

DISCUSSION PAPER SERIES

F-2026-1

Beyond average treatment effects: Heterogeneous impacts of
farmland consolidation projects in Japan using causal machine
learning

Yu Chena and Shinsaku Nakajima

October 2026

***Institute of Social Sciences
Meiji University***

KANDA SURUGADAI, CHIYODA— KU, TOKYO 101- 8301, JAPAN

Beyond average treatment effects: Heterogeneous impacts of farmland consolidation projects in Japan using causal machine learning

Yu Chen^a and Shinsaku Nakajima^{a*}

^a School of Agriculture, Meiji University, 1-1-1 Higashimita, Tama-ku, Kawasaki-shi, Kanagawa, 214-8571, Japan.

Abstract

Introduction: Identifying which rural communities benefit most from farmland consolidation projects is crucial for effective land policies in heterogeneous rural settings. **Methods:** This study provides a nationwide evaluation of farmland consolidation projects in Japan using rural community-level panel data from 2000 to 2010. We combined propensity score matching with difference-in-differences to estimate the average treatment effects and address non-random project placement. We then applied a causal forest model to uncover systematic heterogeneity in policy impacts across communities. **Results:** Farmland consolidation projects increase land rental activity, reduces farmland abandonment, and promotes the expansion of large-scale farms. However, their magnitude varies substantially across local conditions. The elderly-farmer share was the strongest predictor of treatment heterogeneity. Communities with higher shares exhibited larger increases in rental activity and greater reductions in abandoned farmland. **Conclusion:** Prioritizing communities with higher predicted treatment effects yields greater gains than uniform implementation.

JEL Classification: Q15, R14, R52

Keywords: farmland consolidation projects, causal forest model, agricultural structural change, land rental market, policy targeting, rural Japan

* Corresponding author. Tel: +81 (44) 934-7129; Email: anakajim@meiji.ac.jp

1 Introduction

Governments worldwide invest substantial public resources in farmland consolidation and land improvement programmes to promote sustainable land use, improve agricultural productivity, and support rural development. Yet ensuring that these investments generate lasting benefits has become increasingly challenging as demographic decline, population aging, and persistent land fragmentation continue to reshape rural economies. These structural changes constrain mechanization, raise production and transaction costs, and accelerate farmland abandonment, thereby weakening the efficiency and sustainability of agricultural production (Van Dijk, 2003; Hartvigsen, 2015). Consequently, improving the effectiveness of farmland consolidation policies has become an increasingly important priority for governments seeking to enhance agricultural productivity, promote sustainable land use, and improve the efficiency of public investment (OECD, 2019).

These challenges are particularly acute in Japan, where rapid population aging, the continuing decline in the number of farm households, and the predominance of small and fragmented landholdings have fundamentally reshaped agricultural production and land use over the past several decades (OECD, 2019; Otsuka et al., 2016). As increasing numbers of farmers retire without successors, more farmland is abandoned or transferred through rental markets, making the efficient reallocation of farmland an increasingly important policy challenge. In response, the Japanese government has continuously implemented farmland consolidation projects to improve agricultural production conditions and facilitate structural adjustment through land readjustment and agricultural infrastructure development (Akino, 1979). Despite decades of public investment, however, whether these projects have generated their intended structural benefits remains an important empirical and policy question.

Against this background, this study evaluates whether farmland consolidation projects have achieved their intended structural objectives in rural Japan. Specifically, it examines the impacts of farmland consolidation projects on farmland rental, farmland abandonment, and farm enlargement, three key indicators of agricultural structural transformation. In addition, the study investigates whether these impacts vary systematically across rural communities and identifies the demographic, structural, and geographic characteristics associated with stronger policy effects. By addressing these

objectives, the study seeks to provide evidence for improving the effectiveness and targeting of future farmland consolidation policies under conditions of demographic decline and increasing fiscal constraints.

To address these research objectives, this study constructs a nationwide panel dataset of rural communities in Japan covering the period from 2000 to 2010. The analysis focuses on three outcomes that capture key dimensions of agricultural structural transformation: rented-in farmland, abandoned farmland, and the share of farms operating at least five hectares. By exploiting substantial variation in the implementation of farmland consolidation projects across communities, the dataset enables a nationwide assessment of both the structural impacts of the policy and the extent to which these impacts differ across local contexts. This nationwide perspective enables a comprehensive assessment of farmland consolidation beyond the localized case studies that have dominated the existing literature.

The empirical analysis proceeds in three stages. First, difference-in-differences (DID) models are employed to estimate the average effects of farmland consolidation projects by comparing changes in outcomes between treated and untreated communities over time. Second, propensity score matching combined with DID (PSM-DID) is used to reduce potential selection bias arising from the non-random placement of projects. Finally, a causal forest model is applied to estimate heterogeneous treatment effects and identify the demographic, structural, and geographic characteristics associated with larger or smaller policy impacts. This sequential empirical framework enables the study to evaluate both average and heterogeneous policy effects, thereby providing a more comprehensive assessment of farmland consolidation projects across rural communities.

The empirical findings indicate that farmland consolidation projects facilitate agricultural structural transformation in rural Japan, and the estimated effects are robust across alternative identification strategies. The baseline DID estimates indicate that farmland consolidation increases the share of rented-in farmland by 1.62 percentage points, reduces the share of abandoned farmland by 0.47 percentage points, and increases the share of farms operating at least five hectares by 1.35 percentage points. Furthermore, the causal forest analysis reveals substantial heterogeneity in policy impacts across rural communities. In particular, communities with higher shares of elderly farmers experience larger increases in farmland rental and larger reductions in

farmland abandonment. These findings suggest that the effectiveness of farmland consolidation depends not only on improvements in agricultural infrastructure but also on local demographic conditions and the functioning of farmland rental markets.

Existing evaluations of farmland consolidation projects have primarily focused on improvements in agricultural production conditions, farm productivity, and rural community outcomes. Early studies emphasized the role of farmland consolidation in improving agricultural infrastructure and cultivation conditions (Akino, 1979), while subsequent research documented positive impacts on agricultural productivity (Arimoto, 2011). More recently, attention has expanded to broader community-level outcomes, including social capital and collective action in rural areas (Takayama and Nakatani, 2018). While these studies have substantially advanced our understanding of the benefits of farmland consolidation, much less attention has been devoted to its role in facilitating farmland rental, reducing farmland abandonment, and promoting farm enlargement. By shifting the focus from production-oriented outcomes to land reallocation and agricultural structural transformation, this study extends the scope of existing evaluations of farmland consolidation projects.

Agricultural structural transformation has long been recognized as a key pathway for improving agricultural productivity and rural development, with previous studies emphasizing the roles of technological change, farm size growth, and structural change in agriculture (Otsuka et al., 2016; Adamopoulos and Restuccia, 2014). Although land institutions and land markets are widely recognized as important drivers of structural transformation (Deininger and Feder, 2001), relatively little empirical evidence is available on how farmland consolidation projects contribute to this process by facilitating farmland reallocation in aging rural communities. By jointly examining farmland rental, farmland abandonment, and farm enlargement, this study provides new evidence on how farmland consolidation contributes to agricultural structural transformation. In doing so, it demonstrates that the effectiveness of farmland consolidation extends beyond improvements in production conditions to support more efficient land allocation and structural adjustment in rural Japan.

