

Regenerative Agriculture and Farmer Profitability: Modeling Adoption of Scalable Climate Mitigation*

Ryan McWay

University of Minnesota

Jay Coggins

University of Minnesota

Nicholas Gallagher

University of Minnesota

Version: October 19, 2026

Latest Version: [Click Here](#)

Abstract

The climate-mitigation potential of agricultural conservation practices depends in part on whether farmers find them privately profitable. This paper studies the profit motivation for adoption of cover cropping and conservation tillage in Midwestern corn, soybean, and wheat production. I combine farm financial records for 2000–2025 with a machine-learned county-year panel of predicted soil organic carbon (SOC). Using a doubly robust staggered difference-in-differences design, I estimate practice- and crop-specific effects on profitability and soil health. The results reveal a heterogeneous profit premium: estimated net returns increase by \$79.40 per acre for corn cover crops and \$22.50 for soybean reduced tillage, while most other comparisons vary considerably. Under fixed mineral-soil assumptions, the reduced-tillage SOC difference after 15 years reaches 45.88 and 43.21 tonnes CO₂/ha for corn and soybeans. Scaling these point estimates to national harvested acreage yields illustrative stock equivalents of 1.69 and 1.50 billion tonnes CO₂, respectively. To understand long-run soil capital dynamics, I develop a dynamic structural model linking management, production costs, and insurance to soil-stock accumulation and future returns. Together, the empirical analysis and model distinguish observed private returns from agronomic potential and clarify the intertemporal incentives relevant to conservation policy.

Keywords: Regenerative Agriculture; Climate Mitigation; Farmer Profit.

JEL classification: Q12, Q15, Q53, Q57.

*Corresponding Author: Ryan McWay (mcway005@umn.edu). I appreciate feedback from Steve Polasky, Paul Glewwe, Matthew Braaksma, Terry Hurley, Joleen Hadrich, Wallace Pupadapolous, Feyre Archeron, and Rebecca Reilly as well as commentary from participants at NatCap TEEMs, the AAEA Summer Conference, Camp Resources, SumEC Mini-conference, and the WEAI Summer Conference. I received no funding for this research.

1 Introduction

Agriculture faces a dual challenge. It contributes materially to greenhouse gas emissions, while intensive production has also degraded soil health (Vermeulen et al., 2012; Crippa et al., 2021; Lal, 2015; Louwagie et al., 2009; Montgomery, 2007; Jie et al., 2002). This combination has motivated interest in management practices that could sustain agricultural production while improving soil conditions and reducing the environmental footprint of food production. Yet an agronomic opportunity is not by itself a scalable climate policy. Large-scale change requires farmers to adopt and maintain practices whose costs, returns, and risks are experienced privately, even when some of the resulting environmental benefits accrue to society as a public externality.

Policy has consequently pursued both demand-side and supply-side approaches. Demand-side initiatives seek to create markets for products or environmental attributes associated with improved production practices (i.e., certification programs). Supply-side policies, including technical assistance and cost-sharing, seek to reduce the costs and risks of changing farm management. The effectiveness of either approach depends on the private incentives it creates and on whether support addresses the specific barriers farmers face. In particular, conservation practices may require near-term investments or changes in production before soil benefits accumulate. The relevant policy question is therefore not only whether a practice can improve soil, but whether its expected stream of farm returns is sufficient to induce adoption and persistence.

This paper studies that question in U.S. Midwestern row-crop agriculture, focusing on corn, soybean, and wheat production. It examines a suite of regenerative agricultural practices, including cover cropping and a gradient of conservation tillage such as reduced tillage and no-till. Cover crops add living vegetation between cash-crop periods or alongside production, while conservation tillage practices aim to minimize soil disturbance. These practices can alter input requirements and current yields, as well as the rate at which soil stocks accumulate or deplete. Their private returns may therefore differ across crops, farms, and time since adoption.

I ask two principal questions to address agriculture’s dual challenge. First, how does adopting regenerative practices change farm profitability and its components, including revenues, yields, expenses, and crop-insurance outcomes? Second, are adoption and management changes associated with detectable improvements in soil organic carbon (SOC) at scale, an important dimension of soil health? Farm financial records provide detailed economic outcomes at the farmer-year-crop level, whereas the soil outcome available for this study is a machine learning predicted county-year measure. Moreover, soil

capital accumulates gradually, so a short-run response in yields or expenses need not reveal its longer-run contribution to profits. The analysis therefore requires a long-run assessment of the impacts of regenerative practices on both economic and environmental outcomes.

I address these questions by combining two sources of data. The first is the FinBin data set provided by the Center for Farm Financial Management (CFFM) at the University of Minnesota, a farm financial-management database based on producers' farm records and tax-preparation applications. I use its annual records for corn, soybean, and wheat producers in the Upper Midwest over 2000–2025 to construct a panel of management choices and farm-level economic outcomes. The second is a county-year panel of predicted SOC developed from remotely sensed and other spatial features using machine-learning methods. The prediction model is trained on Wisconsin soil-test data and applied to construct county-level estimates over the study period. The two sources provide complementary evidence: FinBin measures private economic outcomes, while the predicted SOC panel supplies a spatially and temporally aligned environmental indicator of soil health and carbon stocks.

Estimating the effect of adoption is challenging because farmers choose their management practices. Producers who adopt may differ from nonadopters in management ability, expectations, resources, soil conditions, or conservation preferences. These same differences can influence profits and soil outcomes. Consequently, a simple comparison of adopters and nonadopters can confound practice effects with selection bias. I address this concern using the Callaway and Sant'Anna (2021) staggered difference-in-differences doubly robust estimator that compare outcome changes for farmers adopting at different times with changes for not-yet-treated or never-treated farmers, under the identifying assumption that, absent adoption, their outcomes would have followed comparable trends conditional on the design. This approach makes treatment timing explicit and permits effects to vary across adoption cohorts and time since adoption.

The reduced-form results show substantial heterogeneity rather than a uniform profit premium. Some crop-practice comparisons have positive estimated net returns, but the pattern differs across outcomes and is not consistently accompanied by higher yields: revenues and expenses can both change, while insurance outcomes are generally less responsive. Several comparisons also show evidence of differential pre-adoption trends, which limits the strength of causal conclusions for those estimates. The SOC analysis likewise finds positive estimates for selected crop-practice comparisons, including reduced tillage for corn and several tillage practices for soybean. These estimates are not uniform across practices or crops, and pre-trend concerns and the county-level predicted outcome

require cautious interpretation.

To illustrate the carbon-sequestration implications of the tillage estimates, I convert predicted concentration differences into stock equivalents assuming a 30 cm mineral-soil layer, bulk density of 1.3 g/cm^3 , and no coarse fragments. Over a 15 year time horizon, reduced tillage implies 45.88 tonnes CO_2/ha for corn and 43.21 for soybeans. Applying these point estimates to national harvested acreage yields 1,688.61 and 1,501.63 million tonnes CO_2 , respectively. Averaged over the 15-year horizon, the per-acre stock gains imply annual removals of 1.24 and 1.17 tonnes CO_2 . Subtracting them from crop-production emissions implied by Table 1 gives net emissions of 0.24 tonnes $\text{CO}_2\text{e}/\text{acre}/\text{year}$ for corn and -0.63 for soybeans. These calculations credit only the SOC change and assume no other practice-induced change in production emissions. At the EPA social-cost benchmark of \$46 per tonne CO_2 for 2025, in 2007 dollars, the implied one-time stock values are \$854.07 per corn acre and \$804.31 per soybean acre. Conservation-tillage and no-till scenarios also yield positive stock point estimates, but all tillage endpoint 95% confidence intervals include zero. The exercise therefore illustrates the potential scale and social value of soil-carbon changes conditional on the conversion and national transfer assumptions.

The gradual adjustment of soil capital motivates a dynamic structural model of farm decision-making. I develop a novel dynamic structural model to describe this relationship. In the model, management affects current production and costs changing with the evolution of soil stocks. Soil stocks, in turn, affect future production and the value of insurance. This intertemporal framework clarifies why a practice with short-run transition costs may nevertheless be privately attractive over a longer horizon, and why the adoption question depends on the persistence of soil improvements, production risk, and farm-specific conditions. I use the model to organize the estimated profitability and soil relationships and to study how soil capitalization shapes adoption incentives. The model complements the reduced-form analysis: the latter estimates outcome changes around adoption, while the former represents the dynamic channels through which those changes can affect choices over time.

This paper makes three contributions. First, it brings farm-level evidence on private returns to the agricultural climate-mitigation debate. Research documents agriculture's emissions and the potential role of soil management in mitigation (Crippa et al., 2021; Smith and Swanson, 2025), but the climate potential of a practice does not establish that farmers have an incentive to adopt it. I combine farmer-year-crop financial records with a county-year panel of predicted SOC and use the group-time doubly robust estimator of Callaway and Sant'Anna (2021) to estimate effects separately by practice and crop under staggered adoption. This design lets me examine a broader set of private outcomes—net

returns, revenues, expenses, yields, and crop-insurance outcomes—alongside an environmental proxy. The contribution is to measure this bundle of outcomes across crops and practices, rather than infer profitability from agronomic potential or from a single production outcome; the county-level SOC estimates remain distinct from direct measurements of field-level carbon stocks.

Second, I connect the economics of soil capital to observed farm decisions and outcomes. Existing dynamic studies provide important proof of concept: Berazneva et al. (2019) link maize production to soil-carbon dynamics in western Kenya and show how time preferences and market frictions can separate observed soil stocks from the model's optimal stocks; Brown et al. (2026) use long-run cotton-experiment data to model joint cover-crop and nitrogen choices, finding that the incentive to adopt depends on the soil stock already accumulated; and Malik (2026) studies no-till and reduced-tillage choices when payments depend on uncertain, heterogeneous SOC gains. Building on this literature, I formulate a dynamic model in which management choices affect current production, input and switching costs, and insurance, while also governing soil-stock transitions that shape future yields and returns. The paper's specific contribution is to place the soil-capital mechanism alongside farmer-level profitability and insurance outcomes for Midwestern row crops, and to use predicted county-year SOC as an additional, explicitly limited measure of soil change.

Third, I contribute to the study of conservation adoption by documenting how private returns vary across crop-practice comparisons and by treating transition costs and production risk as central to persistence. This focus complements evidence that practice use can be heterogeneous, partial, or reversed over time (Newton et al., 2020; Bowman and Pratt, 2025), and the policy analyses of Wu (2000), who shows that behavioral responses can offset intended conservation gains through slippage, and Antle et al. (2003), who demonstrates that spatial heterogeneity and contract design matter for the efficiency of agricultural carbon-sequestration policies. My analysis does not estimate policy additivity or social cost-effectiveness. Instead, it supplies evidence on the farm-level profitability and risk margins that policy design must account for if conservation adoption is to continue beyond initial participation.

2 Regenerative Agriculture

2.1 Principles of Regenerative Agriculture

Regenerative agriculture has no universally accepted definition and can refer to a range of practices intended to improve soil health, biodiversity, and ecosystem services. Commonly cited principles include minimizing soil disturbance, maintaining soil cover, diversifying crops, and integrating livestock. I focus on row-crop production, particularly corn, soybean, and wheat, and considers conservation tillage and cover cropping as management practices associated with these principles.

Conventional field-crop systems in the United States are typically organized around high-input management designed to maximize yield in each production season. In this system, farmers rely on intensive applications of synthetic fertilizer, herbicides, pesticides, tillage passes, and fuel-intensive field operations to sustain high output per acre. The production logic is clear: higher yields increase gross revenue and can help preserve cash flow in the presence of volatile commodity prices and weather risk.

High input use may also reflect the condition of the soil capital. Where soil fertility and organic matter are depleted, producers may rely more heavily on external inputs to sustain production. This can create a reinforcing (vicious) cycle in which low soil quality increases dependence on inputs, while some intensive management practices can further impair soil structure and biological function.

These modern production systems can also generate environmental costs that are not fully reflected in private production decisions. Intensive soil disturbance and input use are associated with soil carbon depletion and greenhouse-gas emissions, including carbon dioxide (CO_2), methane (CH_4), and nitrous oxide (N_2O) (Vermeulen et al., 2012; Crippa et al., 2021). When these social costs are not internalized, private decisions may favor more emissions-intensive management than is socially optimal.

Regenerative agricultural practices are increasingly framed as a climate-smart alternative to this high-input pathway. The proposed mechanisms include reducing soil disturbance, rebuilding soil structure, and improving biological nutrient cycling. In principle, these changes can support carbon sequestration and long-run productivity by improving nutrient availability and retention, including for phosphorus and nitrogen, while building soil organic carbon (Paustian et al., 2016; Lal, 2015; Newton et al., 2020). Regenerative management is therefore relevant both as an environmental intervention and as an economic production decision.

Conservation tillage minimizes or eliminates mechanical soil disturbance before planting. Compared with conventional systems that use multiple tillage passes to prepare

seedbeds and suppress weeds, conservation-tillage systems retain higher crop residue on the soil surface and rely on other weed-control methods. The shift is from managing fields primarily through mechanical disturbance toward management that retains surface cover and supports soil biological processes.

