Empirical Analysis of the Dynamic Effects of Farmland Gift Tax Exemption Policy in South Korea

by

Minjeong Kwak and Kwansoo Kim

Abstract

South Korea has achieved rapid economic development through urbanization and industrialization, but as a result the share of the rural population and agricultural income relative to the total population have steadily declined. To secure a stable rural population and raise farm household incomes, the government has implemented various support programs for agricultural households; among these is the “Gift Tax Exemption for Farming Successors” policy. Under this policy, farming parents may transfer farmland to their children without tax burdens, reducing the cost of land acquisition for successor farmers and thereby enhancing the sustainability of family farming. This study empirically examines whether the gift tax exemption policy for farming successors has a meaningful impact on cultivating the next generation of farmers.

First, we evaluate the policy’s effects on farm population growth and the decline in average farm household age. Second, we analyze how these policy effects evolve dynamically over time. To estimate the policy impact, prior studies have combined propensity score matching (PSM) with Difference-in-Differences (DID) methods using the generalized propensity score (GPS) to address the binary treatment limitation of PSM and comparing pre- and post-policy changes in average household age and a number of new farm households who are engated in agriculture. However, this approach only assesses covariate balance across treatment intervals after the estimation, which does not guarantee balanced farm characteristics at each level of treatment intensity.

To overcome this methodological limitation, our research employs the Covariate Balancing Generalized Propensity Score (CBGPS) approach. Additionally, we use a multiple-period DID framework to trace changes in average age and farm household counts across two intervals (the baseline and comparison periods) by treatment status. Specifically, we divide the period from 2016 to 2024 into four-year segments and examine policy effects at three points in time: the baseline (2010), Period 1 (2011–2015), and Period 2 (2016–2020).

For estimating policy effects, we draw on National Tax Service and Statistics Korea data, incorporating variables such as per-household tax exemption amounts and regional farm characteristics. We expect our findings to provide rigorous empirical evidence on the effectiveness of the gift tax exemption policy in fostering farming successors and to assess whether the policy merits continuation over the long term. Beyond South Korea, the results should offer valuable policy insights for other countries seeking to promote intergenerational succession in agriculture.

Keywords: successor farmers; counteracting rural depopulation; generalized propensity score (GPS); Difference-in-Differences (DID); dynamic treatment effects

1. Introduction

Agriculture is a foundational industry whose continuity must be guaranteed. South Korea's unprecedented rapid urbanization and industrialization, driven by growth concentrated in urban areas, have generated serious structural problems between urban and rural regions. Rural aging, a shortage of farming successors, and widening income gaps between urban and rural populations have disrupted the intergenerational transfer of farms, thereby threatening the sustainability of agriculture. Recently, the risk of regional extinction has reached a critical level. The shortage of agricultural labor across the sector has intensified, prompting a rapid increase in the use of foreign workers as a substitute. These issues extend beyond mere population decline—they jeopardize farmland utilization efficiency, continuity of farm management, and the preservation of local communities and rural culture.

In response, since 1987 the Korean government has initiated the "Gift Tax Exemption for Farming Successors" program. Under this scheme, when senior-generation farmers transfer farmland to their children who are engaged in agriculture, the gift tax is waived subject to certain conditions. The policy is designed to encourage intergenerational transfer of farmland, thereby promoting generational renewal and structural improvement within the agricultural sector.

Although the program has granted an average annual tax exemption of KRW 71.3 billion over the past three years, the nominal size of exemptions alone cannot gauge its effectiveness. It is essential to empirically assess whether the policy has genuinely increased the inflow of successor farmers or mitigated rural aging. In particular, examining which regional and individual characteristics influence the "successor conversion" process—where beneficiaries go on to actively farm after receiving the gift—can offer valuable empirical insights for refining the policy going forward.

This study builds on these concerns by merging National Tax Service gift-tax exemption records with Agricultural Census data at the county (si-gun-gu) level and treating the exemption amount as the continuous treatment variable. We employ the Covariate Balancing Generalized Propensity Score (CBGPS) approach—which centralizes and orthogonalizes covariates—to estimate how varying degrees of policy benefit affect outcome measures such as the number of farms per county and the average age of farm operators. In our analysis, we also address key identification challenges, including the endogeneity of exemption amounts, cross-regional heterogeneity, and the dynamic evolution of treatment effects over time. Through this empirical framework, we evaluate whether the gift-tax exemption policy has substantively promoted generational renewal and the cultivation of successor farmers in Korean agriculture, and we offer evidence-based recommendations for future policy refinement.

2. Method

The program was first enacted in 1987, with eligibility criteria specified for the donated land, the donor, and the recipient. Since its inception, it has undergone incremental reforms—such as adjustments to exemption limits and expansions of eligible land—and over the past decade it has been implemented at an average annual scale of KRW 32.1 billion. In other words, treatment does not occur only once in a given year but is delivered continuously over a long period. Moreover, exemption eligibility is not determined solely by the policy's introduction; it depends on whether donors and recipients meet the statutory requirements and on their voluntary participation decisions. Because the exemption amount is proportional to the size of the transferred land, the treatment intensity varies across observations. Accordingly, in this program the treatment is measured as a continuous variable, and its levels are heterogeneously distributed across units according to their characteristics.

Moreover, this program can be combined with other farm support measures aimed at safeguarding farm household incomes, regardless of whether those other benefits are received. Because treatment intensity is influenced by beneficiaries' self selection, an endogeneity problem may arise, and the existence of similar but separate programs risks biasing the estimated treatment effects—either downward or upward—creating an identification challenge. In light of these program characteristics and potential issues, it is necessary to adopt an appropriate empirical estimation strategy.

Under these circumstances, a simple comparison of group means by treatment status is inadequate for evaluating outcomes. One method for addressing endogeneity is the Propensity Score Method (PSM) (Rosenbaum et al., 1983), which, unlike a randomized trial, constructs comparable groups by matching units with similar propensity scores—each score summarizing observed covariates into a single scalar. However, this approach requires the treatment variable to be binary, whereas in our case treatment intensity is a continuously varying quantity even within the “treated” group. Therefore, PSM is not suitable for our analysis.

In response, KDI (2020) set the size of the gift-tax exemption—the treatment intensity—as a continuous treatment variable rather than a simple binary indicator. Because observations could no longer be cleanly split into treated and control groups, they introduced a dose-response function and estimated outcomes as a continuous function of treatment level using the Generalized Propensity Score (GPS) method (Hirano and Imbens, 2004). Both PSM and GPS construct comparison groups based on observed covariates, making it essential to verify covariate balance across treatment levels. Under PSM, where the treatment is binary, one can directly test balance between treated and control groups by conducting post-matching checks. With GPS, which handles a continuous treatment, one first discretizes the treatment into intervals and then assesses covariate balance within each interval. However, even though balance diagnostics are performed prior to estimating the dose-response function, when GPS values enter the function as covariates, there remains an unavoidable risk that, at certain treatment intensities, covariate distributions will be imbalanced.