Most empirical evaluations of farmland consolidation projects rely on conventional econometric approaches that estimate average treatment effects. Although these methods provide valuable evidence on overall policy effectiveness, they offer

limited insights into how treatment effects vary across rural communities with different demographic, structural, and geographic characteristics. Recent advances in causal inference have demonstrated that machine-learning approaches can improve policy evaluation by estimating heterogeneous treatment effects and identifying the conditions under which policies are most effective (Wager and Athey, 2018; Athey et al., 2019). Building on this literature, this study integrates quasi-experimental methods with a causal forest model to move beyond average treatment effects and systematically identify which rural communities benefit most from farmland consolidation projects. This methodological framework provides a more informative basis for evidence-based policy targeting under increasingly constrained public budgets.

The remainder of this paper is organized as follows. Section 2 describes the data, variable construction, and empirical strategy. Section 3 presents the empirical results, including both average and heterogeneous treatment effects. Section 4 discusses the implications of the findings for farmland consolidation policy and agricultural structural transformation. Section 5 concludes.

2 Materials and Methods

This study evaluates the impact of farmland consolidation projects on agricultural structural transformation in Japanese rural communities by integrating conventional quasi-experimental methods with a causal machine-learning approach. First, we estimated average treatment effects using DID and PSM-DID. We then applied a causal forest model to the matched sample obtained from the propensity score matching procedure to uncover systematic heterogeneity in treatment effects across rural communities. This integrated framework allows us to address two complementary questions: whether farmland consolidation projects generate positive average impacts and whether these impacts differ systematically across rural communities.

This combined empirical strategy is particularly important in Japan, where rural communities differ substantially in demographic structure, geographic conditions, and land market environments. While DID and PSM-DID provide credible estimates of average policy impacts, they are less suited to capturing high-dimensional heterogeneity and nonlinear interactions across communities. To address this limitation, the causal

forest approach was applied to the matched sample generated by PSM, allowing treatment effects to vary flexibly across local structural conditions while maintaining comparability between treated and untreated communities.

2.1 Data and Variables

This study employed nationwide panel data at the rural community level for the years 2000 and 2010, which was constructed by integrating multiple official datasets compiled by Japan's Ministry of Agriculture, Forestry and Fisheries (MAFF, 2010). The primary data sources were the Census of Agriculture and Forestry and the Detailed Survey of Agricultural Infrastructure Information.

The Census of Agriculture and Forestry provides detailed information on agricultural structure, farmland use, labor composition, farm size distribution, and land use conditions at the rural community level. Recently, the number of policy evaluation studies using the Census of Agriculture and Forestry has increased, including those by Ito (2022) and Takayama et al. (2021). The infrastructure survey recorded the implementation and characteristics of land improvement projects, including farmland consolidation projects. The use of nationwide community-level data represents an advantage because farmland consolidation projects are implemented collectively at the rural community level rather than at the individual farm level. By merging these datasets using harmonized rural community identifiers, we constructed a panel dataset that captured structural changes before and after farmland consolidation over a ten-year period. The balanced panel was used for the baseline DID analysis, while matched subsamples were subsequently constructed for the PSM-DID and causal forest analyses.

This analysis focused on rural communities that could be tracked consistently across census waves. Communities affected by administrative boundary changes or lacking consistent identifiers were excluded to ensure comparability over time. Since farmland consolidation projects in Japan are implemented through collective planning, coordinated land readjustment, and community-level infrastructure investment, rural communities constituted an appropriate unit of analysis.

Treatment Variable

The treatment variable captures substantial progress in farmland consolidation between 2000 and 2010. The treatment indicator is constructed based on changes in the share of farmland plots larger than 30 ares (0.3 ha), an indicator widely used to evaluate

farmland consolidation in Japan (Arimoto, 2011; Takayama and Nakatani, 2018). Specifically, a rural community is classified as treated if the increase in the share of farmland plots larger than 30 ares ($\Delta L30$) is at least 10 percentage points between 2000 and 2010. The treatment indicator is defined as follows:

$$Treat_i = \begin{cases} 1, & \text{if } \Delta L30 \geq 10 \\ 0, & \text{otherwise} \end{cases}$$

where $\Delta L30_i$ denotes the change in the share of farmland plots larger than 30 ares between 2000 and 2010, and $L30$ represents the share of farmland plots exceeding 30 ares in rural community i . A threshold of 10 percentage points is adopted to distinguish communities that experienced substantial progress in farmland consolidation from those with only limited changes during the study period.

Outcome Variables

The empirical analysis focuses on three outcome variables that capture different dimensions of agricultural structural transformation. The first outcome variable is the share of rented-in farmland ($rent_in$), which measures the extent of farmland reallocation through rental markets. The second outcome variable is the share of abandoned farmland ($abandoned$), defined as the proportion of farmland left uncultivated. The third outcome variable is the share of farms operating at least 5 ha ($farm_5ha$), which captures farm size expansion. Together, these three outcome variables capture complementary dimensions of agricultural structural transformation and provide a comprehensive assessment of the impacts of farmland consolidation projects.

Covariates

Following Arimoto (2011) and Takayama et al. (2018), the empirical analysis includes a broad set of pretreatment covariates measured in 2000 to account for demographic, geographic, institutional, and agricultural characteristics that may jointly influence farmland consolidation project placement and subsequent agricultural outcomes. Because these covariates are measured prior to treatment, they represent baseline community characteristics and are not affected by farmland consolidation projects. The covariates are used in the propensity score matching procedure to improve the comparability of treated and untreated communities and are subsequently incorporated into the causal forest model to estimate heterogeneous treatment effects.

The demographic variables included the share of elderly farmers and indicators related to agricultural labor composition. Geographic variables captured agricultural area type, terrain conditions, accessibility, and distance to densely inhabited districts. Institutional land-use variables included indicators for agricultural promotion areas and farmland zoning conditions. Baseline agricultural structure variables included farm size composition, farmland utilization patterns, and the number of farm households.

These covariates were first used in the propensity score matching procedure to construct a comparable sample of treated and untreated rural communities. The same set of covariates was subsequently incorporated into the causal forest model to estimate heterogeneous treatment effects and identify the community characteristics associated with larger or smaller policy impacts. Detailed definitions of all variables used in the empirical analysis are provided in Table A-1 in the Appendix.

The baseline DID analysis used a balanced panel of 74,180 rural communities that were consistently observed in both 2000 and 2010, corresponding to 148,360 panel observations.

For the PSM-DID analysis, observations outside the region of common support were excluded prior to matching to improve comparability between treated and untreated communities. The resulting matched sample comprised 54,998 panel observations.

To ensure consistency between the average treatment effect and heterogeneous treatment effect analyses, the causal forest model was estimated using the same matched sample obtained from the propensity score matching procedure. The matched sample comprised 54,998 panel observations, corresponding to 27,499 unique rural communities used in the causal forest analysis.

Descriptive Statistics

Table 1 reports baseline descriptive statistics for treated and control rural communities in the final balanced DID sample observed in both 2000 and 2010. Several systematic differences are observed between the two groups prior to treatment, indicating that farmland consolidation projects were not randomly assigned across rural communities.

In particular, treated communities exhibit higher levels of rented-in farmland and larger shares of farms operating at least 5 ha at baseline, suggesting relatively more

active land rental markets and more advanced structural adjustment even before project implementation. At the same time, treated communities exhibit lower shares of abandoned farmland, indicating comparatively better land-use conditions.

In contrast, treated communities exhibit substantially lower baseline levels of L30, implying that farmland consolidation projects are more likely to be implemented in communities with initially fragmented land structures and greater potential for subsequent consolidation. This pattern is consistent with the institutional objectives of Japanese farmland consolidation policy, which primarily targets areas facing the structural constraints associated with fragmented farmlands.

Several demographic, geographic, and institutional differences were also observed across treatment groups. Treated communities tend to exhibit higher ratios of elderly farmers, larger numbers of farm households, and slightly better accessibility conditions. In addition, treated communities are more likely to be located in agricultural promotion areas and relatively flat terrain, which are generally considered more favorable conditions for land consolidation and mechanized farming. These patterns further suggest that project placement is systematically associated with pre-existing community characteristics related to agricultural production conditions and implementation feasibility.