The primary agronomic benefits are improved soil structure and water management (Nunes et al., 2018; Ogle et al., 2019). Less disturbance can reduce erosion, increase aggregate stability, and improve infiltration and moisture retention. Conservation tillage may also support gradual recovery of soil organic carbon by reducing losses associated with repeated disturbance. These soil-health gains may be especially valuable during periods of stress, when moisture retention and rooting conditions constrain yields.

Production effects vary across time and place. During transition, producers may face yield penalties associated with learning costs, changes in pest pressure, compaction concerns, or a mismatch with local soil and climate conditions (Nunes et al., 2018; Newton et al., 2020). Over longer horizons, these losses may narrow or reverse as soil conditions improve and management is adapted. Cost effects also vary: fewer tillage passes can reduce fuel use, labor, and machinery wear, whereas transition may require equipment changes, more intensive management, or higher herbicide expenditures in some systems.

Conservation tillage has expanded in many row-crop regions, but adoption remains heterogeneous across soil types, climate risks, equipment access, and farmer experience (Newton et al., 2020; Bowman and Pratt, 2025). Partial adoption, intermittent use, and practice reversal are also common (Bowman and Pratt, 2025), suggesting that adoption responds to both short-run profitability and risk-management constraints.

Cover cropping introduces non-cash crops between primary cash-crop seasons. Compared with conventional fallow periods, cover crops maintain living roots in the soil for a larger share of the year and provide additional surface protection. Common row-crop applications include planting winter cover crops after corn or soybean harvest and terminating them before spring planting (Wood and Bowman, 2021; Newton et al., 2020).

The soil benefits are broad: reduced erosion, greater surface protection, improved nutrient retention, and potentially higher soil organic matter through additional biomass inputs (Wood and Bowman, 2021). Cover crops can also reduce nutrient losses by scavenging residual nitrogen and improving nutrient cycling (Pandey et al., 2025; Wood and Bowman, 2021). Over time, these mechanisms may strengthen soil resilience and reduce the sensitivity of production to rainfall and temperature shocks (Wood and Bowman, 2021; Pandey et al., 2025).

As with conservation tillage, production and cost effects involve tradeoffs. Potential benefits include improved long-run soil fertility, better moisture management, and re-

duced need for some inputs (Wood and Bowman, 2021; LaCanne and Lundgren, 2018). Costs may include seed, termination, and additional field operations, as well as short-run yield penalties when timing, species choice, or termination management is suboptimal (Newton et al., 2020; Pandey et al., 2025). Net profitability therefore depends on local agronomic conditions, management quality, and the evaluation horizon.

Adoption of cover cropping has risen in recent years but remains limited in many row-crop systems relative to its technical potential. This gap may reflect transition risk, short-run cash-flow constraints, and uncertainty about private returns under heterogeneous local conditions.

2.2 Soil Health as Productive Capital

Soil health describes the capacity of soil to sustain biological productivity and perform important ecosystem functions. Soil organic carbon (SOC) is a central indicator in this context because it reflects organic matter in the soil and is linked to soil structure, nutrient cycling, and water management. Soil health is nevertheless broader than SOC alone: properties such as structure and nutrient retention also contribute to the soil's productive capacity. An economic representation can therefore treat soil health as a capital stock that enters the production process alongside annual inputs. Management may build or deplete this stock, with consequences for subsequent productivity and input requirements. Unlike purchased inputs whose effects are concentrated in the current season, changes in soil capital accumulate and depreciate gradually. Producers thus face an intertemporal decision: whether the current cost of a soil-building practice is justified by uncertain future returns.

Bioeconomic models make this dynamic relationship explicit by linking agricultural production to a changing soil stock and optimizing management over time. For example, Berazneva et al. (2019) combine an estimated maize production function with soil-carbon dynamics and calibrate the model using agronomic and household data from western Kenya. Their optimal soil-carbon stocks and maize yields exceed observed levels in the study region; the authors associate this gap in part with heterogeneous farmer time preferences, information barriers, and market imperfections. The study also estimates a positive shadow value for soil carbon, illustrating how soil condition can represent both a productive asset and an opportunity cost when it is mismanaged. Its numerical results are specific to the Kenyan production system, but its framework demonstrates why the returns to management should be evaluated across time rather than within a single season.

Brown et al. (2026) extend this soil-capital perspective to the joint choice of cover crops and nitrogen fertilizer. Using a 35-year cotton experiment in West Tennessee, they estimate a yield relationship with accumulated soil capital and fertilizer, then embed the estimates in a dynamic model under biophysical and market uncertainty. In their model, higher soil capital reduces the marginal productivity of fertilizer, while cover-crop adoption becomes optimal only above a threshold level of soil capital. These findings illustrate that adoption incentives can depend on the condition of the soil stock already in place, not only on the practice's average costs and benefits. Because their analysis uses a specific crop, location, and experimental setting, its estimated relationships should not be assumed to transfer directly to all row-crop systems. Together, these studies motivate treating soil condition as a persistent state variable when examining practice adoption and farm profitability.

2.3 Potential of Climate Mitigation via Soil Capture

Agriculture contributes to climate change through several distinct pathways. Production and land-use change emit carbon dioxide (CO_2), fertilizer use can generate nitrous oxide (N_2O), and livestock contribute methane (CH_4). At the same time, soils can remove CO_2 from the atmosphere and store some of that carbon as soil organic carbon (SOC). This potential has motivated interest in soil-building practices as a climate-mitigation strategy, alongside their possible effects on productivity and resilience (Crippa et al., 2021; Smith and Swanson, 2025).

Table 1 places crop and livestock production in a common descriptive frame by reporting estimated land use, land intensity, and emissions per unit of output. These comparisons help illustrate the variation in emissions and land requirements across commodities, but do not estimate how much carbon any specific regenerative practice will sequester. Soil-carbon mitigation depends on changes in SOC relative to a baseline, not simply on the emissions intensity of the commodity produced.

The mitigation case for increasing SOC is plausible but conditional. Practices that increase carbon inputs to soil or reduce soil disturbance can raise SOC under suitable conditions; however, the magnitude and persistence of gains vary with soil, climate, management, and initial carbon stocks. Accumulation also slows as soils approach a new equilibrium, and stored carbon can be lost if management changes. As a result, soil carbon is a finite, reversible sink, not a permanent substitute for reducing fossil-fuel emissions (Ogle et al., 2019; Smith and Swanson, 2025). Credible estimates of climate benefits therefore require a baseline and attention to additionality, measurement, uncertainty, and perma-

nence. The agronomic potential of soil-carbon practices is consequently best understood as a geographically heterogeneous contribution to mitigation, rather than as a uniform benefit of adopting a practice.

Table 1: Commodity Land Intensity and Carbon Emissions

Commodity	Acres Used (Mill.)	Land Intensity (Acres/lb)	MMT CO ₂ e	Avg. lb CO ₂ e/lb
<i>Panel A: Primary Crops</i>				
Corn	91.07	0.0001	134.52	0.461
Cotton	10.90	0.0013	3.64	0.941
Rice	2.75	0.0001	17.33	1.794
Soybeans	75.82	0.0004	40.62	0.466
Wheat	47.27	0.0004	36.04	0.686
<i>Panel B: Secondary Crops</i>				
Barley	3.59	0.0003	4.32	0.909
Hay & Haylage	58.46	0.0002	19.90	0.155
Sorghum	0.90	0.0004	1.22	1.004
Sugarbeet	1.27	0.0000	2.13	0.073
Sugarcane	0.91	0.0000	1.84	0.059
<i>Panel C: Primary Livestock</i>				
Beef Cattle	268.05	0.0169	313.07	43.519
Cattle on Feed	12.48	0.0016	149.19	43.519
Dairy Cows	22.06	0.0048	82.03	39.022
<i>Panel D: Secondary Livestock</i>				
Bison	1.05	0.0036	7.26	55.115
Horses & Ponies	173.46	0.0540	4.89	3.307
Mules, Burros, & Donkeys	29.07	0.1374	0.31	3.307
Goats	36.20	0.1094	5.82	39.022
Mohair & Angora	0.97	0.0321	0.78	39.022
Sheep & Lambs	26.62	0.0308	4.91	12.000
<i>Panel E: Barn Animals</i>				
Chicken	74.26	0.0070	36.76	7.683
Duck	8.59	0.2486	0.18	11.363
Turkey	6.06	0.0020	31.10	22.590
Hogs	23.79	0.0024	102.88	22.590
<i>Panel F: Vegetables, Nuts, and Fruits</i>				
Vegetables	2.48	0.0000	0.85	0.023
Tree Nuts	2.07	0.0004	8.46	2.667
Citrus Fruit	0.88	0.0000	0.92	0.122
Strawberries	0.06	0.0000	0.11	0.084
Blueberries	0.13	0.0000	0.13	0.568

Economic analyses make these biophysical limits central to program design. Malik (2026) models farmers' long-run choices of no-till and reduced tillage under payments linked to SOC gains, allowing sequestration outcomes and farmer returns to vary across soil types and with risk. In the model, the annual payment required to induce adoption ranges from \$8 to \$32 per ton of carbon across the soil types studied. Greater uncertainty in sequestration outcomes also raises the compensation required. These are model-based thresholds rather than universal market prices, and illustrate why a uniform payment per acre or per unit of carbon may not provide the same incentive across farms. Program evaluation must account for both the expected amount and the risk of additional carbon storage.

2.4 Adoption of Regenerative Agriculture

National adoption statistics show both growth and room for further participation. USDA Census of Agriculture data report that cover-crop acreage increased from 10.3 million acres in 2012 to 15.4 million in 2017 and 18.0 million in 2022. The share of cropland under no-till also rose, from 24.8 percent in 2012 to 26.3 percent in 2017 and 27.5 percent in 2022 (USDA National Agricultural Statistics Service, 2022). These national totals describe the extensive margin of practice use, but do not show whether farms maintain a practice consistently or how intensively it is implemented. Indeed, an ERS analysis tracking the same operations across the 2012, 2017, and 2022 censuses finds that only 3.4 percent reported cover-crop use in all three years, even though 25 percent used cover crops in at least one of them (Bowman and Pratt, 2025). Adoption is therefore not simply a one-time transition; persistence and intermittent use are also important margins.

The costs and benefits described above help explain why adoption varies across producers and over time. Seed, equipment, labor, and additional management can raise near-term costs, while soil benefits may accumulate gradually and depend on local conditions. Producers must also manage uncertainty about yield responses, weather, termination timing, and the future returns to improved soil. These considerations can make adoption less attractive when capital or liquidity is constrained, or when the farm's planning horizon is short. Evidence of partial, intermittent, and discontinued use is consistent with adoption decisions that respond to profitability and risk rather than to agronomic potential alone (Newton et al., 2020; Bowman and Pratt, 2025; Brown et al., 2026).

Public conservation programs and private incentives can help address some of these constraints, although they operate through different mechanisms. The USDA's Environmental Quality Incentives Program (EQIP) provides technical assistance and may provide

financial assistance for eligible conservation practices. The Conservation Stewardship Program (CSP) uses multiyear contracts to compensate producers for maintaining and improving conservation on working lands (USDA Natural Resources Conservation Service, 2026c,a). State programs and private supply-chain or carbon-market initiatives provide additional possible channels for reducing transition costs or rewarding environmental outcomes. Carbon payments can link compensation more directly to climate benefits, but require credible measurement and rules for additionality and permanence; practice-based cost-share programs instead subsidize eligible management without necessarily paying in proportion to realized carbon gains. These programs are therefore useful policy counterfactuals for studying whether support changes adoption and whether any resulting profitability effects persist after incentives end.

3 FinBin Data: A Panel of Farmer Finances

3.1 The Data Generating Process

The analysis uses crop- and farm-level records compiled from farm business management programs and made available through FINBIN, a farm financial benchmarking resource maintained by the University of Minnesota’s Center for Farm Financial Management (CFFM). Participating producers work with farm business management educators and advisers using FinPack and related record-keeping and analysis tools to organize whole-farm financial and enterprise information. These records support farm management, financial analysis, and credit assessment; their detailed accounting structure makes it possible to study crop-level revenues, expenses, and management in a way that is uncommon in general-purpose agricultural surveys (University of Minnesota Center for Farm Financial Management, 2026a,b). Because the records originate in a voluntary farm business management setting, however, the participating operations should not be assumed to represent all farms, and the data were not collected as a representative random sample.

Access to the research data was approved through five data collaborators: the Minnesota State Farm Business Management Program; the University of Minnesota Extension Southwest Minnesota Farm Business Management Association; Nebraska Farm Business, Inc.; the Ohio Farm Business Analysis and Benchmarking Program; and the Utah Center for Farm and Ranch Management. The data contains annual records from 2000 through 2025 for corn, soybean, and wheat production. The number of crop observations per year are an unbalanced panel ranging from 9,601 in 2025 to 11,569 in 2007.