To address this limitation, our study employs the Covariate-Balancing Generalized Propensity Score (CBGPS) method (Fong et al., 2018). Like the standard GPS approach, CBGPS estimates propensity scores for a continuous treatment using observed covariates, but it additionally enforces, a priori, that covariate distributions remain balanced at every level of treatment. By incorporating this constraint into the dose-response estimation, CBGPS minimizes covariate imbalance when GPS enters the outcome model, yielding more reliable causal estimates.

Building on and extending prior work, we also expand the analysis period to capture dynamic policy effects. Using data from 2010 through 2023, we implement a multi-period Difference-in-Differences design to assess both the persistence and the trajectory of the policy's impact. Rather than a single before-and-after comparison, we divide the timeline into multiple intervals—Period 1 (2010-2015) and Period 2 (2016-2020)—and examine how estimated treatment effects evolve across these segments. By estimating the slope of treatment effects in each interval, we uncover not only short-term shifts but also longer-term trends in policy effectiveness. This approach, known as DID

-MP (Difference-in-Differences with Multiple Periods), is informed by Egami and Yamauchi (2023).

3. Empirical Model

This study applies the Covariate Balancing Generalized Propensity Score (CBGPS) method to estimate the causal effect of a continuous treatment variable. A continuous treatment variable is used because the magnitude of the gift tax exemption varies across observation units, namely, administrative districts (si/gun/gu). Traditional propensity score matching techniques that dichotomize treatment status assume that the treatment is applied at a relatively uniform level within the treated group, making them unsuitable for capturing the heterogeneity inherent in this policy.

Furthermore, ensuring covariate balance is crucial because only when the association between covariates and treatment levels is minimized can the treatment be regarded as randomly assigned conditional on observed covariates. For example, if regions with a high proportion of successor farmers also tend to receive larger tax exemptions, it becomes difficult to disentangle the effect of the policy from other regional characteristics. In such cases, CBGPS plays a role in statistically adjusting for inter-regional differences in characteristics.

In conclusion, the CBGPS method creates an analytical environment in which the causal effect of a continuous treatment variable on an outcome variable can be assessed by minimizing the association between observed covariates and the treatment level. This section explains the fundamental assumptions and estimation formulas of the CBGPS approach to justify its selection for this study.

While CBGPS offers the advantage of greater flexibility over binary treatment models by accommodating continuous treatments, it has the limitation of identifying only the average treatment effect (ATE) at a single point in time. However, the policy examined in this study is not a one-off intervention implemented in a specific year; rather, it is a long-term policy that has been executed repeatedly over time. The amount of tax exemption also tends to accumulate or fluctuate over the years.

Most importantly, the treatment variable—i.e., the amount of gift tax exemption—is determined voluntarily by the donor, introducing potential endogeneity. Since the characteristics that influence the timing and scope of the gift are unobservable, and if these same characteristics also affect the outcome variable, the estimated causal effect may be biased.

Therefore, estimating the average treatment effect (ATE) at a single point in time is insufficient to capture the overall dynamic effects of the policy, which may accumulate over time or manifest with a time lag. To account for changes in both the treated and untreated groups before and after the policy intervention, this study adopts the Difference-in-Differences (DID) methodology.

The analytical framework of this study consists of three components: (1) a continuous treatment variable T_i assigned to each observational unit i, (2) a vector of pre-treatment covariates for that unit $X_i \in R^K$, and (3) an outcome variable $Y_i(t)$ representing the potential result corresponding to the given level of treatment. The generalized propensity score (GPS) $f(T_i | X_i)$ is defined as the conditional probability density function of receiving a particular treatment level given a specific set of covariates. It serves as a summary measure of the influence of observed covariates on the treatment assignment process. Here, the treatment variable T_i is treated as a random variable, and the treatment level t is specified as an evaluation point—a fixed constant used to distinguish specific treatment intensities.

The covariate balancing generalized propensity score (CBGPS) is estimated by simultaneously modeling this conditional density function and minimizing the correlation between the treatment variable and covariates. This approach ensures that the distribution of covariates is balanced across different treatment levels. In this study, the main analytical model is based on $r(t, x)$ the dose-response function (DRF), which is defined for each level of treatment $Y_i t$. The DRF allows for the estimation of the average expected outcome for a given treatment level, conditional on covariates.

For valid causal inference, the following three assumptions must be satisfied. The first is the Strong Ignorability Assumption.

$$Y_i(t) \perp T_i | X_i \tag{1}$$

This assumption implies that, conditional on covariates X_i, the treatment level T_i is independent of the potential outcome $Y_i(t)$. In the context of causal inference, this serves as a key assumption for ensuring conditional randomization, even in non-experimental settings.

The second is the Common Support Assumption, which states the following:

$$0 < f(T_i = t | X_i = x) < \infty \qquad \forall t, x \tag{2}$$

This assumption requires that the generalized propensity score $r(t,x)$ is strictly positive for all treatment levels t, given covariates X_i. In other words, it ensures sufficient comparability across treatment levels by requiring that each unit has a non-zero probability of receiving any treatment level.

Lastly, the Stable Unit Treatment Value Assumption (SUTVA) states the following:

$$ATE_{[t1,t2]} = E[Y_i(t_2) - Y_i(t_1)] \tag{3}$$

This assumption implies that the treatment received by one unit does not affect the outcomes of other units, meaning there is no interference between observational units. When these assumptions are satisfied, the policy effect can be identified either through the dose-response function (DRF), $r(t)$ or as an Average Treatment Effect (ATE) over a specific interval. Identification via the DRF allows for tracing the entire effect curve—showing how the outcome $Y_i(t)$ changes as the treatment level t increases.

On the other hand, identification via the ATE enables the estimation of the effect by comparing the average outcomes between two specific treatment levels, say from t_0 to t_1, thus capturing the marginal impact across that interval.

This analytical framework, unlike simple mean comparisons or linear regression, explicitly models the treatment assignment process and minimizes potential bias based on covariates, thereby enabling more refined and reliable causal inference. In particular, when estimating the policy effect of a continuous treatment variable such as the amount of tax exemption, the use of CBGPS allows for the identification of average causal effects under a condition of statistical unconfoundedness, as if the treatment were randomly assigned.