Overall, the baseline differences reported in Table 1 provide clear evidence of nonrandom project placement and strongly motivate the use of quasi-experimental approaches to identify the causal effects of farmland consolidation projects. These patterns also suggest that farmland consolidation projects were more likely to be implemented in communities with more favorable demographic, geographic, and agricultural conditions, highlighting the importance of addressing selection bias in the empirical analysis.

Table 1. Baseline Descriptive Statistics (Baseline: Year 2000)

Variable	Full Sample	Treated Communities	Control Communities	Difference
Panel A. Outcome Variables				
Rent-in farmland share (%)	14.406	15.580	14.016	1.564***
Abandoned farmland share (%)	6.582	4.086	6.655	-2.569***

Variable	Full Sample	Treated Communities	Control Communities	Difference
Farms ≥ 5 ha (%)	2.598	3.387	2.121	1.266***
Panel B. Treatment-related Variables				
Share of plots ≥ 30 ares (L30)	0.326	0.111	0.333	-0.222***
Panel C. Demographic and Agricultural Structure Variables				
Elderly farmer ratio	0.567	0.582	0.566	0.016***
Number of farm households	22.271	24.831	22.104	2.727***
Panel D. Geographic and Institutional Variables				
Urbanization promotion area	0.037	0.042	0.037	0.005***
Urbanization control area	0.173	0.188	0.172	0.016***
Nondesignated planning area	0.300	0.276	0.301	-0.025***
Distance to densely inhabited district: <30 min	0.707	0.681	0.709	-0.028***
Distance to densely inhabited district: 30–60 min	0.248	0.271	0.247	0.024***
Distance to densely inhabited district: 60–90 min	0.038	0.041	0.038	0.003**
Distance to densely inhabited district: >90 min	0.007	0.007	0.006	0.001
Agricultural promotion area	0.984	0.989	0.984	0.005***
Agricultural land zone	0.911	0.926	0.910	0.016***
Flat terrain	0.455	0.492	0.454	0.038***
Gentle slope	0.376	0.351	0.377	-0.026***
Steep slope	0.169	0.157	0.169	-0.012***

2.2 Identification Strategy

Farmland consolidation projects are not randomly assigned to rural communities. Project implementation depends on local institutional capacity, geographic conditions, expected returns from land improvement, and the ability of communities to coordinate collective decisions. Consequently, simple comparisons between treated and control communities may produce biased estimates if treatment assignment is correlated with pre-existing structural conditions and outcome dynamics.

The primary empirical challenge is that communities receiving farmland consolidation projects may differ systematically from those that do not, both in observable characteristics and in unobservable factors related to agricultural development. To address this issue, this study combined the DID and PSM-DID approaches to estimate the average treatment effects, while accounting for systematic differences in pre-treatment characteristics. DID controls for time-invariant unobserved heterogeneity through community fixed effects, whereas PSM-DID improves comparability between treated and control communities by constructing a matched sample based on observed baseline characteristics.

However, average treatment effect estimators impose a homogeneous treatment effect assumption and may conceal substantial variation in policy impacts across rural communities. This limitation is pertinent in the Japanese context, where demographic structure, land market conditions, and geographic constraints differ markedly across regions. To relax the homogeneous treatment effect assumption, the analysis applies a causal forest framework that flexibly estimates heterogeneous treatment effects across communities.

2.3 Difference-in-Differences Model

The primary empirical challenge in evaluating farmland consolidation projects is that project implementation is unlikely to be randomly assigned across rural communities. Communities receiving farmland consolidation projects may differ systematically from those that do not in terms of demographic composition, agricultural conditions, institutional capacity, and economic development. Consequently, simple comparisons between participating and non-participating communities may produce biased estimates of policy impacts. This challenge is common in observational policy evaluation settings and motivates the use of quasi-experimental approaches designed to address selection bias (Imbens and Wooldridge, 2009).

To address this issue, this study combines Difference-in-Differences (DID) and Propensity Score Matching Difference-in-Differences (PSM-DID) approaches to estimate the average effects of farmland consolidation projects. Combining multiple identification strategies is increasingly recommended in policy evaluation studies to strengthen causal interpretation when treatment assignment is nonrandom (Wing et al.,

2018). Subsequently, a causal forest framework is employed to investigate heterogeneous treatment effects across rural communities.

For the empirical analysis, let $Outcome_{it}$ denote the outcome variable for rural community i in year t . The analysis considers three outcomes: (i) the share of rented-in farmland, (ii) the share of abandoned farmland, and (iii) the share of farms operating at least 5 hectares. These outcomes capture three key dimensions of agricultural structural transformation: land reallocation, land-use sustainability, and farm size expansion.

The DID framework compares changes in outcomes between treated and untreated communities before and after substantial farmland consolidation. The baseline specification is:

$$Outcome_{it} = \alpha + \beta(Treat_i \times Post_t) + \mu_i + \lambda_t + \varepsilon_{it},$$

where $Outcome_{it}$ denotes one of the three outcome variables for rural community i in year t : the share of rented-in farmland, the share of abandoned farmland, or the share of farms operating at least 5 hectares. $Treat_i$ is a treatment indicator equal to one for communities experiencing substantial farmland consolidation and zero otherwise. $Post_t$ is a post-treatment indicator equal to one for 2010 and zero for 2000. μ_i represents community fixed effects, which control for time-invariant community characteristics, and λ_t denotes year fixed effects, which capture common shocks affecting all communities. ε_{it} is an idiosyncratic error term. The coefficient β captures the average treatment effect of farmland consolidation projects.

The validity of the DID approach relies on the parallel trends assumption, which requires treated and untreated communities to exhibit similar outcome trends in the absence of treatment. Although this assumption cannot be directly tested with only two census waves, the DID estimates are complemented by the PSM-DID analysis, which improves comparability between treated and control communities based on rich pretreatment characteristics.

Because project participation may be systematically related to observable community characteristics, propensity score matching is employed prior to DID estimation. Following Rosenbaum and Rubin (1983), communities are matched on the basis of their estimated probability of receiving treatment using pretreatment covariates measured in 2000. The DID model is then estimated on the matched sample to reduce observable selection bias and improve the credibility of causal inference.

While DID and PSM-DID provide estimates of average treatment effects, they do not reveal whether policy impacts differ across rural communities. This limitation motivates the heterogeneous treatment effect analysis presented in Section 2.5, where a causal forest framework is used to identify the demographic and structural conditions under which farmland consolidation projects are most effective.

2.4 *Propensity Score Matching DID*

As farmland consolidation projects are not randomly assigned across rural communities, simple comparisons between treated and untreated communities may produce biased estimates of policy effects. In Japan, the implementation of farmland consolidation projects depends not only on agricultural conditions, but also on local institutional capacity, collective decision-making, and geographic feasibility. Rural communities with stronger organizational capacities, more favorable terrain conditions, or greater expected returns from land improvement may, therefore, be more likely to adopt consolidation projects. As a result, treatment assignment is potentially endogenous to pre-existing community characteristics.

To address this issue, we employed PSM-DID. The purpose of this approach was to construct a counterfactual comparison group of untreated rural communities that is observationally similar to treated communities prior to project implementation. By combining matching with DID, the empirical strategy reduces selection bias arising from observable pre-treatment differences while simultaneously controlling for time-invariant unobserved heterogeneity.