The analytic sample contains 275,085 farm $\times$ crop $\times$ year observations from 9,506

unique FinBin farm identifiers. Of these crop records, 142,746 are corn, 109,555 soybean, and 18,124 wheat observations. These farmers are sourced from 12 states with records concentrated in Minnesota (241,409 obs.) and Nebraska (27,210 obs.). The remaining ten states together account for a small share of the analytic sample. Accordingly, the data provide a long farm-level panel with substantial depth in these two states rather than uniform national coverage.

The unit of analysis is a farm $\times$ crop $\times$ year as opposed to an ideal field-level analysis. A farm may therefore contribute multiple crop observations in a year. Many conservation decisions are made at the field level, but the available records do not identify which fields within a crop enterprise receive a practice. The crop-level indicators consequently represent reported management for that farm's crop enterprise and cannot capture within-crop differences across fields. Interpreting them as field-level treatment would require assuming uniform practice use across the relevant fields, an approximation that may introduce exposure misclassification and attenuate estimated effects. The panel's state and year coverage, and its unbalanced structure, are accounted for in the descriptive and empirical analyses that follow.

3.2 Farmer Adoption of Regenerative Agriculture

FinBin records tillage system and cover-crop use at the crop level for each farm-year. The tillage variable distinguishes moldboard, reduced till, minimum till, no-till, ridge till, and strip till; records coded as "no answer" are grouped with conventional, moldboard-equivalent tillage in the constructed measure.¹ Cover-crop use is coded as one if the crop record indicates either that the crop was grown with a cover crop (inter-cropping) or that it followed a cover crop (over-wintering). Because all records reporting overwintering also report intercropping, it is difficult empirically to distinguish the difference of these two practices separately. As such, I combine them into a single cover cropping indicator.

I define adoption through the use of binary practice indicators. The binary indicators identify use of a reported practice in a given crop-year and support a direct comparison of adopters and non-adopters.

¹This classification was confirmed through consultation with CFFM data managers.

Table 2: Farmers Practice Adoption and Adoption Paths (N, % of Farmers)

Adoption Paths	Corn	Soybean	Wheat
# of Farmers (N)	8,967	7,960	2,199
Conventional Till	3,879 (43.3%)	3,267 (41.0%)	1,143 (52.0%)
Cover Crop	152 (1.7%)	150 (1.9%)	10 (0.5%)
Compliers	95 (1.1%)	91 (1.1%)	7 (0.3%)
Switchers	57 (0.6%)	59 (0.7%)	3 (0.1%)
Conservation Till	827 (9.2%)	1,072 (13.5%)	95 (4.3%)
Compliers	455 (5.1%)	537 (6.7%)	48 (2.2%)
Switchers	372 (4.1%)	535 (6.7%)	47 (2.1%)
No-Till	366 (4.1%)	700 (8.8%)	66 (3.0%)
Compliers	128 (1.4%)	271 (3.4%)	27 (1.2%)
Switchers	238 (2.7%)	429 (5.4%)	39 (1.8%)
Reduced Till	4,719 (52.6%)	4,220 (53.0%)	1,011 (46.0%)
Compliers	2,753 (30.7%)	2,380 (29.9%)	567 (25.8%)
Switchers	1,966 (21.9%)	1,840 (23.1%)	444 (20.2%)

The table summarizes observed practice use among the unique FinBin farms represented in each crop sample: 8,967 corn farms, 7,960 soybean farms, and 2,199 wheat farms. Cover-crop use is uncommon in these samples (1.7 percent of corn farms, 1.9 percent of soybean farms, and 0.5 percent of wheat farms), while reduced-tillage reports are far more common (approximately 46–53 percent). No-till is reported by 4.1 percent of corn, 8.8 percent of soybean, and 3.0 percent of wheat farms. These figures summarize whether a farm ever reported a practice during its observed panel, not its share of crop acres under that practice or the probability of adoption in a typical year.

The adoption-path rows distinguish farms that continue reporting a practice after its first observed use from those with a less persistent pattern. In the table, Compliers is a descriptive label for farms that keep reporting the practice in all subsequent observed years. Missing practice observations are treated as non-use when classifying paths, so persistence may be understated. Switchers includes the remaining farms that ever report the practice but do not continue reporting it in all subsequent observed years. Persistence varies by practice: among cover-crop adopters, 62.5 percent of corn farms, 60.7 percent of soybean farms, and 70.0 percent of wheat farms continue reporting use after the first observed adoption. The crop-specific farm counts are unbalanced, and farms enter and leave the panel at different times, so these observed paths should not be interpreted as full histories of management.

Practice rows are not mutually exclusive. Cover crops can be used alongside any

tillage system, and the conservation-tillage indicator includes no-till, minimum till, and strip till practices. Consequently, the row percentages should not be added to obtain a total adoption rate. At a finer spatial scale, I also summarize practice use by county, crop, and year. Because FinBin identifies farm locations only at the county level, the comparison with county-level soil-carbon measures cannot match practices to individual fields. Aggregating the crop-level management reports to county-year shares therefore introduces measurement error and ecological aggregation. The county-level summaries are used to describe spatial and temporal coverage and are assessed separately from the farm-level treatment measures.

Finally, crop-specific logit models reported in Appendix section C examine how observed farm, crop, and operator characteristics are associated with practice reports. These estimates are reported as descriptive, conditional associations rather than causal determinants of adoption: adoption is selected by producers, practice indicators are sparse for some crop-practice combinations, and the models do not remove all unobserved differences between farms or fields. They therefore provide context for sample composition and adoption patterns, not a substitute for the causal design developed below.

3.3 Profit Ledgers

To better understand the financial implications of different management practices, we examine crop-level profit ledgers for corn, soybeans, and wheat. These ledgers provide a detailed breakdown of revenue, expenses, and profitability across management types. The three profit ledgers report crop-level means, with standard deviations in parentheses, for revenue, expense, and profit components by crop and management indicator.

Gross return combines main-product value with miscellaneous crop income, including crop-insurance income; total expenses combine direct and overhead expenses. Net return is gross return less total expenses, and net profit margin is net return divided by gross return; margins below -100% are set to -100% in the analysis data. Monetary figures are pooled across years and are not adjusted for inflation, so the dollar means should be read as nominal amounts rather than as real purchasing-power comparisons.

The ledgers show substantial descriptive variation across management groups. For corn, mean net returns range from \$77.67 per acre for conventional tillage to \$130.73 for the cover-crop group; the corresponding mean net profit margins are 4.64% and 9.81%. The higher gross returns in the cover-crop group (\$1,000.22 per acre, compared with \$590.95 for conventional tillage) are accompanied by higher total expenses (\$869.49 versus \$513.28), leaving a more modest difference in net return. For soybeans, reduced-

till observations have the highest mean net return (\$100.94 per acre), while the cover-crop group averages \$84.61 and has a 9.69% mean margin, below conventional tillage's 11.62%. Wheat has lower and more variable mean profitability: conventional tillage averages \$33.77 per acre, reduced tillage \$42.15, and cover crops $-\$38.40$; mean margins are negative for the no-till, cover-crop, and conservation-tillage groups.

Table 3: Profit Ledger: Corn (Mean, SD in parentheses)

Statistic	Conventional Till	No-Till	Cover Crop	Reduced-Till	Conservation Till
Revenues					
Yield (Bu/ Acre)	249.20 (214.07)	201.24 (160.18)	301.27 (256.14)	239.32 (187.32)	215.32 (158.95)
Price (\$/Bu)	3.24 (1.77)	3.97 (1.54)	4.21 (1.72)	3.54 (1.54)	3.96 (1.52)
Crop Revenue (\$/ Acre)	560.01 (393.87)	659.84 (315.89)	924.06 (288.32)	680.79 (283.02)	723.93 (320.85)
Crop Insurance Income (\$/ Acre)	21.10 (74.31)	25.02 (71.34)	58.12 (170.04)	22.56 (67.51)	27.03 (76.66)
Gross Return (\$/ Acre)	590.95 (404.47)	697.84 (316.50)	1000.22 (293.76)	714.55 (285.17)	764.77 (324.59)
Expenses					
Direct Expenses (\$/ Acre)	395.71 (240.91)	469.15 (211.84)	699.18 (188.06)	479.97 (200.18)	510.75 (223.28)
Seed & Plants (\$/ Acre)	78.01 (37.47)	84.06 (36.79)	122.15 (30.76)	91.37 (35.89)	93.36 (37.11)
Fertilizer (\$/ Acre)	103.20 (69.44)	126.68 (73.68)	176.25 (71.46)	119.16 (68.04)	131.07 (72.45)
Crop Chemicals (\$/ Acre)	35.42 (23.08)	37.13 (22.36)	54.41 (24.83)	32.92 (18.04)	37.14 (21.19)
Hired Labor (\$/ Acre)	4.19 (13.75)	5.81 (14.48)	19.00 (19.04)	4.36 (12.76)	7.21 (15.57)
Overhead Expenses (\$/ Acre)	117.58 (93.81)	117.12 (77.68)	170.32 (108.38)	132.10 (90.38)	132.81 (91.39)
Crop Insurance (\$/ Acre)	19.26 (16.12)	19.67 (13.66)	29.40 (14.16)	19.69 (13.12)	21.22 (14.04)
Total Expenses (\$/ Acre)	513.28 (290.41)	586.27 (239.02)	869.49 (190.52)	612.07 (231.13)	643.56 (255.77)
Profit					
Net Return (\$/ Acre)	77.67 (275.33)	111.57 (230.18)	130.73 (297.28)	102.49 (210.90)	121.21 (226.40)
Net Profit Margin (%)	4.64 (32.61)	9.84 (29.89)	9.81 (23.94)	10.04 (26.06)	11.09 (26.68)
Observations					
# of Observations	74,520	2,352	989	61,030	7,196
# of Farmers	7,373	366	152	4,719	827

Table 4: Profit Ledger: Soybean (Mean, SD in parentheses)

Statistic	Conventional Till	No-Till	Cover Crop	Reduced-Till	Conservation Till
Revenues					
Yield (Bu/Acre)	45.30 (15.64)	48.29 (12.47)	51.53 (14.82)	48.01 (12.21)	49.60 (12.42)
Price (\$/Bu)	9.00 (2.97)	9.61 (2.69)	11.82 (2.63)	9.46 (2.64)	9.85 (2.78)
Crop Revenue (\$/Acre)	370.69 (228.30)	448.95 (189.45)	601.68 (194.86)	449.32 (186.56)	475.77 (201.40)
Crop Insurance Income (\$/Acre)	12.62 (71.68)	15.26 (42.73)	32.56 (76.57)	13.01 (40.27)	17.29 (49.82)
Gross Return (\$/Acre)	393.13 (242.44)	476.55 (191.58)	651.66 (181.24)	474.13 (187.38)	504.55 (203.78)
Expenses					
Direct Expenses (\$/Acre)	244.35 (166.51)	288.75 (145.26)	443.46 (141.64)	280.93 (138.42)	309.64 (157.06)
Seed & Plants (\$/Acre)	44.91 (21.96)	47.67 (18.65)	68.10 (21.20)	47.05 (17.78)	47.50 (18.89)
Fertilizer (\$/Acre)	27.16 (34.51)	30.00 (29.96)	58.19 (30.41)	25.68 (26.82)	32.51 (30.64)
Crop Chemicals (\$/Acre)	35.14 (26.12)	38.83 (24.06)	65.88 (24.50)	35.09 (22.00)	41.40 (25.48)
Hired Labor (\$/Acre)	2.82 (10.89)	4.64 (11.60)	24.79 (20.44)	3.62 (11.25)	6.01 (13.35)
Overhead Expenses (\$/Acre)	83.91 (76.11)	90.17 (68.13)	123.59 (90.98)	92.27 (71.51)	96.69 (74.13)
Crop Insurance (\$/Acre)	17.21 (13.17)	17.70 (12.08)	28.06 (15.31)	18.64 (11.58)	19.08 (12.68)
Total Expenses (\$/Acre)	328.26 (201.85)	378.92 (159.83)	567.06 (148.38)	373.20 (151.22)	406.33 (173.79)
Profit					
Net Return (\$/Acre)	64.87 (151.46)	97.63 (134.93)	84.61 (191.77)	100.94 (136.57)	98.22 (145.73)
Net Profit Margin (%)	11.62 (42.44)	17.00 (25.87)	9.69 (27.35)	17.66 (25.19)	16.22 (26.62)
Observations					
# of Observations	54,290	4,474	846	47,392	7,873
# of Farmers	6,399	700	150	4,220	1,072

Table 5: Profit Ledger: Wheat (Mean, SD in parentheses)