Conventional GPS approaches model the distribution of the continuous treatment variable but assess covariate balance only ex post, which limits their ability to ensure unbiased estimation. In contrast, the CBGPS approach adopted in this study estimates the propensity score by explicitly minimizing the correlation between covariates and the treatment level.

This study employs a parametric approach, rather than a nonparametric one, to estimate the parameters of the GPS model using the Generalized Method of Moments (GMM). This procedure enables the construction of stabilized weights, which are then used to perform weighted regression analyses of the outcome variable under conditions where covariate imbalance has been addressed.

The estimation procedure consists of the following three steps.

Step 1: Preprocessing

The first step is the preprocessing stage, in which the continuous treatment variable T_i is centered and standardized to have a mean of 0 and a variance of 1, resulting in a transformed variable T_i^*

$$T_i^* = \frac{T_i - \overline{T}}{s_T} \tag{4}$$

At this stage, the mean treatment level $\overline{T}$ and the sample variance s_T^2 of the treatment variable across the full sample can be calculated using the following formulas.

$$\overline{T} = \frac{1}{N}\sum_{i=1}^{N} T_i \tag{5}$$

$$s_T^2 = \frac{1}{N}\sum_{i=1}^{N}(T_i - \overline{T})^2 \tag{6}$$

The covariates X_i are also centered to have a mean of 0 and are orthogonalized to remove potential multicollinearity.

$$X_i^* = X_i - \overline{X}\ , \quad where\ \overline{X} = \frac{1}{N}\sum_{i=1}^{N} X_i \tag{7}$$

At this stage, an orthogonalization process is performed to eliminate multicollinearity among covariates and transform them into a structure that approximates mutual independence. Let X_i denote the covariate matrix and Σ_X its covariance matrix. The following transformation is applied, and the covariance matrix Σ_X can be decomposed through eigenvalue decomposition as follows:

$$\widetilde{X}_i = \Sigma_X^{-1/2} X_i^* \\ = Q\Lambda^{-1/2} Q' X_i^*, \ where\, \Sigma_X = Q\Lambda Q' \tag{8}$$

Through this transformation, the resulting matrix $\widetilde{X}_i$ has a covariance matrix that approximates the identity matrix, thereby enhancing the linear independence among the covariates.

This step serves multiple purposes: it eliminates the influence of scale differences among variables, clarifies the interactions between the treatment and covariates, and improves the stability and efficiency of propensity score estimation. In addition, this preprocessing lays the groundwork for generating stable weights in the subsequent estimation steps.

Step 2: Model Estimation and Weight Construction

The second step involves specifying the CBGPS model and estimating its parameters. After preprocessing, the treatment variable is transformed into its standardized form T_i^*, and the covariates into mean-centered and orthogonalized variables X_i^*. Using these transformed variables, a conditional normal model is fitted, assuming that the treatment variable follows a normal distribution given the covariates:

$$T_i^* | X_i^* \sim N(X_i^{*\prime}\beta, \sigma^2) \tag{9}$$

From this model, the conditional normal density function $f(T_i | X_i)$ is derived and used as the generalized propensity score (GPS). The conditional density function takes the following form:

$$f_{\theta}(T_i^* | X_i^*) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{ - \frac{(T_i^* - X_i^{*\prime}\beta)^2}{2\sigma^2} \right\} \tag{10}$$

At the same time, the standard normal density function $f(T_i)$ is used as a stabilizing factor. The stabilized weight is then defined as the ratio of these two densities:

$$w_i(\theta) = \frac{\phi(T_i^*)}{f_{\theta}(T_i^* | X_i^*)} \tag{11}$$

This weight plays a crucial role in reducing the variance of the estimator by down-weighting extreme values of the treatment variable. In essence, it re-weights the data such that the treatment assignment appears fair or as if it were randomly assigned.

To obtain the parameters of the model, the parameter vector θ is estimated using the Generalized Method of Moments (GMM). This estimation is conducted under two conditions: (a) the score condition, which ensures the correct specification of the conditional normal model, and (b) the covariate balance condition, which ensures that the weighted covariate distribution is independent of the treatment level.

$$E[\frac{\partial}{\partial\theta}\log f_{\theta}(T_i^* | X_i^*)] = 0 \qquad (12-a)$$

$$E[w_i(\theta) \cdot T_i^* \cdot X_i^*] = 0 \qquad (12-b)$$

These moment conditions jointly ensure that the estimated weights induce covariate balance and enable valid identification of causal effects.

Here, $f(T_i)$ denotes the standard normal density function, which serves as the numerator representing the stabilized distribution of the treatment variable, while the denominator $f(T_i | X_i)$ is interpreted as the generalized propensity score (GPS). The parameters $\hat{\theta}$ of the GPS model—estimated through the conditional normal specification—are used to compute the weights $w_i(\hat{\theta})$.

This weighting process can be interpreted as reweighting the data such that the treatment variable becomes conditionally independent of the covariates. Importantly, this reweighting step marks a key distinction between GPS and CBGPS approaches. While GPS models the adjusted relationship between the outcome and treatment using a single conditional density function $f(T_i | X_i)$, CBGPS goes a step further: even after estimating the GPS, it optimizes the weights to eliminate any remaining imbalance in the distribution of covariates across all levels of treatment. As a result, CBGPS enables the estimation of the outcome–treatment response function under a setting where the distribution of covariates is balanced at every treatment level, thereby offering a stronger basis for causal identification.

Step 3: Causal Effect Estimation.

The third step is the causal effect estimation stage, in which a weighted regression analysis of the outcome variable is conducted using the stabilized weights obtained in the previous step. In this stage, the regression model is specified in the following linear form:

$$\begin{aligned} Y_i &= \delta_0 + \delta_1 T_i + \delta_2{}' X_i + \epsilon_i \\ &= Z_i{}'\delta + \epsilon_i, \ \ where \ \ Z_i = (1 \ \ T_i \ \ X_i{}')' \end{aligned} \tag{13}$$

where $Z_i = (1, T_i, X_i{}')$ may include an intercept, the treatment variable T_i, and a subset of covariates. Since the generalized propensity score (GPS) already adjusts for differences in covariate distributions across treatment levels, it is not necessary to include the full set of covariates in the regression model. This helps prevent overfitting and improves the interpretability of the estimated coefficients. Moreover, this structure aligns with the doubly robust estimation strategy, wherein valid causal inference can be obtained as long as either the treatment model (GPS) or the outcome model is correctly specified.