Following Rosenbaum and Rubin (1983), propensity scores were first estimated using pre-treatment characteristics measured in 2000, which plausibly influenced both project placement and post-treatment agricultural outcomes. These covariates included demographic structure, geographic conditions, land-use institutional variables, and baseline agricultural characteristics. Specifically, demographic variables capture the local labor structure and population aging, geographic variables reflect accessibility and terrain conditions, institutional variables capture land-use regulation and agricultural zoning, and agricultural structure variables describe baseline farming conditions and land-use patterns.

The propensity score is estimated using a probit model:

$$P(Treat_i = 1 | X_i) = \Phi(X_i\gamma),$$

where $Treat_i$ denotes the treatment indicator, X_i represents the vector of pre-treatment covariates, and $\Phi(\cdot)$ is the standard normal cumulative distribution function.

Table 2 reports the probit estimates as average marginal effects. The results indicate that farmland consolidation project placement is systematically associated with pre-treatment demographic, geographic, and institutional characteristics, confirming the presence of nonrandom project placement across rural communities.

Table 2. Determinants of Farmland Consolidation Project Placement

Variable	Marginal Effects S.E.
Agricultural area (urban)	-0.0214*** (0.0048)
Agricultural area (intermediate)	-0.0187*** (0.0023)
Agricultural area (mountainous)	-0.0295*** (0.0027)
Gradient (gentle)	-0.0184*** (0.0020)
Gradient (steep)	-0.0284*** (0.0024)
Distance to densely inhabited district (30–60 min)	0.0050** (0.0022)
Distance to densely inhabited district (60–90 min)	-0.0024 (0.0048)
Distance to densely inhabited district (>90 min)	-0.0004 (0.0107)
Urbanization promotion area	-0.0147*** (0.0047)
Urbanization control area	0.0005 (0.0024)
Not designated area	0.0006 (0.0019)
Agricultural promotion area	-0.0316*** (0.0067)
Agricultural promotion farmland area	0.0178*** (0.0036)
Elderly farmer ratio	-0.0138*** (0.0050)
Number of farm households	0.0004*** (0.0000)
Observations: 53,421	
LR chi2: 416.59	
Log likelihood: -8294.52	
Pseudo R ² (McFadden): 0.024	

Note: Average marginal effects from a probit model estimated using baseline (year 2000) rural community characteristics. Robust standard errors (HC1) are reported in parentheses. The omitted categories were agricultural area (flat/plain), gradient (flat), and distance to densely inhabited district below 30 minutes. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Several patterns were noteworthy. First, rural communities with a larger number of farm households are more likely to adopt farmland consolidation projects. This result

is consistent with the institutional characteristics of Japanese farmland consolidation projects, which require collective coordination and consensus building among local farmers. Communities with larger farming populations may therefore possess a stronger organizational capacity to initiate and implement consolidation projects.

Second, project adoption differs systematically across geographic conditions. Rural communities located on gentle and steep slopes are less likely to adopt consolidation projects than those in flat areas, reflecting the higher engineering costs and physical constraints associated with land improvement in mountainous terrains. Accessibility conditions measured by the distance to densely inhabited districts also influence project placement, suggesting that both geographic feasibility and market accessibility affect the implementation of farmland consolidation projects.

Third, institutional land-use conditions play an important role in project placement. Rural communities located within agricultural land zones exhibit a higher probability of project adoption, indicating that farmland consolidation projects were more actively implemented in areas already prioritized for agricultural use and farmland preservation under Japanese land-use policy.

Finally, the ratio of elderly farmers is negatively associated with project adoption. One possible interpretation is that severely aging rural communities may face difficulties in organizing collective action and mobilizing the financial or labor resources necessary for project implementation, despite the potentially strong demand for land improvement. Overall, these findings are broadly consistent with those of previous studies, such as Arimoto (2011) and Takayama et al. (2018), which emphasize the importance of demographic structure, institutional conditions, and collective organizational capacity in shaping farmland consolidation adoption in Japan.

To assess the common support condition, Figure 1 compares the distributions of estimated propensity scores of treated and untreated communities. Although treated communities exhibited a wider dispersion of estimated propensity scores, a substantial overlap remained between the two groups, indicating that matching was feasible. Observations outside the region of common support were excluded before matching to improve comparability and reduce the influence of observations for which no comparable counterfactual existed.




Figure 1. Overlap of propensity score distribution (common support)

Treated communities were then matched to observationally similar untreated communities using nearest-neighbor matching based on the estimated propensity scores. After constructing the matched sample, the matching quality was evaluated by comparing the standardized mean differences in covariates before and after matching. Figure 2 shows that several covariates exhibited substantial imbalances before matching, whereas these imbalances were considerably reduced after matching. This result indicates that the matching procedure substantially improves the covariate balance and strengthens the comparability between treated and untreated communities.

The DID estimator is subsequently applied to the matched sample to estimate the Average Treatment Effect on the Treated (ATT). Specifically, the PSM-DID estimator compares the changes in outcomes between treated communities and their matched untreated counterparts before and after project implementation:

$$ATT = E[(Y_{i,2010}^T - Y_{i,2000}^T) - (Y_{j,2010}^C - Y_{j,2000}^C) | Treat_i = 1]$$

where $Y_{i,t}^T$ denotes the outcome of treated community i in year t , and $Y_{j,t}^C$ denotes the outcome of the matched untreated community j in year t . By differencing outcomes over time, the DID framework removes time-invariant unobserved heterogeneity, whereas the matching procedure improves comparability in observed baseline characteristics.

Standard errors are heteroskedasticity-robust and clustered at the rural community level.

As an additional robustness check, the DID and PSM-DID models were re-estimated using alternative treatment thresholds based on increases in the share of farmland plots larger than 30 ares. Specifically, treatment indicators were alternatively defined using thresholds of 5 and 15 percentage points. These results are presented in Table A-2. The estimated effects remained qualitatively consistent across alternative specifications, indicating that the main findings are not driven by the particular 10 percentage-point cutoff adopted in the baseline specification.

Overall, the PSM-DID framework provided a more credible estimate of the average causal effects of farmland consolidation projects by combining the two sources of identification. Matching reduces the selection bias associated with observable pre-treatment characteristics, whereas DID controls for time-invariant unobserved heterogeneity across rural communities. Therefore, this combined approach is particularly suitable for evaluating farmland consolidation projects whose implementation depends heavily on pre-existing demographic, institutional, and geographic conditions. Nevertheless, both DID and PSM-DID focus on average treatment effects and therefore provide limited information regarding the distribution of policy impacts across communities. This limitation motivates the heterogeneous treatment effect analysis presented in the next section.




Figure 2. Covariate Balance Before and After Propensity Score Matching

2.5 Causal Forest for Heterogeneous Effects

While the DID and PSM-DID approaches provide estimates of average treatment effects, they do not reveal whether the impacts of farmland consolidation projects differ systematically across rural communities. Such heterogeneity is likely to be particularly important in the Japanese context, where communities vary substantially in demographic composition, agricultural structure, geographic conditions, and land market environments. Consequently, understanding who benefits most from farmland consolidation projects is essential not only for explaining policy effectiveness but also for improving the targeting of future public investments.

To investigate heterogeneous treatment effects, this study employs the causal forest approach developed by Athey and Imbens (2016), Wager and Athey (2018), and Athey et al. (2019). Causal forests extend conventional random forest algorithms to estimate conditional average treatment effects (CATEs) while maintaining desirable statistical properties for causal inference. Unlike traditional regression approaches that require researchers to specify interaction terms *ex ante*, causal forests flexibly capture nonlinear relationships and high-dimensional interactions among covariates in a data-driven manner. This feature is particularly useful in the present setting because the

determinants of policy effectiveness may involve complex interactions among demographic, geographic, and agricultural characteristics that are difficult to model using conventional parametric approaches.