Statistic	Conventional Till	No-Till	Cover Crop	Reduced-Till	Conservation Till
Revenues					
Yield (Bu/Acre)	50.83 (18.17)	52.83 (15.03)	60.68 (14.01)	57.36 (16.14)	55.32 (15.81)
Price (\$/Bu)	5.06 (1.82)	5.55 (1.82)	7.03 (1.59)	5.88 (1.72)	6.09 (1.96)
Crop Revenue (\$/Acre)	248.40 (234.61)	303.58 (182.02)	422.73 (152.62)	340.68 (147.07)	345.65 (179.11)
Crop Insurance Income (\$/Acre)	5.71 (28.35)	8.21 (31.16)	10.90 (40.38)	8.55 (31.80)	10.73 (35.96)
Gross Return (\$/Acre)	266.68 (238.01)	323.17 (184.25)	448.47 (162.70)	363.56 (146.90)	372.62 (182.45)
Expenses					
Direct Expenses (\$/Acre)	181.00 (108.80)	252.33 (113.19)	380.41 (104.83)	248.91 (101.66)	282.79 (116.57)
Seed & Plants (\$/Acre)	18.30 (10.57)	22.20 (12.41)	44.18 (31.33)	22.12 (10.73)	25.07 (14.16)
Fertilizer (\$/Acre)	56.66 (38.91)	71.54 (42.37)	117.83 (45.10)	78.25 (39.25)	86.21 (46.52)
Crop Chemicals (\$/Acre)	22.10 (14.08)	26.61 (18.63)	41.60 (24.29)	27.00 (14.22)	31.40 (19.28)
Hired Labor (\$/Acre)	0.88 (4.52)	2.11 (11.80)	2.44 (7.75)	2.44 (7.75)	2.17 (11.75)
Overhead Expenses (\$/Acre)	51.90 (41.70)	58.23 (43.53)	106.47 (59.87)	72.51 (49.58)	76.04 (60.96)
Crop Insurance (\$/Acre)	10.90 (8.06)	12.16 (9.28)	12.20 (6.67)	14.43 (8.24)	13.86 (9.48)
Total Expenses (\$/Acre)	232.90 (128.92)	310.57 (129.25)	486.88 (122.20)	321.42 (116.92)	358.83 (144.66)
Profit					
Net Return (\$/Acre)	33.77 (201.75)	12.60 (159.75)	-38.40 (192.36)	42.15 (114.05)	13.79 (153.85)
Net Profit Margin (%)	4.98 (32.47)	-4.68 (36.11)	-17.42 (42.69)	5.89 (30.47)	-3.43 (34.76)
Observations					
# of Observations	10,876	231	72	6,911	337
# of Farmers	1,839	66	10	1,011	95

4 Soil Data: Predicting Soil Organic Carbon using Machine Learning

4.1 Difficulty in Measuring Spatio-temporal Soil Organic Carbon

Evaluating whether regenerative practices are associated with soil-carbon accumulation requires repeated measurements that can be linked to management at a sufficiently fine spatial scale. This is difficult because soil organic carbon (SOC) changes gradually, soil sampling is costly, and soil tests are not routinely collected from the same fields at regular intervals. A difference-in-differences design therefore requires more than a spatial map of baseline SOC: it needs comparable measurements across locations and years, with treatment timing and location measured consistently. In this study, FinBin farm locations are available only at the county level due to confidentiality concerns, so the outcome must also be constructed at the county-year level. This spatial aggregation limits the analysis to county-level changes and does not reveal changes on the specific fields where practices were adopted.

Existing soil maps provide useful baseline information but do not directly solve this measurement problem. U.S. soil survey products such as SSURGO and STATSGO2 characterize mapped soil properties and are valuable for controlling for persistent differences in soil characteristics, but are not repeated annual observations of SOC for the same farms (USDA Natural Resources Conservation Service, 2026b). SoilGrids provides global predictions of soil properties, including SOC, at 250-m resolution and six standard depth intervals, with pixel-level uncertainty (ISRIC - World Soil Information, 2026). These model-based surfaces support spatial comparison, but a mapped estimate is not a panel of observed county-year changes aligned to FinBin management histories. Other national or global SOC mapping products face the same basic challenge for this application: spatial resolution, measurement years, target depth, and estimation methods may differ from the farm panel, and repeated surfaces need not represent comparable measurements through time. I therefore use these products as context rather than as the outcome for the causal analysis.

4.2 Prediction Using Remotely Sensed Machine Learning

To construct a time-varying proxy, I combine soil-test observations with remotely sensed and other environmental features in a supervised machine-learning model. Optical reflectance and vegetation indices do not measure SOC directly; instead, they provide in-

formation on vegetation, surface conditions, and land cover that can be predictive when combined with soil, terrain, and climate covariates. Prior work demonstrates the feasibility of this approach in specific settings. For example, Chen et al. (2024) compare random forest, XGBoost, and support-vector methods for surface SOC prediction in Shandong using soil samples and multiple environmental data sources. Their results support testing flexible models and combining covariate families, but accuracy in a different region does not establish accuracy for the Upper Mississippi River Basin. I therefore evaluate candidate models against Wisconsin soil-test data rather than assume that published performance transfers to this application.

The target in my training data is county-year mean soil organic matter (OM), as reported by the Wisconsin Soil and Forage Laboratory. I convert predicted OM to approximate SOC using the conventional factor of 0.58; as described in the appendix. The resulting measure is a predicted concentration, not a soil-carbon stock or a field-level measurement. Because FinBin reports farm location only by county, I aggregate the prediction target and spatial predictors to county-year. This aligns their geographic units but cannot identify within-county variation.

4.3 Ground Truth Data

The Wisconsin Soil and Forage Laboratory provides a state-wide soil-testing service and publishes county-year summaries of submitted samples. The usable training panel spans 1995–2019 and includes 1,799 county-year observations across all 72 Wisconsin counties, representing 5,336,425 submitted samples, or about 2,967 samples per county-year on average. The county-year mean OM values range from 0.10% to 21.80%. As discussed in the appendix, the distribution of submissions is not a random sample of Wisconsin cropland, the submitted samples are not necessarily distinct fields or farms, and some county-year averages are based on few samples. The labels should consequently be interpreted as laboratory summaries of submitted tests, not error-free or representative measurements of every field in a county.

I have not obtained equivalent county-year ground-truth data for Minnesota despite efforts to access data from the University of Minnesota Soil Testing and Research Analytical Laboratory (MOSH). The prediction model is therefore trained in Wisconsin and applied to Minnesota without a Minnesota test set. This cross-state application assumes that the relationship between predictors and OM transfers across states.

4.4 Selecting the Best Performing Machine Learning Model

I compare fifteen candidate machine learning regressors rather than selecting a model a priori. The comparison includes a mean-prediction baseline, regularized linear models, support-vector regression, neural-network models, random forests and extremely randomized trees, gradient-boosted trees, and a random-forest/neural-network hybrid. Predictor families include satellite reflectance and vegetation-index summaries, mapped soil characteristics, ecoregion, time terms, and indicators for missing or lagged features. The appendix describes the feature construction and model specifications. The fitted expanded feature set contains 150 predictors after features with no observed values are removed.

The usable Wisconsin panel contains 1,799 county-year observations. I reserve a stratified random sample of 20% of the county observations within each year for the holdout (test) set, leaving 1,449 training observations and 350 holdout observations in the saved comparison. The same county may occur in training in some years and in the holdout in other years, so this split evaluates prediction for held-out county-year observations within the observed period, not generalization to entirely unseen counties or future years. To hypertune these models, I perform a grid search over hyperparameter values using the training set and select the combination that minimizes cross-validated error. This is evaluated using three-fold cross-validation grouped by year.

XGBoost ranks first among the fifteen models on holdout RMSE, narrowly ahead of random forest (0.595 versus 0.605 OM percentage points). For XGBoost, the holdout root mean squared error (RMSE) is 0.595 percentage points, mean absolute error (MAE) is 0.309, and R^2 is 0.587; its year-grouped cross-validation RMSE is 1.198 and R^2 is -0.188 , indicating substantially weaker performance when years are held out.

XGBoost builds an additive ensemble of decision trees, fitting successive trees to reduce the remaining prediction errors and allowing nonlinear relationships and interactions among predictors. I use it to generate the reported predictions because it performs best at prediction out-of-sample. The same holdout set is used to rank the candidate models, so the winning holdout score is a model-selection statistic rather than an independent final validation estimate. The figure below plots the observed and predicted OM values for the complete Wisconsin panel within the training set. Table 6 reports the holdout performance metrics for all fifteen candidate machine learning models in descending order of holdout RMSE.

Table 6: Organic Matter (OM) Machine Learning Models Out-of-Sample Prediction

Model	Test Group	Performance Metrics		
		Root Mean Sq. Error (RMSE)	Mean Abs. Error (MAE)	R2
Extreme Gradient Boosting	Hold Out Group	0.595	0.309	0.587
Random Forest (RF)	Hold Out Group	0.605	0.325	0.573
Random Forest (Tuned)	Hold Out Group	0.606	0.327	0.572
Extra Trees	Hold Out Group	0.608	0.328	0.568
Gradient Boosting	Hold Out Group	0.610	0.362	0.566
RF + MLP Hybrid	Hold Out Group	0.613	0.335	0.562
Extra Trees (Tuned)	Hold Out Group	0.616	0.334	0.557
Histogram Gradient Boosting	Hold Out Group	0.682	0.373	0.457
Support Vector Machines	Hold Out Group	0.768	0.446	0.312
Lasso	Hold Out Group	0.795	0.541	0.263
Elastic Net	Hold Out Group	0.795	0.540	0.262
Multilayer Perceptron (Tuned)	Hold Out Group	0.803	0.505	0.248
Long Short-Term Memory	Hold Out Group	0.837	0.565	0.182
Dummy Classifier (Base)	Hold Out Group	0.928	0.611	-0.005
Multilayer Perceptron	Hold Out Group	0.944	0.552	-0.040

Note: Details on the machine learning model specifications are discussed in Appendix Section [D.3](#). To rank order model selection, my primary consideration is to minimize RMSE for the out-of-sample model prediction. Additionally, MAE and R2 are considered as an explanation for the model explaining variation.




Figure 1: Observed and Predicted County-Year Organic Matter in Wisconsin. Predictions are from the selected XGBoost model; the plotted R^2 is calculated over the complete panel and is not the holdout score.

4.5 Predicting SOC for Minnesota Farmers

I apply the Wisconsin-trained XGBoost model to county-year features for Minnesota, where FinBin farm records are available. The analysis is restricted to Minnesota because I do not have comparable ground-truth samples for the other FinBin states. Predictions for Nebraska or other states would require additional validation. Application to Minnesota is an extrapolation across states, not an estimate calibrated to Minnesota soil tests.

The constructed Minnesota panel used in the FinBin analysis contains 1,392 county-year records for 87 counties over 2010–2025. Its predicted SOC varies across most counties, with 79–80 distinct predictions per year, but only 25.3%–38.7% of the model features are nonmissing per county-year. The county maps and crop-by-practice time series below show spatial and temporal patterns and associations in the linked data; they do not establish that the practices caused SOC changes.




Figure 2: County-Level Predicted SOC and Its Descriptive Association with Adoption in Minnesota




Figure 3: Mean Predicted SOC Over Time by Crop and Management Group in Minnesota. Practice groups are not mutually exclusive; trends are descriptive and are not causal estimates.

5 Causal Identification

5.1 Treatment Definition

I define treatment constructing separate binary management indicators from the tillage and cover cropping responses recorded in FinBin. Rather than combining practices into a single regenerative-agriculture indicator, I estimate separate models for cover-crop use, no-till, reduced tillage, and conservation tillage. The cover-crop indicator equals one when the enterprise is reported as either intercropped with a cover crop or grown after a cover crop. Reduced tillage corresponds to the chisel/reduced-till, and no till corresponds to reported no-till tillage systems. Conservation tillage combines minimum, no-, ridge, and strip tillage to capture a broader range of regenerative tillage practices. Because no-till is included within this broader conservation-tillage definition, these treatment categories overlap and are not mutually exclusive.

For practice p , let A_{it}^p indicate whether farmer i reports practice p in year t , and define the first recorded adoption year as $G_i^p = \min\{t : A_{it}^p = 1\}$. The treatment status used in each event study is then

$$D_{it}^p = \mathbf{1}\{t \geq G_i^p\}, \quad (1)$$

with $G_i^p = \infty$ for farmers who never report the practice during the observed panel. I estimate a separate event study for each practice and crop outcome. In the data construction, however, the first reported year is calculated at the FinBin farmer-ID level across that farmer's crop records; it is not a field-level or crop-specific adoption date. The crop-specific analyses therefore compare crop outcomes according to the farmer's first observed report of the practice, even where the first report may concern another crop.

This treatment definition captures the extensive margin of first reported use, not the intensity, acreage share, or duration of use. The event-study treatment remains in effect after the first report even if a farmer later stops reporting the practice. In addition, crop rotation means that an outcome for a given crop-year may not refer to the same physical field over time. These features can blur the timing and location of exposure and should temper interpretation of the estimated effects as impacts of sustained, field-level adoption.

5.2 Endogeneity Concern

Adoption is not randomly assigned. Farmers may select practices in response to expected yields, soil conditions, weather risk, financial resources, equipment accessibility, experience, or conservation preferences. Some of these factors can also affect farm profitability and soil health as detected through soil organic carbon. A simple comparison of adopters and non-adopters may therefore combine practice effects with pre-existing differences between the groups; the direction of this selection bias is not known a priori.