The coefficient vector δ is estimated by minimizing the following weighted least squares (WLS) objective function:

$$\hat{\delta} = \arg\min_{\delta} \sum_{i=1}^{N} w_i (Y_i - Z_i{}'\delta)^2 \tag{14}$$

Alternatively, this estimator can also be derived by solving the corresponding moment condition, which ensures the weighted residuals are orthogonal to the regressors:

$$E[\widehat{w_i}(Y_i - Z_i{}'\delta) Z_i] = 0 \tag{15}$$

That is, the estimates obtained from the WLS minimization and those satisfying the moment condition are numerically equivalent when stabilized weights are used.

In the context of CBGPS, the selection of the regressor vector $Z_i = (1, T_i, X_i{}')$ requires particular care. The covariates included should contain only the necessary information to model the conditional distribution of the treatment variable T_i, without being overly collinear with it. At the same time, the selected covariates should retain interpretability with respect to the outcome variable Y_i, exhibit low multicollinearity, and contribute to stable post-estimation covariate balance. These considerations help ensure that the resulting estimates are both statistically stable and substantively meaningful.

Once the coefficients $\hat{\delta}$ are estimated, the dose - response function(DRF), defined as $\hat{r}(t)$ can be approximated as follows:

$$\hat{r}(t) = \delta_0 + \delta_1 t + \delta_2' E[X_i] \tag{16}$$

This specification represents the expected outcome Y_i at treatment level t, conditional on the average values of the covariates. Here, δ_0 captures the baseline outcome, δ_1 reflects the marginal effect of the treatment, and $\delta_2' E[X_i]$ accounts for the contribution of covariates evaluated at their means. By holding covariate characteristics constant, this formulation allows for the construction of a dose - response curve that isolates the effect of varying treatment intensities.

Using this estimated function, the average treatment effect (ATE) between two treatment levels t_0 and t_1 can be calculated as:

$$ATE_{[t1,t0]} = \hat{r}(t_1) - \hat{r}(t_0) \tag{17}$$

This provides a flexible and interpretable measure of how changes in the continuous treatment variable influence the expected outcome, independent of variation in covariate profiles. Despite this flexibility, the validity of causal inference still depends on the assumption that the treatment is exogenously assigned—which may not hold in observational data.

In this framework, Y_i denotes the observed outcome variable, where $Y_i(1)$ refers to the potential outcome for unit i under treatment $(T=1)$, and $Y_i(0)$ refers to the potential outcome under no treatment$(T=0)$. Since each observational unit can only be exposed to either the treatment or the control condition—not both simultaneously—the counterfactual outcome is inherently unobservable. Therefore, the treatment effect for unit i is not defined as $Y_i(1) - Y_i(0)$ directly, but rather as the expected difference:

$$ATE = E[Y_i(1) - Y_i(0)] \tag{18}$$

When treatment is a simple yes/no indicator, the average treatment effect (ATE) can be estimated by the basic regression model:

$$Y_i = \beta_0 + \beta_1 T_i + u_i \tag{19}$$

Here, T_i is a binary dummy (1 if unit i is treated, 0 otherwise), so the estimated treatment effect is $\widehat{\beta}_1^{LS}$. However, if assignment T_i is not random—if unobserved factors affect both T_i and the error term u_i —then and the OLS estimate $\widehat{\beta}_1^{LS}$ is biased.

Under the Conditional Independence Assumption, once we control for observed covariates X_i, treatment assignment T_i is as good as random. In regression form:

$$Y_i = \beta_0 + \beta_1 T_i + \beta_2 X_i + u_i \tag{20}$$

we require $T_i \perp u_i | X_i$ to get unbiased $\widehat{\beta}_1^{LS}$.

With panel data, we can difference out unobserved, time invariant factors. Let P_t be a post treatment time dummy. The DID model is:

$$Y_{it} = \beta_0 + \beta_1 T_i + \beta_2 P_{it} + \beta_3 T_i P_{it} + u_{it} \tag{21}$$

By comparing the outcomes of treated and untreated groups before and after the intervention period, one can difference out time-invariant unobserved factors and isolate the pure policy effect. This effect can be identified through the coefficient on the interaction term in the following regression specification, and is typically presented as follows:

Table 1. Identification of the Pure Policy Effect in a Binary DID Model

Group	Time: before $(P=0)$	Time: after $(P=1)$	diff
Treated $(T=0)$	$\widehat{\beta}_0$ $E[Y\|T=0,P=0]$	$\widehat{\beta}_0+\widehat{\beta}_2$ $E[Y\|T=0,P=1]$	$\widehat{\beta}_2$
Control $(T=1)$	$\widehat{\beta}_0+\widehat{\beta}_1$ $E[Y\|T=1,P=0]$	$\widehat{\beta}_0+\widehat{\beta}_1+\widehat{\beta}_2+\widehat{\beta}_3$ $E[Y\|T=1,P=1]$	$\widehat{\beta}_2+\widehat{\beta}_3$
diff	$\widehat{\beta}_1$	$\widehat{\beta}_1+\widehat{\beta}_3$	$\widehat{\beta}_3$

The Difference-in-Differences (DID) method has the advantage of being intuitively easy to interpret and capable of estimating the pure causal effect of a policy intervention. However, a key limitation lies in its inability to capture whether the randomness of treatment assignment remains valid over time—particularly when heterogeneous factors may evolve and influence treatment assignment in later periods.

To address this, an extended approach known as the Multiple-Period Difference-in-Differences (DID-mp) has been developed. This method leverages repeated observations of both treatment status and outcome variables across multiple time periods, allowing for more dynamic modeling of treatment effects over time. The corresponding estimation equation is specified as follows:

$$Y_{it} = \alpha + \lambda_t + \mu_i + \beta T_{it} + \gamma' X_{it} + \epsilon_{it} \tag{22}$$

In this specification, λ_t represents time fixed effects, which control for time-specific shocks that are common to all units. μ_i denotes unit fixed effects, capturing time-invariant characteristics of each observational unit. The variable T_{it} is the treatment indicator (or intensity) at time t, and the coefficient β captures the average marginal effect of a one-unit increase in the treatment level on the outcome variable.

This approach allows for the identification of time-varying treatment effects and helps mitigate violations of the parallel trends assumption by utilizing repeated observations across time. The parallel trends assumption is a key identifying condition in DID analysis, and can be expressed formally as:

$$E[Y_{it} - Y_{it-1} | T_i = 1] = E[Y_{it} - Y_{it-1} | T_i = 0] \tag{23}$$

This condition implies that, in the absence of treatment, the treated and control groups would have followed similar trends in outcomes over time. Ensuring the plausibility of this assumption is crucial for the validity of DID-based causal inference.

In other words, the parallel trends assumption implies that, in the absence of treatment, the change in outcomes over time would have been the same for both treated and control groups. That is, aside from the treatment itself, there should be no structural difference in how outcomes evolve over time across groups. This assumption is fundamental to the identifiability of causal effects in DID frameworks.