Following the matching procedure described above, the causal forest model is estimated using the matched sample obtained from propensity score matching. This approach ensures that treatment and control communities are comparable in terms of observed pretreatment characteristics before heterogeneous treatment effects are estimated. Consistent with the identification strategy of the PSM-DID analysis, all covariates are measured in the baseline year (2000) to avoid post-treatment bias.

Let Y_i denote the observed outcome for rural community i , W_i denote the treatment indicator, and X_i denote a vector of pre-treatment covariates. The parameter of interest is the Conditional Average Treatment Effect (CATE), defined as:

$$\tau(x) = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = x]$$

where $Y_i(1)$ and $Y_i(0)$ denote the potential outcomes with and without substantial farmland consolidation, respectively. The CATE captures how treatment effects vary systematically across rural communities with different observable characteristics.

The causal forest estimation is based on the following partially linear causal model:

$$Y_i = \mu(X_i) + \tau(X_i)W_i + \varepsilon_i$$

where $\mu(X_i)$ represents the baseline outcome function and $\tau(X_i)$ captures the treatment effect heterogeneity conditional on the observed community characteristics.

In practice, the generalized random forest framework estimates treatment effects using orthogonalized residuals based on the following local moment condition:

$$E = \left[\left(Y_i - \hat{m}(X_i) - \tau(X_i)(W_i - \hat{e}(X_i)) \right) (W_i - \hat{e}(X_i)) \middle| X_i = x \right] = 0$$

where $\hat{m}(X_i) = E[Y_i | X_i]$ denotes the estimated outcome regression, and $\hat{e}(X_i) = \Pr(W_i = 1 | X_i)$ denotes the estimated propensity score. This orthogonalization reduces sensitivity to nuisance parameter estimation and improves robustness when estimating heterogeneous treatment effects in high-dimensional settings.

Identification relies on the assumptions of conditional unconfoundedness and overlap:

$$\{Y_i(1), Y_i(0)\} \perp W_i | X_i$$

$$0 < P_r(W_i = 1|X_i) < 1$$

These assumptions imply that, conditional on observed pretreatment characteristics, treatment assignment is independent of potential outcomes and that comparable treated and control communities exist throughout the covariate space.

The set of covariates used in the causal forest analysis includes demographic characteristics, agricultural structure variables, geographic conditions, land-use regulations, and baseline outcome measures. Specifically, the model incorporates the share of elderly farmers, the number of farm households, agricultural area classifications, terrain conditions, urban planning designations, agricultural promotion area indicators, and baseline values of the outcome variables. These variables are selected because previous studies suggest that they influence both agricultural structural transformation and the effectiveness of farmland-related policies (Adamopoulos and Restuccia, 2014; Otsuka et al., 2016; Ito et al., 2016; Takayama et al., 2021).

To interpret the sources of treatment heterogeneity, the analysis proceeds in four steps. First, the distribution of estimated CATEs is examined to assess the extent of treatment heterogeneity across rural communities. Second, variable importance measures are used to identify the baseline characteristics most strongly associated with variation in treatment effects. Third, the estimated CATEs are linked to key community characteristics to explore the mechanisms underlying heterogeneous policy impacts. Particular attention is given to the share of elderly farmers because population aging is closely associated with farm exits, retirement, and the potential supply of transferable farmland in rural Japan. Following the land-market mechanism emphasized by Shogenji (1998) and Ito et al. (2016), the elderly farmer ratio is interpreted as an important proxy for local land reallocation opportunities. Finally, the estimated treatment effects are used to conduct a policy-targeting exercise that evaluates whether prioritizing communities predicted to experience larger policy gains could improve the efficiency of farmland consolidation programs relative to uniform allocation approaches.

By complementing DID and PSM-DID with the causal forest framework, this study relaxes the homogeneous treatment effect assumption and provides a more comprehensive evaluation of farmland consolidation projects across Japan's heterogeneous rural communities. The causal forest analysis not only identifies who benefits most from farmland consolidation projects but also provides insights into the

demographic and structural conditions under which such policies are likely to be most effective.

3 Results

3.1 Average Treatment Effects: DID and PSM-DID

We estimated the average effects of farmland consolidation projects using DID and PSM-DID. Table 3 presents the baseline results.

The DID estimates indicate that consolidation projects increase the share of rented-in farmland and large-scale farms, but reduce the share of abandoned farmland. This suggests that consolidation facilitates land reallocation and improves land-use efficiency, while contributing to structural adjustments toward larger and potentially more efficient farm operations.

To improve comparability between treated and untreated communities, we further estimated a PSM-DID specification based on the matched sample constructed using propensity score matching. The PSM-DID estimates remained quantitatively and qualitatively similar to the DID results, suggesting that the estimated impacts are robust to observable differences between treated and control rural communities.

Quantitatively, the PSM-DID estimates indicate that farmland consolidation projects increase the share of rented-in farmland by approximately 1.49 percentage points and reduce the share of abandoned farmland by approximately 0.48 percentage points. Although these magnitudes may appear modest, they are economically meaningful, given that the analysis is conducted at the rural community level and captures average effects across a nationwide sample. Collectively, these results indicate that consolidation promotes land reallocation and structural adjustment.

Table 3. Average Treatment Effects of Farmland Consolidation Projects

	(1) Rent-in Farmland	(2) Abandoned Farmland	(3) Large-scale Farms
Panel A. DID			
Treat × Post	0.0162*** (0.0028)	-0.0047*** (0.0011)	0.0135*** (0.0016)
Panel B. PSM-DID			
Treat × Post	0.0149*** (0.0043)	-0.0048** (0.0019)	0.0110*** (0.0022)
Observations			

DID observations	148,360	148,360	148,360
PSM-DID observations	54,998	54,998	54,998
Community FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Note: Standard errors are clustered at the rural community level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.2 Heterogeneous Treatment Effect

To examine whether the impacts of farmland consolidation projects differ across rural communities, we estimated the conditional average treatment effects (CATE) using a causal forest model. This approach allows treatment effects to vary flexibly according to local demographic, geographic, and institutional characteristics.

Figures 3 and 4 present the distributions of the estimated CATE for rented-in and abandoned farmlands, respectively. Substantial heterogeneity was observed in both outcomes. For rented-in farmland, most of the estimated treatment effects were positive, indicating that consolidation projects generally promote land rental activity across rural communities. For abandoned farmland, most of the estimated effects are negative, suggesting that consolidation projects tend to reduce farmland abandonment. Approximately 88.8% of rural communities exhibit positive treatment effects for rented-in farmland, while 86.4% exhibit negative treatment effects for abandoned farmland.




Figure 3. Distribution of Estimated CATE: Rented-in Farmland




Figure 4. Distribution of Estimated CATE: Abandoned Farmland

The substantial dispersion indicates that effects vary considerably across communities. The impacts of farmland consolidation projects are highly context-dependent and cannot be fully captured by average treatment effects. This finding suggests that average treatment effects may conceal substantial variation in policy effectiveness across communities.

The substantial heterogeneity suggests that local demographic, geographic, and institutional conditions play important roles in shaping policy effectiveness. The following analysis therefore investigates which community characteristics are most strongly associated with variations in treatment effects.

3.3 Key Drivers of Treatment Heterogeneity

To identify the baseline characteristics most strongly associated with treatment effect heterogeneity, we examined variable importance measures derived from the causal forest model. Figures 5 and 6 present the relative importance of baseline demographic, geographic, institutional, and agricultural structural characteristics in explaining variation in treatment effects across rural communities. Notably, the characteristics associated with larger treatment effects differ from those associated with project placement, suggesting that communities most likely to receive projects are not necessarily those that benefit the most.