Figure 4 summarizes the hypothesized relationships among adoption, farm and farmer characteristics, and the outcomes through a directed acyclic graph (DAG). The fixed-effects and staggered difference-in-differences designs address time-invariant differences and common year shocks, while the controlled specifications condition on a limited set of observed covariates.




Figure 4: Directed Acyclic Graph of Regenerative Agriculture’s Relationship to Farm Profit and GHG Storage

Note: The diagram uses the causal-inference notation described by Pearl (2009). Arrows indicate hypothesized causal directions; the dotted arrow denotes a relationship involving a factor that is not observed in the analysis data.

5.3 Identification Strategy

To recover causal effects of adopting regenerative agricultural practices on farm profitability and soil health as measured by soil organic carbon, I employ a difference-in-differences (DiD) methodological framework. I use two-way fixed effects (TWFE) as a benchmark. For each treatment definition and outcome, the event-study specification is

$$Y_{it} = \alpha_i + \lambda_t + \sum_{\substack{k=-5 \\ k \neq -1}}^{15} \beta_k \mathbf{1}\{t - G_i^p = k\} + \gamma' \mathbf{X}_{it} + \varepsilon_{it}, \quad (2)$$

where Y_{it} is the outcome for farmer i in year t , α_i and λ_t are farmer and year fixed effects, and event time $k = -1$ is the omitted reference period. The event window includes leads from -5 to -2 and lags from 0 to 15 . The coefficients β_k describe changes relative to the year immediately before first reported adoption. Equation 2 is estimated separately for each of the three main crop commodity samples: corn, soybean, and wheat. α_i and λ_t capture unobserved time-invariant farmer characteristics and common year shocks, respectively.

The identifying assumption is that, absent adoption, outcomes for farmers first reporting the practice in different years would have evolved in parallel, conditional on the fixed effects and observed covariates. This assumption cannot be established from the data, but may be diagnostically evidenced via event-study leads. Under parallel trends and no anticipation, pre-treatment coefficients provide evidence about the plausibility of the design, while post-treatment coefficients describe dynamic effects relative to the reference period. The appendix reports pre- and post-adoption event-time estimates and alternative staggered-DID specifications.

The controlled TWFE specifications add a vector of covariates $\mathbf{X}_{it}$ including crop-specific growing and killing degree days, the number, age, and farming experience of operators, sole-proprietorship and partnership indicators, and the farm's percent of crop acres owned. Standard errors are clustered at the level of treatment assignment – the farmer ID – to allow errors to be correlated within farmer over time.

TWFE is a useful benchmark, but under staggered adoption its event-time coefficients may combine comparisons between newly treated units, never-treated units, and already-treated units. When treatment effects vary across cohorts or over event time, those comparisons can receive non-convex or negative weights. The resulting coefficients need not equal the cohort-specific average treatment effects. I therefore interpret TWFE as a comparison specification and perform my DID estimation using the cohort-time estimates from the Callaway–Sant'Anna approach below.

5.4 Staggered Difference-in-Differences

To estimate a staggered DiD event study strategy, I create staggered treatment cohorts based on the first reported adoption year for each practice. This complements TWFE with the group-time difference-in-differences estimator of Callaway and Sant'Anna (2021). For each practice, let $G_i^p = g$ denote the first year farmer i reports practice p , and let $G_i^p = \infty$ denote a farmer who never reports that practice in the observed panel. For cohort g and year $t \geq g$, the target is the group-time average treatment effect

$$ATT(g, t) = E[Y_{it}(g) - Y_{it}(\infty) \mid G_i^p = g], \quad (3)$$

where $Y_{it}(g)$ is the outcome under adoption beginning in year g , and $Y_{it}(\infty)$ is the outcome under no adoption during the panel. The estimator compares each adoption cohort with never-treated farmers rather than using already-treated farmers as controls. It thus avoids the particular already-treated comparisons that can make conventional TWFE estimates difficult to interpret when effects vary across cohorts or event time.

I estimate these group-time effects using the doubly robust (DR) variant, with never-treated farmers as the comparison group and a universal base period corresponding to event time $t = -1$ (Callaway and Sant'Anna, 2021). The DR procedure combines an outcome-regression adjustment for the untreated outcome path with propensity-score weighting to balance observed covariates between each treated cohort and its comparison group. Under conditional parallel trends, and overlap identification conditions, the estimator is consistent if either the propensity-score model *or* the outcome-regression model is correctly specified. This property does not address unobserved confounding or guarantee greater precision. But it does construct a broadly flexible estimator that can accommodate heterogeneous treatment effects.

The group-time effects estimated by the doubly robust estimator are expressed as

$$\widehat{ATT}^{DR}(g, t) = \frac{1}{N_g} \sum_{i: G_i = g} [\Delta Y_{it} - \widehat{m}_{0,g,t}(X_i)] - \frac{1}{N_g} \sum_{i: G_i = \infty} \frac{\widehat{p}_g(X_i)}{1 - \widehat{p}_g(X_i)} [\Delta Y_{it} - \widehat{m}_{0,g,t}(X_i)], \quad (4)$$

and aggregated as an average treatment effect on the treated (ATT) as

$$ATT^{CS}(k) = \sum_g w_{g,k} ATT^{CS}(g, g + k). \quad (5)$$

Under this specification, the covariates included for the set of controls in the doubly robust estimation approach are growing degree days (GDD), killing degree days

(KDD), number of farm operators, operator age and years of farming experience, sole-proprietorship and partnership indicators, and percent of crop acres owned. To avoid conditioning on post-adoption changes in these variables, the implementation holds each control at the farmer’s first observed value (baseline-level) in the crop sample. Farmers first observed adopting in the initial sample year are excluded because there is no pre-treatment period (always treated observations). Because farmers enter the panel at different times, I estimate an unbalanced panel.

The dynamic event study effects aggregate the cohort-time $ATT(g, t)$ estimates by event time $k = t - g$. The reported event-study window uses leads through $t \in [-5, -2]$, omits $t = -1$ as the reference period, and follows effects from $t = 0$ through $t = 15$ (a 16-period post-treatment horizon), subject to the cohorts and observations available at each event time. Because questions on cover cropping begin in 2021, this limits the post-treatment horizon. The event-time estimates are weighted averages across the cohorts contributing at that horizon. This is analogous to the TWFE β_k coefficients.

For the results, I report the Callaway and Sant’Anna (2021) doubly robust (DR) estimates. Complementary specifications for TWFE and Sun & Abraham DiD estimates are reported in Appendix F. Their differences help assess sensitivity to the estimator and comparison group.

5.5 Power Calculation

Low adoption rates, particularly for cover crops, may limit the ability to detect economically or agronomically meaningful effects. To characterize this constraint, I calculate the minimum detectable standardized difference for each crop–practice pair at a two-sided significance level of 5% and target power of 80%. For each combination, I reduce the crop panel to one observation per FinBin farmer, averaging the observed outcome across that farmer’s crop records with weights proportional to normal acres. Farmers with a recorded adoption cohort are classified as treated; never-adopters form the comparison group. Let n_1 and n_0 denote the resulting treated and comparison-farmer counts. The minimum detectable effect in Cohen’s d units is the value satisfying

$$\Pr(|T_{n_0+n_1-2}| > t_{1-\alpha/2, n_0+n_1-2} \mid \delta = d, n_0, n_1) = 0.80, \quad \alpha = 0.05, \quad (6)$$

using the unequal-sample-size, two-sample t -test calculation. Cohen’s d standardizes the mean difference by the pooled standard deviation.

Across FinBin-reported outcomes, the 80%-power MDEs range from $d = 0.059$ to $d = 0.418$ across crops and practices. For SOC, the range is $d = 0.080$ to $d = 0.431$. The largest

thresholds occur in the sparsest practice groups, including wheat cover-crop adopters (46 treated farmers). Sample-size configurations are shared across FinBin-reported outcomes with the same observation coverage, while SOC and insurance coverage θ have outcome-specific samples. I display the resulting minimum detectable effects in the figures in Appendix E. These figures plot power against Cohen's d , by crop and management practice, for the FinBin outcomes, predicted SOC, and θ .

6 Results: Impact on Profitability

Is regenerative agriculture profitable for farmers? I begin with net returns per acre, then examine yields, crop revenue, and total expenses to describe the components underlying profitability. The tables report the Callaway–Sant’Anna doubly robust average treatment effect (ATT) across post-adoption cohort-years for each crop and practice, with standard errors and p-values. Monetary estimates are nominal dollars per acre.


Table 7: Callaway and Sant’Anna doubly robust DID summary for Net Return (\$/Acre)

Statistic	Cover Crop	No-Till	Reduced Till	Conservation Till
Corn				
ATT	79.396 (21.060) [≤ 0.001]	-5.243 (5.636) [0.352]	6.869 (3.206) [0.032]	1.513 (5.242) [0.773]
N obs.	51496	50774	47086	49831
N farmers	8967	8943	8787	8919
Mean DV	88.576	88.867	89.087	88.661
Pretrend p-value	0.526	0.140	0.005	0.579
Anticipation p-value	0.788	0.094	0.260	0.570
Soybean				
ATT	23.594 (15.727) [0.134]	-0.235 (3.718) [0.950]	22.503 (2.317) [≤ 0.001]	2.536 (3.525) [0.472]
N obs.	45725	44936	41511	44019
N farmers	7958	7933	7783	7907
Mean DV	83.280	83.639	82.001	83.522
Pretrend p-value	0.015	0.224	0.573	0.933
Anticipation p-value	0.995	0.325	0.948	0.584
Wheat				
ATT	31.458 (32.188) [0.328]	0.780 (5.861) [0.894]	-5.831 (5.530) [0.292]	-1.838 (5.579) [0.742]
N obs.	9840	9713	9124	9537
N farmers	2195	2181	2162	2169
Mean DV	31.654	31.813	32.925	31.630
Pretrend p-value	0.960	0.024	0.682	0.506
Anticipation p-value	0.758	0.130	0.279	0.517


The net-return estimates are positive for cover crops in corn (\$79.40 per acre, $p < 0.001$) and reduced tillage in soybeans (\$22.50, $p < 0.001$). Other point estimates are less conclusive: cover-crop returns are \$23.59 higher for soybeans ($p = 0.134$) and \$31.46 higher for wheat ($p = 0.328$); reduced-till returns are \$6.87 higher for corn ($p = 0.032$) but \$5.83 lower for wheat ($p = 0.292$). The no-till and conservation-tillage net-return estimates are not statistically distinguishable from zero in any of the three crops. Thus, the results do not support a general claim that each practice raises net returns across crops.

The event-study plots show the cohort-aggregated effects over event time and their uncertainty; the average post-treatment estimates should be read alongside this dynamic


evidence as the benefits from maintained regenerative practices should accrue over time with soil capitalization. The paths are noisy and do not show a uniform, sustained increase in returns across crops and practices. The pre-trend tests also flag departures from zero leads for corn reduced tillage ($p = 0.005$), soybean cover crops ($p = 0.015$), and wheat no-till ($p = 0.024$). These diagnostics weaken a causal interpretation for those comparisons.




(a) Cover Crop



(b) Conservation Tillage



(c) No Till



(d) Reduced Tillage

Figure 5: Net-return event studies by regenerative practice.