However, even with repeated measurements, additional adjustments are required to account for the endogenous distribution of treatment intensity. To address this, one can incorporate the Covariate Balancing Generalized Propensity Score (CBGPS) approach into the DID-mp framework. Specifically, by first estimating the generalized propensity score (GPS) and deriving the stabilized weights, the method proceeds with a weighted multiple-period DID, commonly referred to as the Double Robust DID.

This approach enables the identification of outcome changes under the assumption that treatment is as-if randomly assigned, while simultaneously adjusting for both selection bias and temporal dynamics. When covariate balance is preemptively achieved through CBGPS, the parallel trends assumption becomes more plausible, enhancing the internal validity of DID even in non-randomized observational settings.

The procedure follows a sequential structure: First, the conditional density function of the continuous treatment variable is estimated given the covariates X, as shown in Equation (10). Next, the stabilized weight is calculated using Equation (12).

Finally, these weights are applied to a weighted DID-mp regression model to estimate the causal effect of the treatment. The weighted regression model is specified as follows:

$$Y_{it} = \alpha_i + \lambda_t + \phi T_i + \delta D_t + \theta(T_i \times D_t) + \gamma' X_{it} + \epsilon_{it} \tag{24}$$

The double robust DID approach compares the post‐pre outcome differences of treated units with those of untreated units to approximate the counterfactual outcomes that the treated units would have experienced in the absence of treatment. As a result, this method does not identify the average treatment effect (ATE) for the entire population, but rather the average treatment effect on the treated (ATT). In this framework, the coefficient of interest that identifies the policy effect corresponds to θ.

To test for treatment effect heterogeneity, a triple interaction term is introduced into the double robust DID model. By incorporating a heterogeneity variable, the model allows for the evaluation of how the marginal effect of treatment intensity varies depending on specific subgroups or characteristics. The extended model is specified as follows:

$$\begin{aligned} Y_{it} = {} & \alpha_i + \lambda_t + \phi T_i + \delta D_t + \eta Z_i + \theta(T_i \times D_t) \\ & + \psi(T_i \times Z_i) + \rho(D_t \times Z_i) + \varkappa(T_i \times D_t \times Z_i) + \gamma' X + \epsilon \end{aligned} \tag{25}$$

Each coefficient in the model identifies a distinct causal or structural effect:

α_i captures unit fixed effects, accounting for time-invariant characteristics of each observational unit (e.g., land conditions, topography at the county level). λ_t represents time fixed effects, controlling for national-level shocks or common events occurring at specific time points (e.g., policy changes, climate fluctuations). ϕ denotes the main effect of treatment intensity, identifying the average effect (ATE) of a one-unit (e.g., one million KRW) increase in the treatment variable, regardless of group or period. δ represents the period (comparison group) effect, capturing the average shift in the outcome variable for all units between the baseline and comparison periods. η identifies the main effect of the heterogeneity variable, capturing the average difference in the outcome between groups (e.g., high vs. low farming experience), irrespective of treatment or time.

The interaction terms capture more nuanced causal relationships:

θ, the baseline ATT, measuring the average treatment effect on the treated during the comparison period, independent of the heterogeneity group.

ψ, the interaction between treatment intensity and the heterogeneity variable, identifying whether the effect of treatment differs by heterogeneity group in the baseline (or untreated) period. For example: “In high-experience regions, the effect of each additional unit of tax exemption is ψ larger.”

ρ, the interaction between time and the heterogeneity variable, identifying whether the time-related change in outcomes differs by heterogeneity group. For instance: “Regions with high farming experience saw a ρ-unit larger shift in the outcome during the comparison period than regions with low experience.”

$\varkappa$, the triple interaction coefficient, representing the heterogeneous ATT. It measures how the effect of treatment intensity differs across heterogeneity groups in the post-treatment period. This can be interpreted as: “In high-experience regions, the ATT per unit of tax exemption increased (or decreased) by $\varkappa$ compared to the baseline ATT.”

γ, represents the effects of time-varying covariates, capturing auxiliary influences of changing contextual factors on the outcome variable.

4. Data

In this study, the treatment variable is defined as the five-year cumulative average amount of gift tax exemption per capita. This definition reflects the institutional feature of the policy, which limits the total exemption amount to 100 million KRW over a five-year period. Given the self-selection nature of both the decision to gift farmland and the level of exemption received, we opted for a continuous treatment variable—namely, the exemption amount—instead of a binary indicator. This approach better captures the heterogeneous policy exposure across observations.

To evaluate the generational turnover effects of the policy, we consider two outcome variables: (1) the number of farm households and (2) the average age of farm household heads. Rather than focusing on the absolute levels of these variables, the analysis centers on their temporal changes. This strategy is grounded in the hypothesis that, in regions suffering from rural aging and population decline, larger exemption amounts may contribute to slowing the decline in farm households and mitigating the rise in the average age of farm heads. Such a pattern would indicate that the policy has positively affected generational renewal in the agricultural sector.

To ensure valid comparisons between treated and control groups, we selected four covariates that influence both the treatment assignment and outcome variables, while also satisfying the covariate balance condition. These are: (1) the urbanization ratio (urban area over total area), (2) the log-transformed total farmland area, (3) the proportion of full-time farms, and (4) average agricultural income. These variables were incorporated into the analysis model only after confirming their balancing effectiveness in the estimation process of the Covariate Balancing Generalized Propensity Score (CBGPS).

In addition to the core covariates, we included several control variables that may confound the relationship between the treatment and outcomes. These include farming experience, the proportion of successor children among farm heads, the ratio of farmland to total area, the farm employment rate (number of long-term agricultural workers per cultivated area), farm mechanization level (number of machines per hectare), and the proportion of facility farming area. We also incorporated land market characteristics such as the total area of traded farmland, average transaction price, number of transactions, and number of gifted parcels.

To construct the dataset, we relied on several government data sources. The treatment variable is based on administrative records from the National Tax Service (NTS) on gift tax exemptions. The outcome variables were drawn from the “Census of Agriculture, Forestry, and Fisheries” conducted by Statistics Korea. Covariates and other control variables were compiled from the “Population Census” and household-level agricultural census data from Statistics Korea, land use data from the Ministry of Land, Infrastructure and Transport, and land market statistics from the Korea

Rural Community Corporation.

Due to the limitations of the available data, the unit of analysis was set at the county (Si/Gun/Gu) level. This decision stems from the fact that the gift tax exemption data were anonymized and thus could not be linked at the individual level. Consequently, the outcome and explanatory variables were also aggregated to the county level. Variables such as farmland area, agricultural income, and farming experience were transformed into county-level averages or ratios, allowing for valid cross-regional comparisons and policy effect estimations.