The results indicate that the elderly farmer ratio is the most important predictor of treatment heterogeneity for both rented-in farmland and abandoned farmland. Across both outcomes, this variable exhibits substantially higher importance scores than all other covariates, suggesting that demographic structure plays a central role in shaping the effectiveness of farmland consolidation projects. This finding implies that the impacts of farmland consolidation are closely related to local land supply conditions, which are strongly influenced by the age composition of farm operators.

The number of farm households and agricultural area type also emerge as important sources of treatment heterogeneity. For rented-in farmland, the number of farm households is the second most influential predictor, followed by agricultural area type. For abandoned farmland, agricultural area type ranks second, while the number of farm households remains highly important. These results suggest that both community scale and regional agricultural conditions influence the effectiveness of farmland consolidation projects. Larger farming communities may provide greater opportunities for land reallocation and coordination among landowners, while differences across urban, intermediate, and mountainous agricultural regions may affect land market functioning and the potential for structural adjustment.

The prominence of the elderly farmer ratio is particularly informative for understanding the mechanism underlying treatment heterogeneity. Population aging is closely associated with local land supply conditions because farm retirement, declining labor capacity, and succession difficulties tend to increase the amount of farmland available for transfer. Under such circumstances, farmland consolidation projects may facilitate land transfers by improving plot accessibility, reducing operational inefficiencies, and lowering transaction costs. Consequently, communities with a higher share of elderly farmers are more likely to experience stronger impacts from farmland consolidation projects through enhanced land market adjustment.




Figure 5. Key Drivers of Treatment Heterogeneity: Rented-in Farmland




Figure 6. Key Drivers of Treatment Heterogeneity: Abandoned Farmland

Other variables, including accessibility conditions, institutional land-use characteristics, and farmland-related factors, also contribute to treatment heterogeneity, although their relative importance is considerably smaller. Overall, the results suggest that heterogeneous policy impacts are jointly shaped by demographic structure, community scale, geographic conditions, and local land-use institutions.

Variable importance does not by itself establish causality, but identifies characteristics associated with treatment-effect variation. The consistency of the elderly farmer ratio as the most important predictor across both outcomes suggests that population aging represents a fundamental source of treatment heterogeneity rather than

an outcome-specific correlation. Given its theoretical relevance and empirical importance, the following analysis focuses on the elderly farmer ratio as a proxy for local land supply conditions. In the Japanese agricultural context, older farmers are more likely to reduce or cease farming activities because of retirement, labor constraints, and succession difficulties. As a result, communities with a higher share of elderly farmers tend to have a greater potential supply of farmland available for transfer through rental markets. Previous studies have highlighted the role of aging farm populations in facilitating farmland reallocation and structural adjustment through land rental transactions (Shogenji, 1998; Arimoto, 2011). Therefore, the elderly farmer ratio serves as a theoretically meaningful indicator of local land supply conditions in the subsequent heterogeneity analysis.

3.4 Demographic Structure and Treatment Heterogeneity

To better understand the sources of treatment heterogeneity identified by the causal forest model, this section examines how estimated treatment effects vary with local demographic characteristics, particularly the share of elderly farmers. Figures 7 and 8 indicate that treatment effects vary systematically with the share of elderly farmers across rural communities. In communities with higher shares of elderly farmers, farmland consolidation projects generate larger increases in farmland rental activity and larger reductions in farmland abandonment.

For rented-in farmland, the estimated treatment effects increase steadily with the share of elderly farmers, indicating that farmland consolidation projects generate larger increases in land rental activity in older rural communities. For abandoned farmland, the estimated treatment effects become increasingly negative as the elderly farmer ratio rises, implying larger reductions in farmland abandonment in communities experiencing greater population aging.




Figure 7. Treatment Heterogeneity by Elderly Ratio: Rented-in Farmland




Figure 8. Treatment Heterogeneity by Elderly Ratio: Abandoned Farmland

While the overall relationship is driven primarily by demographic structure, the magnitude of the effects also varies across geographical conditions. In Figure 7, the positive effect on rented-in farmland is consistently stronger in hilly and mountainous areas than in flat areas, suggesting that the positive effects of farmland consolidation on land rental activity are stronger in areas facing greater geographic constraints. In contrast, Figure 8 indicates that reductions in abandoned farmland are particularly pronounced in flat areas with higher levels of aging.

These findings suggest that local demographic and land market conditions are closely associated with the effectiveness of farmland consolidation projects. One possible explanation is that retirement and declining labor capacity in aging communities increase the amount of farmland available for transfer. Improvements in plot conditions and accessibility associated with farmland consolidation may make land transfers easier and contribute to more active land use.

Therefore, the findings are consistent with the view that farmland consolidation projects are more effective in communities where demographic conditions create greater opportunities for land reallocation.

To formally examine the relationship suggested by the graphical evidence, we conducted a slope test by regressing the estimated treatment effects on the share of elderly farmers:

$$\hat{\tau}_i = \alpha + \beta \cdot Elderly_i + \epsilon_i$$

where $\hat{\tau}_i$ denotes the estimated CATE for rural community i and $Elderly_i$ is the share of elderly farmers.

Table 4 reports the results. The estimated coefficients on the elderly farmer ratio are positive and statistically significant for rented-in farmland and negative and statistically significant for abandoned farmland. These findings confirm that demographic structure is systematically associated with variation in treatment effects across rural communities. The consistency between the graphical evidence and the regression-based analysis further suggests that the observed patterns reflect systematic heterogeneity rather than random variation. Together with the variable importance analysis, these results provide consistent evidence that population aging is a key source of treatment heterogeneity in farmland consolidation projects.

Table 4. Slope Test of Treatment Heterogeneity

	(1) Rent-in CATE	(2) Abandoned CATE
Share of elderly farmers	0.0321*** (0.0006)	-0.0194*** (0.0005)
Constant	-0.0002 (0.0004)	-0.0041*** (0.0003)
Observations	27,499	27,499
R-squared	0.103	0.037

Note: Dependent variables are estimated CATEs obtained from the causal forest model. Robust standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.5 Policy Targeting Implications

Finally, we examine the policy implications of treatment heterogeneity through a targeted analysis. Specifically, we evaluated whether directing farmland consolidation projects toward rural communities with higher predicted treatment effects could improve policy outcomes relative to uniform implementation.

Figures 9 and 10 show that targeting high-impact communities yields substantially larger gains than non-targeted allocations. In particular, prioritizing communities with higher predicted treatment effects produces larger increases in rented-in farmland and larger reductions in abandoned farmland compared to uniform project implementation.




Figure 9. Targeting Gain Curve: Rented-in Farmland

The gain curves indicate that a relatively small share of communities accounts for a disproportionately large share of potential policy gains, suggesting that targeting based on predicted treatment effects could improve the marginal returns to public investment. Under limited fiscal resources, incorporating information on treatment heterogeneity

into project selection may improve policy efficiency by directing resources toward communities in which consolidation is expected to generate larger impacts.




Figure 10. Targeting Gain Curve: Abandoned Farmland

More broadly, the targeting analysis demonstrates that treatment heterogeneity is not only an empirical feature of policy impacts, but also a potentially important source of information for policy design. Rather than adopting a uniform nationwide allocation strategy, policymakers should improve the effectiveness of farmland consolidation programs by considering local demographics and land market conditions.

Taken together, the results demonstrate that farmland consolidation projects generate positive average effects while also exhibiting substantial treatment heterogeneity. The findings further suggest that demographic structure, particularly population aging, plays a central role in shaping policy effectiveness and should therefore be considered when evaluating and prioritizing farmland consolidation policies.