Table 8: Callaway and Sant'Anna doubly robust DID summary for Yield (Bu/ Acre)

Statistic	Cover Crop	No-Till	Reduced Till	Conservation Till
Corn				
ATT	9.924 (9.314) [0.287]	3.554 (3.559) [0.318]	-8.136 (3.692) [0.028]	0.266 (3.351) [0.937]
N obs.	51496	50774	47086	49831
N farmers	8967	8943	8787	8919
Mean DV	227.464	228.523	228.125	228.747
Pretrend p-value	0.447	0.569	0.035	0.886
Anticipation p-value	0.815	0.412	0.152	0.525
Soybean				
ATT	0.560 (1.254) [0.655]	1.126 (0.457) [0.014]	1.364 (0.358) [;0.001]	1.553 (0.406) [;0.001]
N obs.	45725	44936	41511	44019
N farmers	7958	7933	7783	7907
Mean DV	45.878	45.866	45.471	45.837
Pretrend p-value	0.065	0.683	0.055	0.711
Anticipation p-value	0.152	0.775	0.470	0.944
Wheat				
ATT	2.327 (5.165) [0.652]	0.862 (1.235) [0.485]	3.110 (1.087) [0.004]	0.379 (1.103) [0.731]
N obs.	9840	9713	9124	9537
N farmers	2195	2181	2162	2169
Mean DV	51.950	52.025	51.528	51.987
Pretrend p-value	0.607	0.012	0.330	0.081
Anticipation p-value	0.701	0.496	0.094	0.659

Table 9: Callaway and Sant'Anna doubly robust DID summary for Crop Revenue (\$/Acre)

Statistic	Cover Crop	No-Till	Reduced Till	Conservation Till
Corn				
ATT	105.622 (27.755) [;0.001]	13.974 (7.805) [0.073]	73.715 (5.678) [;0.001]	29.339 (7.476) [;0.001]
N obs.	51496	50774	47086	49831
N farmers	8967	8943	8787	8919
Mean DV	587.058	587.425	577.698	586.795
Pretrend p-value	0.699	0.052	0.429	0.041
Anticipation p-value	0.933	0.149	0.346	0.044
Soybean				
ATT	41.627 (21.733) [0.055]	17.750 (5.275) [;0.001]	53.477 (4.401) [;0.001]	26.185 (4.936) [;0.001]
N obs.	45725	44936	41511	44019
N farmers	7958	7933	7783	7907
Mean DV	395.343	395.691	387.569	395.290
Pretrend p-value	0.147	0.034	0.923	0.578
Anticipation p-value	0.518	0.149	0.527	0.233
Wheat				
ATT	-0.585 (29.178) [0.984]	8.829 (8.801) [0.316]	47.766 (8.152) [;0.001]	11.285 (8.024) [0.160]
N obs.	9840	9713	9124	9537
N farmers	2195	2181	2162	2169
Mean DV	276.893	278.028	272.795	277.825
Pretrend p-value	0.692	0.110	0.009	0.353
Anticipation p-value	0.344	0.205	0.138	0.502

Table 10: Callaway and Sant'Anna doubly robust DID summary for Total Expenses (\$/Acre)

Statistic	Cover Crop	No-Till	Reduced Till	Conservation Till
Corn				
ATT	21.832 (19.794) [0.270]	19.516 (6.662) [0.003]	67.372 (5.142) [;0.001]	29.523 (6.173) [;0.001]
N obs.	51496	50774	47086	49831
N farmers	8967	8943	8787	8919
Mean DV	528.140	528.345	517.828	527.947
Pretrend p-value	0.092	0.555	0.433	0.350
Anticipation p-value	0.608	0.563	0.831	0.129
Soybean				
ATT	22.725 (15.438) [0.141]	19.204 (4.739) [;0.001]	33.183 (3.921) [;0.001]	25.765 (4.266) [;0.001]
N obs.	45725	44936	41511	44019
N farmers	7958	7933	7783	7907
Mean DV	334.285	334.358	327.168	334.105
Pretrend p-value	0.552	0.085	0.393	0.263
Anticipation p-value	0.328	0.423	0.765	0.236
Wheat				
ATT	-11.934 (38.506) [0.757]	17.011 (7.796) [0.029]	54.873 (6.673) [;0.001]	20.566 (6.909) [0.003]
N obs.	9840	9713	9124	9537
N farmers	2195	2181	2162	2169
Mean DV	267.995	268.845	261.385	268.886
Pretrend p-value	0.709	0.710	0.004	0.900
Anticipation p-value	0.492	0.408	0.327	0.361

We can decompose net returns into the three major components: crop revenue, total expenses, and yields. The component outcomes suggest that revenue and cost responses differ across crop-practice pairs. Crop revenue is higher for corn cover-crop (\$105.62 per acre, $p < 0.001$), reduced-till (\$73.72, $p < 0.001$), and conservation-till (\$29.34, $p < 0.001$) observations, and for soybean no-till (\$17.75), reduced tillage (\$53.48), and conservation tillage (\$26.19) ($p < 0.001$ for each). The soybean cover-crop estimate is also positive (\$41.63, $p = 0.055$). The wheat revenue increase is concentrated in reduced tillage (\$47.77, $p < 0.001$). Total expenses also rise for several groups: corn no-till (\$19.52), reduced tillage (\$67.37), and conservation tillage (\$29.52); soybean no-till (\$19.20), reduced tillage (\$33.18), and conservation tillage (\$25.77); and wheat no-till (\$17.01), reduced tillage (\$54.87), and conservation tillage (\$20.57) (all $p < 0.05$). For yields, soybean no-till, reduced-till, and conservation-till estimates are positive (1.13, 1.36, and 1.55 bushels per acre, respectively; all $p \leq 0.014$), as is the wheat reduced-till estimate (3.11 bushels, $p = 0.004$); the corn reduced-till estimate is negative (-8.14 bushels, $p = 0.028$). The component pattern therefore points to both revenue and expense changes, rather than a uniform yield or cost-saving pathway. Some component estimates also have significant pre-trend tests, including soybean no-till crop revenue and wheat reduced-till crop revenue and expenses.

In Appendix [G](#), I present some secondary related outcomes indicative of profitability related to regenerative agriculture. The secondary insurance outcomes provide limited evidence of a common insurance response. Crop-insurance income per acre is higher for reduced tillage in corn (\$2.28, $p = 0.009$), soybeans (\$2.01, $p < 0.001$), and wheat (\$2.72, $p = 0.001$); most other crop-practice estimates are not statistically significant. Insurance-coverage θ is higher for corn cover crops (0.045, $p = 0.036$) and reduced tillage in corn (0.033, $p < 0.001$), soybeans (0.030, $p < 0.001$), and wheat (0.094, $p < 0.001$). This is suggestive that regenerative farmers are more likely to have higher insurance coverage, although the evidence is not uniform across all practices and crops.

7 Results: Impact on Soil Organic Carbon (SOC)

The estimated effects on SOC concentration are heterogeneous across crops and practices. In corn, reduced tillage is associated with a 0.177-percentage-point (p.p.) increase ($p < 0.001$); the estimated effects for cover crops, no-till, and conservation tillage are not statistically distinguishable from zero. In soybeans, the estimates are positive for no-till (0.134 p.p., $p = 0.048$), reduced tillage (0.180, $p < 0.001$), and conservation tillage (0.105, $p = 0.031$), while the cover-crop estimate is not significant. The wheat estimates are

also not statistically distinguishable from zero: the event-study summaries are 0.060 p.p. for cover crops, 0.039 for no-till, -0.246 for reduced tillage, and 0.134 for conservation tillage. The negative reduced-tillage estimate is imprecise ($p = 0.076$). Overall, the estimates provide some evidence of higher predicted SOC for particular corn and soybean practice combinations, but not a consistent increase across all crops and practices.

Table 11: Callaway and Sant’Anna doubly robust DID summary for Soil Organic Carbon (%)

Statistic	Cover Crop	No-Till	Reduced Till	Conservation Till
Corn				
ATT	-0.011 (0.046) [0.817]	0.047 (0.061) [0.441]	0.177 (0.039) [$\hat{0}$.001]	0.029 (0.044) [0.507]
N obs.	28784	27021	15621	26213
N farmers	5357	5120	3642	5065
Mean DV	2.307	2.321	2.339	2.326
Pretrend p-value	0.056	0.719	0.692	0.833
Anticipation p-value	0.005	0.957	0.274	0.381
Soybean				
ATT	0.040 (0.051) [0.433]	0.134 (0.068) [0.048]	0.180 (0.048) [$\hat{0}$.001]	0.105 (0.049) [0.031]
N obs.	25333	23585	12834	22819
N farmers	4742	4511	3139	4455
Mean DV	2.440	2.462	2.483	2.471
Pretrend p-value	0.008	0.629	0.095	0.352
Anticipation p-value	$\hat{0}$.001	0.838	0.554	0.116
Wheat				
ATT	0.060 (0.142) [0.673]	0.039 (0.113) [0.731]	-0.246 (0.139) [0.076]	0.134 (0.118) [0.256]
N obs.	4978	4761	2441	4674
N farmers	1206	1149	730	1137
Mean DV	3.192	3.219	3.152	3.230
Pretrend p-value	0.519	$\hat{0}$.001	$\hat{0}$.001	0.027
Anticipation p-value	0.708	0.112	0.616	0.154

The event-study profiles reveals the dynamic patterns of SOC concentration though noisy at longer horizons where fewer cohorts contribute. Pre-trend tests reject the zero-lead null for soybean cover crops ($p = 0.008$) and for wheat no-till, reduced tillage, and

conservation tillage ($p < 0.001$, $p < 0.001$, and $p = 0.027$, respectively). These departures make a causal interpretation of those comparisons less secure. The estimated outcome is a county-year SOC concentration, not a direct measurement of carbon on treated fields.

7.1 Implied Carbon Sequestration

To express the predicted-SOC effects in carbon-stock units, I apply a fixed-depth conversion under explicit mineral-soil assumptions: a 30 cm soil layer, dry bulk density of 1.3 g/cm³, and zero coarse-fragment volume fraction. These assumptions are coarse approximations of typical mineral-soil conditions for Minnesota agriculture. Let $\hat{\tau}_{c,m}(k)$ denote the event-study effect on predicted SOC concentration, in percentage points, for crop c , practice m , and event time k , relative to the pre-adoption reference period. Multiplying this effect by bulk density, depth, and the fine-soil fraction yields tonnes of carbon per hectare. And multiplying by 44/12 converts carbon to its CO₂ equivalent. Under these assumptions, one percentage point of SOC corresponds to 39 tonnes C/ha, or 143 tonnes CO₂/ha. This calculation converts a concentration difference at a particular horizon into an implied stock difference relative to the counterfactual management practice. It does not integrate the event-study curve, whose integral would have units of percentage-point-years, or identify the year-by-year sequestration path. To compare these stocks with annual production emissions, I convert each endpoint stock difference into an average annual flow by dividing by the number of years elapsed since the reference period: $k + 1$ years at event time k .

The conversion is conditional on the predicted concentration representing the assumed soil layer and on comparable soil mass across management regimes. Management-induced changes in bulk density would require an equivalent-soil-mass comparison, and the assumed density is not appropriate for organic soils. The calculations should therefore be read as illustrative stock equivalents, not verified removals of atmospheric carbon. Their confidence intervals propagate the event-study standard errors but omit uncertainty in SOC prediction, soil depth, bulk density, and coarse fragments.

Through the event-study horizon of $k = 0, \dots, 14$, corn's implied carbon-stock differences are 15.81 tonnes C/ha under conservation tillage, 7.91 under no-till, and 12.51 under reduced tillage, equivalent to 57.97, 28.99, and 45.88 tonnes CO₂/ha, respectively. The corresponding soybean estimates are 11.23, 6.03, and 11.78 tonnes C/ha, or 41.17, 22.09, and 43.21 tonnes CO₂/ha. These positive point estimates are imprecise: every corn and soybean endpoint's 95% confidence interval includes zero. For example, reduced tillage in corn has an interval of $[-0.59, 92.35]$ tonnes CO₂/ha, and in soybeans an interval of

$[-5.72, 92.13]$. At $k = 4$, the cover-crop endpoints are negative for corn and soybeans (-22.03 and -20.11 tonnes CO_2/ha), also with intervals spanning zero.

Table 12 reports these average annual flows per acre. For crop c and practice m , the implied annual removal is

$$R_{c,m} = 0.40468564224 \times \frac{\Delta\text{CO}_{2,c,m}(k)}{k+1},$$

where $\Delta\text{CO}_{2,c,m}(k)$ is the CO_2 -equivalent stock difference per hectare and 0.40468564224 converts hectares to acres. Positive values denote removals. I construct baseline annual production emissions per acre from Table 1 as 134.52 MMT CO_2e divided by 91.07 million acres for corn, or 1.477 tonnes $\text{CO}_2\text{e}/\text{acre}/\text{year}$, and 40.62 MMT CO_2e divided by 75.82 million acres for soybeans, or 0.536 tonnes $\text{CO}_2\text{e}/\text{acre}/\text{year}$. Implied net emissions are this baseline minus $R_{c,m}$; negative values indicate an implied net removal.

The net calculation holds production emissions at their Table 1 intensity and therefore credits adoption only with the estimated SOC change. Relative to conventional management, the implied change in net emissions is $-R_{c,m}$. It does not account for possible practice-induced changes in fuel use, fertilizer, N_2O , yields, or other production emissions. The baseline is also an all-acre commodity average, not an estimate for the counterfactual conventional acres in the sample.

For national scale, I apply the per-acre flows to the USDA NASS harvested area benchmarks used in this exercise: 90.948 million corn acres and 85.881 million soybean acres. I apply Table 1's per-acre emissions to the same acreage so that both terms in the net calculation use a common area base; the resulting baselines are 134.34 and 46.01 MMT $\text{CO}_2\text{e}/\text{year}$. These are separate full-area scenarios for each practice, assuming that the estimated Minnesota response applies uniformly to national acreage and that all acres experience the same adoption horizon. They are not estimates of feasible additional adoption: the acreage includes existing adopters, and soils, climate, baseline SOC, and management differ across regions.

Table 12: Implied Annual SOC Removals and Net Crop Emissions

Practice	Crop	Endpoint k	Per acre (t CO ₂ e/year)			National (MMT CO ₂ e/year)	
			Removal	Baseline	Net	Removal	Net
Cover crops	Corn	4	-1.783	1.477	3.260	-162.15	296.49
Cover crops	Soybean	4	-1.628	0.536	2.163	-139.77	185.78
Conservation tillage	Corn	14	1.564	1.477	-0.087	142.24	-7.90
Conservation tillage	Soybean	14	1.111	0.536	-0.575	95.38	-49.37
No-till	Corn	14	0.782	1.477	0.695	71.14	63.20
No-till	Soybean	14	0.596	0.536	-0.060	51.19	-5.18
Reduced tillage	Corn	14	1.238	1.477	0.239	112.57	21.77
Reduced tillage	Soybean	14	1.166	0.536	-0.630	100.11	-54.10

Notes: Scenarios use the Callaway Sant’Anna DR endpoint estimates, a 30 cm depth, bulk density of 1.3 g/cm³, and no coarse fragments. Removals are average annual flows equal to endpoint stock differences divided by $k + 1$; positive removals reduce net emissions. Baselines are Table 1 annual emissions per acre. Net equals baseline minus removal and assumes no other practice-induced change in production emissions. National columns use 90.948 million corn acres and 85.881 million soybean acres. All endpoint 95% confidence intervals include zero.


