit goes up by 1%), but it also distorts input choices away from the polluting input which is now relatively more expensive. As a consequence the ratio of chemical-use-to-output is reduced for the two crops (resp. 1.2% for wheat, and 3.9% for the other crop). The total use of chemicals is decreased by 3.3%. I decompose this total change in chemical use, using a modification of the decomposition formula of De Loecker and Collard-Wexler (2015) (see Appendix E).[60] This decomposition reveals that decreases in total chemical use are in part driven by changes in farm-crop-specific input mixes holding output constant (-2.3%), and in part by decreases in total production and reallocations (-1%). Focusing on consumer surplus and total farm profit, this change in subsidization is only evaluated as positive if we account for the environmental gains from reduced chemical use. Losses in farm profit are about twice as large as gains in consumer surplus from reduced output prices. This calculation does not account for any additional policy objectives, such as supporting the revenue of small farmers or guaranteeing food security. If we account for the decreases in chemical use, evaluated using low and high marginal costs of agricultural pollution from the French CGDD (2011), these welfare gains become positive, and are raised by a value between 274 M€ and 809 M€.[61]

A change in the model introducing crop-specific fixed costs would give rise to more specialized producers in the model, and would likely accentuate reform-driven price decreases for both crops. The actual MacSharry reform also included an end to export subsidies and import levies, which should have further decreased output prices, and which I do not model. Finally, the fact that I only model the French market, rather than the whole EU market, might hide reallocations in production happening across countries, where larger reform-induced farm exit of low efficiency producers might also decrease output prices and costs of production. The French grain market has higher yields than other European markets, and one could expect more significant farm exit in other EU countries, in turn softening the impact on the French market.

[60] The set of farm types that participate in the market is equal across the two equilibria. As such the last two terms shown in Appendix E are equal to zero in this counterfactual exercise, and are bundled with the reallocation term in the table. However, while this set remains constant, there is still a decrease in the overall mass of active producers.

[61] This report from the French government estimates total damages from agricultural pollution to be ranging between 828 and 2430 euros per ha. I borrow an estimate for the chemical expenditure per ha from Guillermet (2015) of 415 euros/ha, and I translate these damages into damages per euro of chemical expense (resp. 2 or 5.9).