4 Discussion

The results show that farmland consolidation projects increase farmland rental activity, reduce farmland abandonment, and promote the expansion of large-scale farms. Furthermore, the findings reveal substantial heterogeneity in policy impacts across rural communities, suggesting that the effectiveness of farmland consolidation depends not only on physical land improvements but also on local demographic and structural conditions.

The positive average effects observed in this study are consistent with the institutional objectives of farmland consolidation policies in Japan. By enlarging plot size and improving agricultural infrastructure, consolidation projects reduce operational constraints and facilitate more efficient agricultural production. Improved accessibility and cultivation conditions may also lower the transaction costs associated with farmland transfers, thereby promoting land reallocation through rental markets. These factors may help explain the observed increase in rented-in farmland and expansion of large-scale farming operations. At the same time, by improving cultivation conditions and facilitating the transfer of land to active farmers, farmland consolidation projects can reduce the likelihood that farmland remains idle or becomes abandoned. These findings are broadly consistent with previous studies emphasizing the role of farmland consolidation in promoting land mobility and structural adjustment in Japanese agriculture (Kunimitsu, 2005; Arimoto, 2011; Ito et al., 2016). The consistency of the DID and PSM-DID estimates further strengthens the credibility of these findings by suggesting that the estimated impacts are unlikely to be driven solely by observable selection bias.

While average effects are important, the most notable finding of this study is that policy impacts vary systematically across rural communities. The causal forest analysis identifies the elderly farmer ratio as the strongest predictor of treatment heterogeneity, indicating that demographic structure plays a central role in shaping policy effectiveness. In particular, the effects of farmland consolidation projects are substantially larger in communities with a higher share of elderly farmers.

A possible explanation is that aging rural communities tend to have greater potential farmland available for transfer with increasing farm retirement and succession difficulties, as well as declining labor capacity. Under such conditions, farmland

consolidation projects may facilitate land reallocation by improving plot accessibility, reducing operational constraints, and lowering transaction costs associated with farmland transfers. Because individual land transactions are not observed, however, the heterogeneous treatment effects identified in this study are consistent with the view that farmland consolidation is more effective where demographic conditions create greater opportunities for land reallocation.

These findings are broadly consistent with the broader structural adjustment literature in Japanese agricultural economics. Shogenji (1998) argued that agricultural transformation depends not only on technological improvements but also on the transfer of farmland from less active to more productive operators. Similarly, Ito et al. (2016) highlighted the importance of land transactions and rental institutions in preventing farmland abandonment. The present study extends this literature by showing that the effectiveness of farmland consolidation projects depends not only on physical land improvements but also on local demographic conditions that influence the availability of land for reallocation.

The findings also contribute to a growing literature on heterogeneous policy effects. Most previous evaluations of farmland consolidation projects have focused on average treatment effects and implicitly assumed that policy impacts are relatively homogeneous across regions. The present results suggest that such an assumption may obscure important differences in policy effectiveness. Communities with similar projects may experience substantially different outcomes because demographic structures, land market conditions, and agricultural organization differ across locations. These findings underscore the importance of incorporating heterogeneity into the evaluation of rural development policies and highlight the value of causal machine-learning approaches for uncovering systematic variation in policy impacts.

The presence of substantial treatment heterogeneity has important implications for policy design. The targeting analysis suggests that considerable efficiency gains may be achieved by prioritizing communities predicted to experience larger treatment effects. This finding is particularly relevant in the context of fiscal constraints and population decline, where governments face increasing pressure to allocate public resources more efficiently. Rather than assuming that farmland consolidation projects generate identical

benefits everywhere, policymakers may improve policy performance by considering local demographic and structural conditions when determining project priorities.

At the same time, the results highlight an important trade-off between efficiency and equity. Prioritizing communities with the largest expected gains may maximize aggregate policy effectiveness but could also lead to an unequal distribution of public investment across regions. Consequently, policy targeting should not be interpreted as a purely technical optimization exercise. Instead, it should be viewed as one input into broader policy decisions that must balance efficiency objectives with concerns regarding regional equity, rural sustainability, and social inclusion.

Several limitations of this study should nevertheless be acknowledged. First, although the empirical strategy addresses selection bias using quasi-experimental methods, unobserved community-level factors may still influence both project implementation and agricultural outcomes. Second, the analysis relies on benchmark-year observations and therefore cannot fully capture short-term adjustment dynamics following project implementation. Third, although the evidence suggests that farmland consolidation is more effective in communities with greater opportunities for land reallocation, the analysis does not directly observe individual land transactions or dynamic interactions within local rental markets. Future research using detailed micro-level transaction data and higher-frequency panel datasets would provide more direct evidence regarding the channels through which farmland consolidation projects influence agricultural structural transformation.

Overall, the results suggest that farmland consolidation projects contribute not only to improving agricultural production conditions but also to facilitating land reallocation and supporting structural transformation in aging rural societies. More importantly, the findings demonstrate that policy effectiveness depends critically on local demographic conditions. Recognizing and incorporating such heterogeneity may be essential for designing more effective and sustainable rural development policies in the future. Although this study focuses on Japan, the analytical framework developed here may also provide useful insights for evaluating land consolidation policies in other aging agricultural regions facing similar challenges of demographic decline and farmland fragmentation.

5 Conclusions

This study examined the impacts of farmland consolidation projects on agricultural structural adjustment in rural Japan using nationwide rural community-level panel data covering the period from 2000 to 2010. By combining difference-in-differences (DID), propensity score matching combined with difference-in-differences (PSM-DID), and causal forest methods, the analysis evaluated both average treatment effects and heterogeneous treatment effects across rural communities.

The results indicate that farmland consolidation projects significantly increase farmland rental activity, reduce farmland abandonment, and promote the expansion of large-scale farming. These findings suggest that farmland consolidation contributes not only to improving agricultural production conditions but also to facilitating land reallocation and supporting structural adjustment in rural areas.

More importantly, the analysis reveals substantial heterogeneity in policy impacts across rural communities. The causal forest results identify the share of elderly farmers as the most important predictor of treatment heterogeneity, indicating that demographic structure plays a central role in shaping policy effectiveness. Communities with larger proportions of elderly farmers experience greater increases in farmland rental activity and larger reductions in farmland abandonment following farmland consolidation. These findings are consistent with the view that the effectiveness of farmland consolidation depends not only on physical land improvements but also on local demographic and land market conditions.

This study contributes to the literature by linking farmland consolidation to land reallocation and by identifying systematic heterogeneity in policy effects. The findings also have important policy implications. The results suggest that farmland consolidation projects are not equally effective across all rural communities. While project implementation tends to be associated with favorable agricultural and institutional conditions, larger treatment effects are observed in communities with higher shares of elderly farmers. This implies that the communities most likely to receive projects are not necessarily those that benefit the most. Under conditions of limited fiscal resources and continuing population decline, incorporating information on treatment heterogeneity into project selection may improve the effectiveness of farmland consolidation policies.

Several limitations should be acknowledged. First, although the empirical strategy addresses selection bias using quasi-experimental methods, unobserved community-level factors may still influence both project implementation and agricultural outcomes. Second, the analysis relies on benchmark-year observations and therefore cannot fully capture short-term adjustment dynamics following project implementation. Third, although the findings are consistent with a land market explanation, the analysis does not directly observe individual land transactions. Future research using detailed farm-level or transaction-level data would provide a more direct understanding of the mechanisms through which farmland consolidation influences rural structural change.

Overall, the results suggest that farmland consolidation projects play an important role in facilitating land reallocation and supporting agricultural adaptation under conditions of population aging and rural decline. More broadly, the study demonstrates that understanding who benefits from rural development policies can be as important as estimating their average effects. Incorporating treatment heterogeneity into policy evaluation may therefore contribute to the design of more effective and sustainable rural development policies in aging societies.