<Table 1> Variable Definitions and Summary Statistics

var.	Variable Name			Mean	Variance	Min	Max
Outcome Var.	Average Age of Farm Head (years)		2010	60.59	9.52	53.63	66.27
			2015	64.15	4.04	57.93	68.66
			2020	65.56	3.74	59.48	69.46
	Number of Farm Households		2010	4708.0	1.87e+7	5	20808
			2015	4318.8	1.50e+7	22	18670
			2020	4140.4	1.13e+7	25	17783
Treatment Var.	Avg. Per Capita Exempted Amount (Million KRW)		2010	12.29	422.49	0	100
			2015	16.33	392.88	0	137
			2020	21.63	428.25	0	100
Control Var.	Covariates	Urban Area Ratio	2010	0.487	0.187	0.0052	1
			2015	0.482	0.177	0.0052	1
			2020	0.481	0.177	0.0033	1
		Total Farmland Area (ha)	2010	6008.3	4.29e+7	3.52	31837.8
			2015	5392.8	3.50e+7	7.15	28812.6
			2020	4632.5	2.50e+7	12.70	26779.5
		Full-Time Farmer Ratio	2010	0.44	0.022	0	0.78
			2015	0.46	0.020	0.13	0.79
			2020	0.52	0.021	0.12	0.87
		Avg. Agricultural Income (Million KRW)	2010	13.23	3.82e-7	149.23	4256.37
			2015	12.68	4.74e-7	118.64	4683.57
			2020	12.02	6.15e-7	121.49	5525.23
	Avg. Farming Experience (years)		2010	27.28	81.78	7.22	40.32
			2015	27.49	67.14	10.14	40.11
			2020	26.34	68.60	10.98	38.95
	Successor Child Ratio		2010	0.02	0.000721	0	0.35
	Farmland Ratio		2010	0.287	0.070	0	0.83
			2015	0.278	0.067	0	0.80
			2020	0.278	0.067	0	0.80
	Employment Level		2010	0.19	0.066	0	2.28
			2015	0.18	0.03	0	1.57
			2020	0.27	0.15	0	4.13
	Farm Machinery Ownership (machines/ha)		2010	0.703	0.070	0.11	1.88
			2015	0.731	0.079	0.12	1.95
			2020	0.735	0.081	0.08	1.95
	Facility Farming Share		2010	0.096	0.007	0	0.63
			2015	0.088	0.006	0	0.62
			2020	0.075	0.005	0	0.61

	Total Transaction Area (ha)	2010	1839	2757517	159	7433
		2015	5145	25246890	278	25266
		2020	9048	69515942	240	37646
	Avg. Market Land Price (Million KRW/transaction)	2010	0.109	0.017	0.007	0.67
		2015	0.115	0.015	0.008	0.56
		2020	0.146	0.026	0.011	0.75
	Number of Transactions	2010	1226	1016750	1	4540
		2015	3435	9414692	1	14936
		2020	6081	26710183	2	21895
	Number of Gifted Parcels	2010	26.5	1175.6	0	231
		2015	129.3	26710.9	0	1293
		2020	155.0	56740.2	0	2171

주: 1) a: Statistics Korea, 2010/2015/2020 Census of Agriculture
b: National Tax Service, 2010-2023 Gift Tax Exemption Data
c: Korea Rural Community Corporation, 2011-2023 Farmland Market Data
d: Korea Rural Community Corporation, 2014 Agricultural Promotion Zone Statistics
2) The policy imposes a cap of 100 million KRW in tax exemption over 5 years. Thus, this study defines the treatment variable as the 5-year per capita exemption amount.

5. Estimation Results

Inthe first stage of CBGPS estimation, we conducted an ordinary least squares (OLS) regression where the treatment variable—per capita tax exemption—was regressed on a set of covariates. This preliminary step allowed us to evaluate the statistical significance of each covariate. Notably, covariates with excessively high explanatory power can undermine covariate balance in the propensity score model. To address this concern, we selected a minimal set of four centered covariates: urbanization level, adjusted farmland scale (log-transformed), full-time farm ratio, and average agricultural income. These were chosen to maintain balance while preventing overfitting in the GPS estimation process.

Following this, we assessed whether the selected covariate combinations ensured common support across treatment levels. Given that multicollinearity may destabilize GPS estimation, we verified the stability of the covariate sets by examining the variance inflation factor (VIF) for each combination. Once covariates passed this diagnostic step, weights were derived to ensure that, conditional on the covariates, the distribution of the continuous treatment variable T became independent of the covariate profile—achieving covariate balance at each treatment level t.

The dataset was structured into three cross-sectional panels: the baseline year (2010), the first period (2011–2015), and the second period (2016–2020). Each cross-section was constructed using time-aggregated, cleaned variables. For each period, we used the CBGPS-derived weights to estimate fixed effects linear regression (FE-OLS) models that examined the causal effect of treatment T on outcomes Y. The FE-OLS framework was appropriate in this context, as it controlled for unobserved, time-invariant heterogeneity across counties (Si/Gun/Gu), while leveraging the balanced covariate distribution ensured by CBGPS.

(1) Period-Specific Treatment Effects

Table 2 presents the estimated treatment effects of per capita tax exemption on the average age of farm heads for each time period. At baseline (2010), an increase of one million KRW in the per capita exemption amount was associated with a statistically significant reduction of approximately 0.00427 years in the average age of farm heads. In the first period (2011-2015), the estimated reduction was approximately 0.0009 years per million KRW, and in the second period (2016-2020), the marginal effect was 0.00048 years—both statistically significant.

These results provide empirical support for the hypothesis that regions with higher exemption amounts experienced relatively greater inflows of younger farmers, suggesting that the gift tax exemption policy may have contributed to generational renewal in the agricultural sector. However, the estimated marginal effect of the exemption on average age diminishes over time—from the baseline to the second period—indicating a potential temporal weakening of the policy's effectiveness.