Over 15 years, reduced tillage implies average removals of 1.238 tonnes CO₂/acre/year in corn and 1.166 in soybeans. These flows would offset 83.8% of corn’s baseline production emissions and 217.6% of soybeans’. The resulting net emissions are 0.239 tonnes CO₂e/acre/year for corn and -0.630 for soybeans, implying a net removal in soybeans. Conservation tillage also yields point estimates below zero for both crops, at -0.087 and -0.575 tonnes CO₂e/acre/year. No-till implies 0.695 for corn and -0.060 for soybeans. These net-negative point estimates do not establish that production removes more carbon than it emits. All the underlying endpoint intervals include zero; for example, the 95% intervals for reduced-tillage net emissions are [-1.014, 1.493] for corn and [-1.950, 0.690] for soybeans. Short-horizon cover-crop flows are negative, raising net emissions in the point estimates, but are similarly imprecise.

Nationally, reduced tillage implies 112.57 and 100.11 MMT CO₂/year of average removals for corn and soybeans. Net emissions would fall from 134.34 to 21.77 MMT CO₂e/year for corn and from 46.01 to -54.10 MMT CO₂e/year for soybeans. Their respective 95% intervals are [-92.25, 135.78] and [-167.46, 59.27] MMT CO₂e/year. Average annual flows are not constant sequestration rates: SOC gains may be concentrated in some years, may slow as soils approach equilibrium, and may be reversed. Accordingly, these figures summarize average flows over each observed horizon rather than permanent annual mitigation.


To illustrate the potential social value, I use the archived EPA social cost of CO₂: \$46 per metric tonne for emissions in 2025 at the central 3% discount-rate estimate, expressed

in 2007 dollars.² Multiplying each stock equivalent by this fixed benchmark gives illustrative one-time stock values. Reduced tillage implies \$77.68 billion for corn and \$69.08 billion for soybeans, or \$854.07 and \$804.31 per acre, respectively. Under the same assumptions, the other tillage scenarios imply \$411.30–\$1,079.17 per acre. These are one-time stock-equivalent valuations, not annual compensation rates; the negative cover-crop values similarly reflect imprecise negative endpoints rather than established climate damages. A valuation of the actual sequestration path would require year-specific changes in stocks, year-specific social costs, discounting to a common date, and accounting for subsequent reversals.


²See https://19january2017snapshot.epa.gov/climatechange/social-cost-carbon_.html. This is the historical benchmark specified for the exercise, not a current EPA estimate or an observed carbon-market price.




(a) Cover Crop



(b) Conservation Tillage



(c) No Till



(d) Reduced Tillage

Figure 6: SOC event studies by regenerative practice.

8 Modeling Soil Capitalization into Farmer Decision-Making

8.1 Model Motivation

Soil-capital accumulation is gradual process. Even a 25-year panel may not fully capture the long-term effects. Agronomists argue that meaningful change from regenerative agricultural practices may take several decades to manifest – if not a century. So the 25-year panel may capture only part of the response to a change in management. I therefore develop a novel model mapping how soil capitalization influences farmer management adoption with the primary motivation of profit maximization. I model soil stocks as a state variable that links contemporaneous production and management choices to future profitability, rather than treating the observed panel horizon as sufficient to reveal long-run effects. The objective of this model is to recover the parameters governing profit responses given the state of soil capital.

The model proceeds from the farmer’s objective to maximize both contemporaneous profits (e.g., yields this season) as well as the selection of crop insurance to hedge against unpredictable market outcomes. Both of these profit factors are determined by production yields by the farmer. These yields are determined by a mixture of the inputs (additives) as well as the quality of the soil. Farmers can modify the soil capital through varying management practices (e.g., regenerative practices) which may incur short-run production losses to provide long-run stability in the soil. This model allows us to predict why farmers may choose to adopt certain management practices considering the financial and soil implications of that decision. As evidenced by low adoption of regenerative practices in the FinBin and NASS data, this model can assist in illuminating frictions and incentives for adoption.

8.2 Profit Maximization

I embed the decision to select a management practice within an intertemporal profit maximization framework. Farmer i selects a management practice m_t each period to maximize profit. The farmer starts with the default management practice m_0 (conventional tillage, no cover cropping) at the beginning of the planning horizon. Because we wish to model long-run impacts of adoption, once a farmer adopts a new practice I force them to continue with that practice in subsequent periods.

The farmer maximizes expected discounted utility over a finite horizon T , choosing management practices m_t , variable inputs $\vec{z}_t$, and insurance coverage θ_t each period. Contemporaneous profit for seasonal production under management practice m_t is:

$$\pi_t(m_t, m_{t-1}; \vec{z}_t) = \underbrace{p_t \cdot y_t(m_t, X_t, \vec{z}_t) + \iota_t(\theta_t, m_t)}_{\text{Revenue}} - \underbrace{P(m_t) - \sum_{j=1}^J w_j c_{jt}}_{\text{Cost}} - \kappa_t(m_t, m_{t-1}) \quad (7)$$

where p_t is the output price, $y_t(m_t, X_t, \vec{z}_t)$ is yield, $\iota_t(\theta_t, m_t)$ is the crop insurance payout, $P(\cdot)$ is the insurance premium paid, w_j are input prices for inputs c_{jt} , and $\kappa_t(m_t, m_{t-1})$ is the switching cost of adoption. The state vector includes the soil nutrient stocks $\mathbf{X}_t = (X_{1t}, \dots, X_{Kt})'$, which evolve endogenously with the farmer's decisions.

For farmer i , the intertemporal problem is to choose the sequence $\{m_t, \{x_{jt}\}_{j=1}^J, \{f_{kt}\}_{k=1}^K, \theta_t\}_{t=0}^T$ to maximize expected discounted profits, $\mathbb{E}_{i,0} \left[\sum_{t=0}^T \beta^t \pi_t(m_t, m_{t-1}; \mathbf{z}_t) \right]$, where $\beta \in (0, 1)$ is the endogenous discount factor. At the start of period t , the relevant state is $(\mathbf{X}_t, m_{t-1})$: soil stocks affect current production, while the previous practice determines whether a change in management incurs a switching cost. The choices $(m_t, \{x_{jt}\}, \{f_{kt}\}, \theta_t)$ determine current profit and, through the evolution of $\mathbf{X}_t$, the production opportunities available in subsequent periods. Consistent with the adoption restriction stated above, after farmer i adopts a practice, the management choice is held at that practice in later periods; thus, the model does not allow subsequent switching back to the initial practice or to another practice. The partial-budget comparison holds market and other farm conditions fixed (viz., *certius paribus*) when evaluating differences in production across management choices.

8.3 Adoption Costs

I distinguish normal production inputs $\vec{z}_t$ from the adoption-specific management costs captured by $\kappa_t(m_t, m_{t-1})$. These fixed and variable adoption costs only incur when a farmer adopts or switches management practices. The fixed component $\kappa^{\text{fixed}}(m_t)$ captures up-front or amortized costs such as tillage equipment. The fixed component $\kappa^{\text{fixed}}(m_t)$ is incurred only in the period a farmer switches practices — e.g., first-time adopters — zero otherwise. The variable component $\kappa^{\text{var}}(m_t)$ reflects how practice alters input expenditure: tillage management may reduce/increase fuel and labor costs, while cover cropping can substitute for herbicide application and requires cover crop seed purchases. This allows κ^{var} to be either positive or negative. The variable component $\kappa^{\text{var}}(m_t)$ reflects how the ongoing practice alters input expenditure and recurs every period regardless of switching.

The practice cost $\kappa_t(m_t, m_{t-1})$ captures these two types of costs as

$$\kappa_t(m_t, m_{t-1}) = \mathbf{1}\{m_t \neq m_{t-1}\} \kappa^{\text{fixed}}(m_t) + \kappa^{\text{var}}(m_t). \quad (8)$$

8.4 Crop Insurance

Inside the profit function, I include the cost of crop insurance as a deduction from revenue. The omnipresent risk of yield loss makes crop insurance a standard component of the farmer's decision-making framework. The farmer purchases a yield-protection contract at coverage level $\theta \in [0, 1]$ against the ten-year Actual Production History (APH) benchmark.

The insurance premium paid by the farmer is determined by

$$P(m_t) = \lambda \cdot \theta_t \cdot p^* \cdot \bar{y}_{APH,t}, \quad (9)$$

where λ is premium rate for coverage $\theta_t \in [0, 1]$ at projected price p^* at historical production yield $\bar{y}_{APH,t}$. The term $p^* \cdot \bar{y}_{APH,t}$ anchors the premium to the expected value of insured production: farmers with higher demonstrated productivity face proportionally larger nominal premiums.

The Actual Production History (APH) is a standard metric used to determine a farmer's historical yield performance over the past ten years. The APH provides a conservative estimate of expected production, which in turn affects insurance premiums and indemnity calculations.

$$\bar{y}_{APH,t} = \frac{1}{10} \sum_{t=10}^{t-1} y_t(X_t, m_t). \quad (10)$$

When realized yield falls below the guarantee, the indemnity payment is:

$$\iota_t(\theta_t, m) = \max(0, \theta_t \cdot \bar{y}_{APH,t} - y_t(m)) \cdot p^*, \quad (11)$$

where p^* is the projected price. An actuarially fair insurance contract is one where the expected indemnity equals the premium. Farmers choosing actuarially fair insurance coverage (NVI) are indifferent at the following level.

$$\text{NVI}(\theta_t, m) = \mathbb{E}[\iota_t(\theta_t, m_t)] - P(m_t). \quad (12)$$

8.5 Production Function

At the core of the profit function is a production function driven by yield response to soil health and other inputs managed by the farmer. I model the production function from the agronomic notion of Liebig's law of the minimum. Functionally, it takes the Mitcherlich-Baule multiplicative form.

The yield response curve is described by

$$y_t(m_t, X_t, \theta_t; \vec{Z}_t) = a(X_t, m_t) \prod_{k=1}^K (1 - \beta_k(X_t) \exp(-\gamma_k(X_t) \cdot Z_{kt})), \quad (13)$$

where X_{kt} is the stock of macronutrient k (e.g., soil organic carbon) at the start of period t . $\beta_k(m) \in (0, 1)$ is the maximum yield penalty from complete absence of nutrient k ; and $\gamma_k(m) > 0$ governs how quickly productivity recovers as the stock is replenished. Each factor $(1 - \beta_k(X_t) \exp(-\gamma_k(X_t) \cdot Z_{kt}))$ is a saturation function bounded in $(1 - \beta_k(m), 1)$. Z_{kt} are weather conditions for rainfall and growing degree days moderating the yield response. The technology parameter $a_t(m)$ is a level indicator for baseline productivity of a given field, capturing the inherent productive capacity independent of the current nutrient stocks. The product structure $(\prod_{k=1}^K)$ ensures that a deficiency in any single nutrient constrains overall soil productivity regardless of the abundance of others.

8.6 Nutrient Stock Dynamic

At the heart of the dynamic yield response is a nutrient stock dynamic. The following state equation describes the evolution of soil organic carbon as the limiting soil health factor driving soil productivity over time. The state equation is

$$X_{t+1}(m) = \delta(m) X_t + \phi(m) y_t, \quad (14)$$

where $\delta(m) \in (0, 1)$ is the nutrient retention rate, and $\phi(m) y_t$ represents soil capitalization, which scales with yield. Figure 7 illustrates the yield response of this nutrient stock dynamic under different management regimes: a conventional practice (m_0) versus a regenerative practice (m_1). In the bottom panel I describe the yield response production function along a range of growing conditions (Z). The realized yield maps into the top panel displaying the distribution of potential yield outcomes conditional of the underlying state of soil health (X).

The conventional regime has a higher ceiling ($a_0 > a_1$) but larger downside spread ($\beta_0 > \beta_1$), while the regenerative regime has a higher intercept ($A_1(0) > A_0(0)$) and a

narrower yield range. The right-angle slope markers show that initial marginal responsiveness is steeper under conventional management ($c_0 > c_1$). This highlights two major insights from the construction of the state equation and the production function.

First, conventional practices have the potential for both higher and lower yields. Regenerative practices are not production improving technology but rather a risk-reducing technology. The narrower spread in potential yield response results implies that, given similar expected growing conditions, a regenerative farmer has lower expectation of extreme outcomes, both positive and negative, compared to a conventional farmer. Removing bad seasons should lead to lower crop insurance coverage.