References

- Akino, M., 1979. Land infrastructure improvement in agricultural development: The Japanese case, 1900–1965. *Economic Development and Cultural Change* 28(1), 97–117.
- Adamopoulos, T., Restuccia, D., 2014. The size distribution of farms and international productivity differences. *American Economic Review* 104(6), 1667–1697.
- Arita, H., Kimura, K., 2003. Land consolidation of paddy field. Association of Agriculture and Forestry Statistics, Tokyo.
- Arimoto, Y., 2011. The impact of farmland readjustment and consolidation on structural adjustment: The case of Niigata, Japan. Center for Economic Institutions Working Paper Series 2011-3, Hitotsubashi University.
- Athey, S., Imbens, G.W., 2016. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences* 113(27), 7353–7360.
- Athey, S., Tibshirani, J., Wager, S., 2019. Generalized random forests. *Annals of Statistics* 47(2), 1148–1178.
- Deininger, K., Feder, G., 2001. Land institutions and land markets. In: Gardner, B., Raussier, G. (Eds.), *Handbook of Agricultural Economics*, Vol. 1. Elsevier, Amsterdam, pp. 287–331.

- Hartvigsen, M., 2015. Experiences with land consolidation and land banking in Central and Eastern Europe after 1989. Land Tenure Working Paper 26, FAO.
- Hashimoto, S., Nishi, M., 2016. Policy evolution of land consolidation and rural development in postwar Japan. *Geomatics, Land management and Landscape* 30, 57–75.
- Imbens, G.W., Wooldridge, J.M., 2009. Recent developments in the econometrics of program evaluation. *Journal of Economic Literature* 47(1), 5–86.
- Ito, J., Nishikori, M., Toyoshi, M., Feuer, H.N., 2016. The contribution of land exchange institutions and markets in countering farmland abandonment in Japan. *Land Use Policy* 57, 582–593.
- Ito, J., 2022. Program design and heterogeneous treatment effects of payments for environmental services. *Ecological Economics* 191, 107235.
- Kunimitsu, Y., 2005. Impacts of paddy-field consolidation projects on farmland rental transactions: Application of discrete choice model. *Japanese Journal of Rural Economics* 7, 49–60.
- MAFF, 2010. Annual report on food, agriculture and rural areas in Japan. Ministry of Agriculture, Forestry and Fisheries, Tokyo.
- Nakashima, Y., 1989. Structure of rice production and farmland capital stocks. *Journal of Rural Economics* 61(1), 19–28.
- OECD, 2019. Innovation, agricultural productivity and sustainability in Japan: OECD food and agricultural reviews. OECD Publishing, Paris.
- Otsuka, K., Liu, Y., Yamauchi, F., 2016. Growing advantage of large farms in Asia and its implications for global food security. *Global Food Security* 11, 5–10.
- Rosenbaum, P.R., Rubin, D.B., 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), 41–55.
- Shogenji, S., 1998. Economic analysis of modern agricultural policy. University of Tokyo Press, Tokyo.
- Takayama, T., Nakatani, T., 2018. Impact of participatory development projects on social capital: Evidence from farmland consolidation projects in paddy areas of Japan. *Japan Agricultural Research Quarterly* 52(3), 273–283.
- Takayama, T., Nakatani, T., Senda, T., Fujie, T., 2021. Less-favoured-area payments, farmland abandonment and farm size: Evidence from hilly and mountainous areas in Japan. *Australian Journal of Agricultural and Resource Economics* 65(3), 658–678.
- Van Dijk, T., 2003. Dealing with Central European land fragmentation: A critical assessment on the use of land consolidation. *Geoforum* 34(4), 487–498.
- Wager, S., Athey, S., 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523), 1228–1242.
- Wing, C., Simon, K., Bello-Gomez, R.A., 2018. Designing Difference in Difference Studies: Best Practices for Public Health Policy Research. *Annual Review of Public Health* 39, 453–469.

Table A-1. Definition of Variables

Variable	Definition
Panel A. Treatment Variables	
$\Delta L30$	Change in the share of farmland plots larger than 30 ares between 2000 and 2010.
L30	Share of farmland area in plots larger than 30 ares.
Treatment indicator	Equals 1 if $\Delta L30$ increased by at least 10 percentage points between 2000 and 2010, and 0 otherwise.
Panel B. Outcome Variables	
Rent-in farmland share	Share of farmland rented in by farm households (%).
Abandoned farmland share	Share of abandoned farmland area in total farmland area (%).
Farms ≥ 5 ha	Share of farm households operating at least 5 ha (%).
Panel C. Geographic Variables	
Agricultural area (urban)	Dummy variable equal to 1 if the rural community is classified as an urban agricultural area.
Agricultural area (intermediate)	Dummy variable equal to 1 if the rural community is classified as an intermediate agricultural area.
Agricultural area (mountainous)	Dummy variable equal to 1 if the rural community is classified as a mountainous agricultural area.
Distance to densely inhabited district: <30 min	Indicator related to accessibility to densely inhabited districts within 30 minutes.
Distance to densely inhabited district: 30–60 min	Indicator related to accessibility to densely inhabited districts within 30–60 minutes.
Distance to densely inhabited district: 60–90 min	Indicator related to accessibility to densely inhabited districts within 60–90 minutes.
Distance to densely inhabited district: >90 min	Indicator related to accessibility to densely inhabited districts beyond 90 minutes.
Flat terrain	Dummy variable equal to 1 if the rural community is classified as flat terrain.
Gentle slope	Dummy variable equal to 1 if the rural community is classified as gentle slope terrain.
Steep slope	Dummy variable equal to 1 if the rural community is classified as steep slope terrain.

Variable	Definition
Panel D. Institutional Land-use Variables	
Urbanization promotion area	Dummy variable equal to 1 if the rural community is located within an urbanization promotion area.
Urbanization control area	Dummy variable equal to 1 if the rural community is located within an urbanization control area.
Nondesignated planning area	Dummy variable equal to 1 if the rural community is located outside designated city planning areas.
Agricultural promotion area	Dummy variable equal to 1 if the rural community is designated as an agricultural promotion area.
Agricultural land zone	Dummy variable equal to 1 if the rural community is designated as an agricultural land zone within agricultural promotion areas.
Panel E. Demographic and Agricultural Structure Variables	
Number of farm households	Total number of farm households in the rural community.
Elderly farmer ratio	Share of agricultural workers aged 65 years or older.

Table A-2. Robustness to Alternative Treatment Thresholds

Method	Threshold Definition	Rent-in Farmland	Abandoned Farmland	Farms ≥ 5 ha
DID	$\Delta L30 \geq 5pp$	0.0118*** (0.0029)	-0.0036*** (0.0011)	0.0094*** (0.0015)
DID	$\Delta L30 \geq 10pp$	0.0162*** (0.0028)	-0.0047*** (0.0011)	0.0135*** (0.0016)
DID	$\Delta L30 \geq 15pp$	0.0198*** (0.0031)	-0.0055*** (0.0013)	0.0159*** (0.0018)
PSM-DID	$\Delta L30 \geq 5pp$	0.0107** (0.0045)	-0.0034* (0.0019)	0.0089*** (0.0024)
PSM-DID	$\Delta L30 \geq 10pp$	0.0149*** (0.0043)	-0.0048** (0.0019)	0.0110*** (0.0022)
PSM-DID	$\Delta L30 \geq 15pp$	0.0181*** (0.0048)	-0.0057*** (0.0021)	0.0138*** (0.0025)

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$