<Table 2> Estimated Results by Period: Dependent Variable

- Average Age of Farm Household Heads(Y_{i1})

Category	Variable	Baseline (2010) Estimate (s.e.)	Period 1(2011-15) Estimate (s.e.)	Period 2(2016-20) Estimate (s.e.)
Outcome Variable 1 (Average Age of Farm Household Heads)	Intercept	53.1553 **** (0.7573)	57.6868 **** (0.9448)	59.3578 **** (0.6951)
	Per Capita Exemption Amount	-0.00427 *** (0.0140)	-0.00089 *** (0.0003)	-0.00040 ** (0.0001)
	Average Market Price of Farmland	-0.0038 (0.0025)	-0.0044 **** (0.0012)	-0.0014 (0.0013)
	Average Agricultural Income	-0.0139.33 (0.02078)	-0.0143 (0.0218)	-0.0059 (0.0181)
	Average Farming Experience	0.3282 **** (0.0204)	0.3051 **** (0.02817)	0.2842 **** (0.0226)
	Total Cultivated Area	0.000033 * (0.000015)	0.000040 ** (0.000018)	-0.000007 (0.000020)
	Average Cultivated Area	-0.8913 ** (0.3453)	-1.4312 **** (0.3362)	-0.7775 ** (0.3406)

	per Household			
	Urban Area Ratio	0.6368 * (0.3268)	0.5376 (0.3817)	-0.0278 (0.4464)
	Level of Farm Mechanization	0.2617 (0.3507)	-0.3687 (0.4187)	-0.2697 (0.4128)
	Agricultural Employment Level	1.7784 * (0.9268)	1.3929 * (0.7228)	-0.2625 ** (0.1135)
	Proportion of Greenhouse Farming Area	-1.9006 (1.2607)	-2.5383 ** (1.0177)	-0.2549 (1.1824)

Notes:
*, **, ***, and **** indicate statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively. Numbers in parentheses represent robust standard errors. The time periods are defined as follows: Baseline = 2010; Period 1 = 2011–2015; Period 2 = 2016–2020.

Regarding the average farming experience variable, results show a consistent and statistically significant positive association with the average age of farm heads across all periods. At the baseline year (2010), an increase of one year in average farming experience was associated with an increase of approximately 0.33 years in the average age. This effect remained statistically significant in the first period (2011-2015) and the second period (2016-2020), with estimated increases of 0.35 and 0.28 years, respectively. These findings suggest that counties where older, more experienced farmers continue farming tend to have higher average ages among farm heads.

As for regional characteristics, the urban area ratio—measuring the proportion of urban land within a county—showed a negative relationship with average age at the baseline. That is, more urbanized counties tended to have younger farm heads. However, this relationship was not statistically significant, indicating that the level of urbanization alone may not be a decisive factor in explaining variation in average age at the county level.

We now turn to the estimation results for the second outcome variable: the number of farm households, as shown in Table 3. The treatment variable—per capita exemption amount—shows a positive effect on the number of farm households across all time periods. At the baseline (2010), an increase of one million KRW in per capita exemption was associated with an estimated increase of 7.09 households at the county level. In the first period (2011-2015), the estimated effect was 0.96 households, and in the second period (2016-2020), it was 0.41 households. All estimates were statistically significant.

These results suggest that the policy of gift tax exemption consistently contributed to maintaining or increasing the number of farm households throughout the study period. However, the size of the effect has gradually declined over time, indicating a diminishing marginal impact of the exemption policy on farm household retention or formation as time progresses.

Table 3. Estimated Results by Period: Dependent Variable

– Number of Farm Households(Y_{i2})

Category	Variable	Baseline (2010) Estimate (s.e.)	Period 1(2011–15) Estimate (s.e.)	Period 2(2016–20) Estimate (s.e.)
Outcome Variable 2 (Number of Farm Households)	Intercept	5451.849 **** (1097.966)	3991.587 **** (987.5539)	6654.0923 **** (968.1170)
	Per Capita Exemption Amount	7.09163 ** (3.5063)	0.9642 ** (0.4071)	0.4101 * (0.2192)
	Average Market Price of Farmland	-7.49526 *** (2.6916)	-3.7794 *** (1.2531)	-5.1509 **** (1.307)
	Average Agricultural Income	140.4296 *** (52.6436)	66.0522 * (34.6909)	48.4886 * (26.3643)
	Average Farming Experience	62.9877 * (35.3564)	88.3056 ** (39.5883)	-13.1219 (39.4560)
	Total Cultivated Area	0.5980 **** (0.05765)	0.6067 **** (0.06263)	0.7074 **** (0.0562)
	Average Cultivated Area per Household	-5658.415 **** (819.6957)	-4911.1362 **** (675.9651)	-4643.9168 **** (594.6565)
	Urban Area Ratio	-1149.015 * (675.9870)	-1716.3296 **** (455.1300)	-1169.1881 ** (582.0593)
	Level of Farm Mechanization	-731.9759 (737.8942)	-562.7768 (513.8046)	-601.7355 (484.8584)
	Agricultural Employment Level	-1592.203 (1236.244)	-345.4245 (1022.411)	-144.1501 (214.3226)
	Proportion of Greenhouse Farming Area	1032.960 (2345.521)	1225.3899 (1799.825)	471.2364 (1725.527)

Notes:

*, **, ***, and **** indicate statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively. Numbers in parentheses represent robust standard errors. The time periods are defined as follows: Baseline = 2010; Period 1 = 2011–2015; Period 2 = 2016–2020.

The analysis also examined the effects of key covariates on the number of farm households. Notably, the average farmland transaction price was found to have a statistically significant negative effect across all three time periods. Specifically, a one-million KRW increase in average land price was associated with a decrease of approximately 7.5 households in the baseline period, 3.8 households in the first period (2011–2015), and 5.2 households in the second period (2016–2020). These results suggest that higher land prices may act as a barrier to entry for new farmers or as a push factor for existing farms to exit the sector, especially in regions with active land markets.

Similarly, average agricultural income showed a significant and positive association with the number of farm households. At the baseline, a one-million KRW increase in income corresponded to an increase of approximately 140.4 households, while in the first and second periods the estimated effects were 66.1 and 48.5 households, respectively. Although still statistically significant in the second period (at the 10% level), the declining magnitude of the effect suggests that the initial boost in income may have already driven much of the entry and retention, and that the marginal effect of further income growth has since diminished.

The effect of land area on farm household numbers displayed a dual pattern. On one hand, an increase of one hectare in total farmland was associated with a statistically significant increase of 0.6 to 0.7 households, indicating that greater land availability facilitates agricultural activity and household retention. On the other hand, an increase of one hectare in average farmland per household was associated with sharp declines in the number of farm households—estimated at 5,658, 4,911, and 4,643 fewer households in the baseline, first, and second periods, respectively. This likely reflects a consolidation process, whereby small farms are absorbed into larger-scale operations, resulting in a reduction in the overall number of farming units.

In summary, increases in per capita exemption amounts were consistently associated with increases in the number of farm households. However, the magnitude of this effect declined over time, suggesting a diminishing marginal impact of the policy. Across other covariates, average farming experience was positively associated with average age, agricultural income positively influenced farm household numbers (though the effect weakened over time), and farmland variables showed a complex pattern—total land area supported household growth, while land concentration per household reduced it. Taken together, these findings imply that the gift tax exemption policy had a generally positive impact on generational turnover and the maintenance of farm households, though its effectiveness appears to have waned over time.