Second, conventional practices will be far more sensitive to growing conditions. Note in the lower panel, as growing conditions weaken (move to the left) $A_1(0) > A_0(0)$. This implies that under poor growing conditions, regenerative practices provide a yield advantage due to their reduced sensitivity to environmental stress, highlighting their role as a risk-reducing strategy rather than a yield-maximizing one.

This figure describes the heart of the trade off and friction with adoption of regenerative practices. This requires farmers to forgo short-run windfalls from high input production to reap minimized variance of returns in poorly performing seasons.




Figure 7: Yield response under conventional high-input management (m_0) versus regenerative low-input management (m_1)

8.7 Bellman Equation

Because management affects the soil state $\mathbf{X}_t$ through the retention and uptake parameters, and because $\mathbf{X}_t$ propagates forward via equation (14), the choice of m generates a path of soil health that influences yield, insurance payouts, and premiums over the full planning horizon. The farmer's problem is therefore a dynamic program in which the state vector $\mathbf{X}_t$ links current practice choices to future payoffs through both the production function and the APH benchmark in the insurance contract.

The Bellman equation formalizing the value function for this dynamic programming problem is given by

$$V(m_t, m_{t-1}, \mathbf{X}_t) = \max_{m_t, m_{t-1}, \theta_t} \pi_t(m_t, m_{t-1}) + \delta V(\mathbf{X}_{t+1}, m_t). \quad (15)$$

m_{t-1} enters the state because the fixed switching cost $\kappa^{\text{fixed}}(m_{t-1} \rightarrow m_t)$ is incurred only in the period of transition.

To characterize the solution to (15) in the steady state, fix management at m^* and hold

the distribution of exogenous production conditions constant over time. A deterministic steady state is a soil stock $\mathbf{X}^*$ and yield y^* that are invariant under the nutrient transition and production function:

$$\mathbf{X}^* = \delta(m^*) \odot \mathbf{X}^* + \boldsymbol{\phi}(m^*)y^*, \quad y^* = y(m^*, \mathbf{X}^*, \theta^*; \mathbf{Z}^*), \quad (16)$$

where $\odot$ denotes element-by-element multiplication, and $\delta(m^*)$ and $\boldsymbol{\phi}(m^*)$ collect the nutrient-specific retention and capitalization parameters. If $0 < \delta_k(m^*) < 1$, the transition equation implies $X_k^* = \phi_k(m^*)y^*/[1 - \delta_k(m^*)]$ for each nutrient k . Under stationary insurance and production conditions, the APH benchmark is stationary as well; insurance coverage θ^* is optimal when it maximizes expected discounted profit given the steady-state soil stock and yield distribution. Thus, the steady state is a fixed point of the soil transition under the farmer's optimal stationary choices, together with the corresponding production and insurance outcomes. It is a stationary farm-level allocation, not a market-clearing equilibrium; no market-clearing conditions are imposed in this model.

The steady-state conditions characterize the model's long-run benchmark, while the moment conditions connect its parameters to the observed farm data. Let $\boldsymbol{\psi}$ collect the unknown production, soil-transition, management-cost, and insurance parameters, and let $\mathbf{q}_{it}(\boldsymbol{\psi})$ denote the model-implied vector of outcomes for farmer i in period t under the model's optimal policy. For observed data $\mathbf{d}_{it}$, let $\mathcal{T}_i$ denote the periods observed for farmer i and let $T_i = |\mathcal{T}_i|$. Define that farmer's moment contribution as

$$\mathbf{g}_i(\boldsymbol{\psi}) \equiv \frac{1}{T_i} \sum_{t \in \mathcal{T}_i} [\mathbf{h}(\mathbf{d}_{it}) - \mathbf{h}(\mathbf{q}_{it}(\boldsymbol{\psi}))], \quad (17)$$

and the population moment restrictions are $\mathbb{E}_i[\mathbf{g}_i(\boldsymbol{\psi}_0)] = \mathbf{0}$, where $\boldsymbol{\psi}_0$ is the parameter vector that makes the model reproduce the selected data moments and $\mathbf{h}(\cdot)$ maps outcomes into the moments used for estimation.

9 Conclusion

Agricultural conservation presents a private investment problem with potential public benefits. Cover crops and conservation tillage can change current production costs and yields while influencing soil conditions that matter for future returns. Their agronomic promise, however, does not establish that adoption will be privately profitable across crops and farms. This paper examines that distinction by combining FinBin farmer-year-crop financial records with a county-year panel of predicted SOC and estimating practice-

specific responses under staggered adoption.

The financial results show heterogeneous returns rather than a general profitability advantage. Estimated net returns are higher for corn cover crops and soybean reduced tillage, but most other crop-practice estimates are not statistically distinguishable from zero. Yield, revenue, and expense responses also differ across comparisons, so higher net returns cannot be attributed to a common yield mechanism. Insurance outcomes provide an additional perspective on private incentives, although the reported effects do not establish that adoption consistently reduces production risk. The event-study evidence likewise does not reveal a uniform, sustained increase in returns, and differential pre-adoption trends weaken causal interpretation for some comparisons.

The soil results provide complementary but more limited evidence. Positive estimated changes in predicted SOC occur for selected crop-practice comparisons, including corn reduced tillage and several soybean tillage practices. These estimates concern county-year SOC concentration, not repeated measurements of soil on adopters' fields. Spatial aggregation, prediction uncertainty, and pre-trend concerns limit the conclusions that can be drawn about practice-induced soil improvements. To illustrate their possible carbon implications, I convert the tillage endpoints into stock equivalents under fixed assumptions about soil depth, bulk density, and coarse fragments. These calculations are conditional scenarios, not direct measurements of cumulative sequestration or actual emissions reductions.

For the study time horizon, reduced tillage implies 45.88 tonnes CO₂/ha for corn and 43.21 for soybeans. Scaling to national harvested acreage gives stock equivalents of 1,688.61 and 1,501.63 million tonnes CO₂. Averaged over 15 years, these stock differences imply annual removals of 112.57 and 100.11 MMT CO₂. Holding Table 1 production-emission intensities fixed, net emissions fall to 21.77 MMT CO₂e/year for corn and -54.10 for soybeans. These are average-flow scenarios, not evidence that production is persistently net-negative. Using the archived EPA benchmark of \$46 per tonne CO₂ for 2025, in 2007 dollars, gives illustrative social values of \$77.68 billion and \$69.08 billion, or \$854.07 and \$804.31 per acre. Across the tillage scenarios, the implied per-acre values range from \$411.30 to \$1,079.17, but every endpoint 95% confidence interval includes zero. Uncertainty in the SOC predictions, soil conversion, and transfer from Minnesota to national acreage further limits these calculations.

The dynamic structural model supplies an intertemporal framework for interpreting these economic and soil relationships. Management choices affect current production, input and switching costs, and insurance outcomes, while soil-stock transitions connect those choices to future productivity and returns. This formulation makes explicit the dis-

tion between the cost of changing practices and the value of maintaining them as soil conditions adjust. It does not establish that soil accumulation causes the observed profit gains, nor does the current analysis report estimated structural parameters or completed policy counterfactuals. Those applications require identification of the model's parameters and validation of the soil transition and production relationships.

The paper thus contributes farm-level evidence on private returns to the agricultural climate-mitigation debate, places those outcomes within a dynamic soil-capital framework, and documents heterogeneity relevant to conservation adoption and policy design. The policy implication is not that a uniform subsidy will make every practice profitable, but that incentives should be evaluated against crop- and practice-specific returns, transition costs, and the time required for soil benefits to develop. Establishing the case for public support additionally requires evidence on environmental additionality, persistence, and social benefits relative to program costs. Farmers' nonpecuniary motivations, including land stewardship, also remain outside the measured financial outcomes. Sustainable conservation policy must account for both private incentives and verified public benefits, rather than equate adoption with profitability or climate mitigation.

References

- Antle, J. M., S. M. Capalbo, S. Mooney, E. T. Elliott, and K. H. Paustian (2003). Spatial Heterogeneity, Contract Design, and the Efficiency of Carbon Sequestration Policies for Agriculture. *Journal of Environmental Economics and Management* 46, 231–250.
- Berazneva, J., J. M. Conrad, D. T. Güereña, J. Lehmann, and D. Woolf (2019). Agricultural productivity and soil carbon dynamics: A bioeconomic model. *American Journal of Agricultural Economics* 101(4), 1021–1046.
- Bowman, M. and B. Pratt (2025). *Producers vary in consistent cover crop usage*. USDA Economic Research Services.
- Brown, Z. S., L. Chen, C. Cho, and R. M. Rejesus (2026). Soil capital investment and optimal cover crop choice. Selected paper prepared for presentation at the 2026 AAEA Annual Meeting.
- Callaway, B. and P. H. C. Sant’Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2), 200–230.
- Chen, Q., Y. Wang, and X. Zhu (2024). Soil organic carbon estimation using remote sensing data-driven machine learning. *PeerJ* 12, e17836.
- Crippa, M., E. Solazzo, D. Guizardi, F. Monforti-Ferrario, and F. N. T. an A Leip (2021). Food Systems are responsible for a third of global antropogenic GHG emissions. *Nature Food* 2, 198–209.
- ISRIC - World Soil Information (2026). Soilgrids: Global soil information. ISRIC.
- Jie, C., C. Jing-zhang, T. Man-zhi, and G. Zi-tong (2002). Soil degradation: A global problem endangering sustainable development. *Journal of Geographical Science* 12, 242–252.
- LaCanne, C. E. and J. G. Lundgren (2018). Regenerative Agriculture: merging farming and natural resource conservation profitably. *Peer J*.
- Lal, R. (2015). Restoring Soil Quality to Mitigate Soil Degradation. *Sustainability* 7, 5875–5895.
- Louwagie, G., S. H. Gay, and A. Burrell (2009). Addressing soil degradation in EU agriculture: relevant processes, practices and policies. *JRC Scientific and Technical Reports*.

- Malik, K. (2026). Farmer adoption and payment design under risk: Variability in soil carbon sequestration across conservation practices. NBER Working Paper 34902, National Bureau of Economic Research. Revised June 2026.
- Montgomery, D. R. (2007). Soil erosion and agricultural sustainability. *Proceedings of the National Academy of the Sciences* 104, 13268–13272.
- Newton, P., N. Civita, L. Frankel-Goldwater, K. Bartel, and C. Johns (2020). What Is Regenerative Agriculture? A Review of Scholar and Practitioner Definitions Based on Processes and Outcomes. *Frontiers in Sustainable Food Systems* 4.
- Nunes, M. R., H. M. Van Es, R. Schindelbeck, A. J. Ristow, and M. Ryan (2018). No-till and cropping system diversification improve soil health and crop yield. *Geoderma* 328, 30–43.
- Ogle, S. M., C. Alsaker, J. Baldock, M. Bernous, F. J. Breidt, B. McConkey, K. Regina, and G. G. Vazquez-Amabile (2019). Climate and Soil Characteristics Determine Where No-till management can store carbon in soils and mitigate greenhouse gas emissions. *Scientific Reports* 9.
- Pandey, S. R., X. Liu, A. Racelis, and P. Soti (2025). Estimating economic benefits of cover cropping on carbon and nitrogen storage on organic farms. *Agroecology and Sustainable Food Systems*, 1–19.
- Paustian, K., J. Lehmann, S. Ogle, D. Reay, G. P. Robertson, and P. Smith (2016). Climate-smart soils. *Nature Perspective* 532, 49–57.
- Pearl, J. (2009). *Causality*. Cambridge University Press.
- Smith, A. and A. Swanson (2025). The economic and policy challenges of climate-smart agriculture. *Review of Environmental Economics and Policy* 19(1), 90–117.
- University of Minnesota Center for Farm Financial Management (2026a). Agricultural credit analysis and finpack. University of Minnesota.
- University of Minnesota Center for Farm Financial Management (2026b). Finbin: Farm financial and benchmarking data. University of Minnesota.
- USDA National Agricultural Statistics Service (2022). Quick stats database: Census of agriculture. USDA National Agricultural Statistics Service.

- USDA Natural Resources Conservation Service (2026a). Conservation stewardship program. USDA Natural Resources Conservation Service.
- USDA Natural Resources Conservation Service (2026b). Description of the state soil geographic database (statsgo2). USDA Natural Resources Conservation Service.
- USDA Natural Resources Conservation Service (2026c). Environmental quality incentives program. USDA Natural Resources Conservation Service.
- Vermeulen, S. J., B. M. Campbell, and J. S. I. Ingram (2012). Climate Change and Food Systems. *Annual Review of Environment and Resources* 37, 195–222.
- Wood, S. and M. Bowman (2021). Large-scale farmer-led experiment demonstrates positive impact of cover crops on multiple soil health indicators. *Nature Food* 2, 97–103.
- Wu, J. (2000). Slippage Effects of the Conservation Reserve Program. *American Journal of Agricultural Economics* 82(4), 979–992.

Supplementary Appendix

The supplementary material related to this manuscript is available in the [separate supplementary appendix](#) and at <https://mcwayrm.github.io/research.html>.