(2) Causal Effect Estimation via Comparison between the Reference and Comparison Periods

To more rigorously estimate the policy effect, we extended the baseline CBGPS-weighted fixed effects (FE-OLS) model by incorporating interaction terms between treatment intensity and time periods. Specifically, we introduced dummy variables to differentiate between the baseline period (2010, D=0) and the comparison period (2016–2020, D=1), and interacted them with the treatment variable—per capita exemption amount.

In this specification, the main coefficient on the treatment variable represents the baseline effect of tax exemption on outcomes (average age, number of farm households). The interaction term between treatment and the comparison period captures the additional change observed in high-treatment areas during the latter period and can be interpreted as the average treatment effect on the treated (ATT) in dynamic form.

We also included interaction terms between key covariates and the time dummies, based on earlier findings that some covariates exhibited period-specific effects. This allowed for a more comprehensive assessment of the dynamic policy impact. While the cross-sectional time-series models discussed earlier may have overestimated treatment effects due to time-invariant unobservables and shared shocks, the difference-in-differences (DID) approach mitigates this concern by differencing out such fixed effects. As a result, the magnitude of the estimated coefficients in the DID specification may be smaller but reflect a purer estimate of post-treatment change.

To further refine the policy effect estimation, we extended the CBGPS-weighted fixed effects linear regression (FE-OLS) model by incorporating both a period dummy and its interaction with the continuous treatment variable. In this specification, the first period (2011–2015) serves as the reference period (D=0), and the second period (2016–2020) is designated as the comparison period (D=1). The coefficient on the per capita tax exemption variable captures its baseline effect on the outcome variables—average age of farm heads and the number of farm households—during the reference period.

The interaction term between the tax exemption variable and the period dummy represents the incremental average treatment effect on the treated (ATT) observed in high-exemption regions during the comparison period. In other words, the sum of the main effect and the interaction term reflects the net treatment effect observed in the later period, conditional on the exemption level.

To capture additional sources of treatment heterogeneity, we also introduced interaction terms between time and covariates that had shown statistical significance in previous analyses. This allowed for a more dynamic and comprehensive discussion of how the policy effect evolved over time depending on regional characteristics.

It is important to note that earlier cross-sectional time-series models may have overstated treatment effects by conflating them with time-invariant regional characteristics and common shocks. By contrast, the DID approach controls for such unobserved fixed effects, isolating the net post-treatment change. While this leads to more conservative (and potentially smaller) coefficient estimates, it enhances the credibility and causal validity of the findings.

Table 4. Estimation and Comparison of Outcome Variables between the Reference and Comparison Periods

Category	Variable	Outcome Variable 1 (Average Age of Farm Household Heads)	Outcome Variable 2 (Number of Farm Households)
Control Variables	Per Capita Exemption Amount	-0.000694 * (0.000353)	-0.1388240 (0.4675)
	Average Market Price of Farmland	0.001096 (0.002149)	1.010566 (1.5307)
	Average Agricultural Income	0.2336 *** (0.0829)	-261.7531 *** (90.0104)
	Average Farming Experience	0.334364 **** (0.036458)	-158.3074 **** (42.96680)
	Average Cultivated Area per Farm	-1.464858 *** (0.448806)	-3059.884 **** (728.7245)
	Urban Area Ratio	6.202741 (9.377590)	4135.475 (15383.21)
	Number of Traded Land Parcels	-0.000031 (0.000027)	-0.208559 **** (0.050995)
	Farm Mechanization Ratio	-1.793879 *** (0.459551)	-1652.055 ** (699.7594)
	Agricultural Employment Ratio	0.039168 (0.054053)	13.09849 (95.66662)
	Number of Gifted Parcels	0.000601 (0.000741)	-0.471864 (1.314147)
	Exemption Amount × Period Dummy	0.000695 *** (0.000237)	-0.07518000 (0.3112300)
	Avg. Agricultural Income × Avg. Farming Experience	-70.405574 *** (25.4333)	102042.3 *** (30853.59)
	Urban Area Ratio × Number of Traded Land Parcels	-0.000362 **** (0.000092)	0.15806 (0.172576)

Notes:
*, **, ***, and **** indicate statistical significance at the 10%, 5%, 1%, and 0.1% levels, respectively. Numbers in parentheses represent robust standard errors. The time periods are defined as follows: Baseline = 2010; Period 1 = 2011–2015; Period 2 = 2016–2020.

The interaction term between the per capita tax exemption and the comparison period dummy yields a positive and statistically significant coefficient for the outcome variable of average age. This implies that the marginal effect of the exemption amount on reducing average age became smaller during the comparison period (2016-2020) than in the reference period (2011-2015). Specifically, while the base trend indicates a decrease of approximately 0.000694 years in average age per additional one million KRW in exemption, the interaction term nearly offsets this effect, suggesting that the policy's marginal effectiveness in lowering the average age has diminished over time. In contrast, for the number of farm households, no statistically significant change was observed across periods.

The interaction between average agricultural income and average farming experience also produced statistically significant effects for both outcome variables, though these require more careful interpretation. These terms do not indicate the effect of simultaneous one-unit increases in both variables. Instead, they show how the marginal effect of one variable changes depending on the level of the other. For instance, the marginal effect of a one-million KRW increase in agricultural income depends on the value of average farming experience (A), and is calculated as $0.2336-70.4\times A$. Conversely, a one-year increase in farming experience yields a marginal effect of $-158+102042\times(\text{avg. income})$. The results suggest that counties with both higher income and greater farming experience tend to have younger farm heads and larger numbers of farm households, possibly indicating successful generational turnover via successor farm entries in those areas.

Additional significant interactions were found between the urban area ratio and the number of land parcel transactions. These were relevant only for the average age outcome. More urbanized areas with a greater number of land transactions tended to experience further declines in average age during the comparison period, implying that active land markets in urban regions may facilitate the entry of younger farmers.

Among the control variables, average farmland area per household showed particularly strong and consistent effects. Controlling for treatment intensity and period effects, a one-hectare increase in farmland per household was associated with a 1.4-year decrease in average age and a reduction of approximately 3,059 farm households. This finding suggests that while larger-scale farming operations are often maintained by younger generations, they may also be accompanied by structural consolidation—absorbing smaller farms and thereby reducing the total number of farm units.

(3) Dose–Response Function and Scenario-Based Policy Simulation

The dose-response function (DRF), estimated using the CBGPS methodology, offers a visual and analytical representation of how changes in the continuous treatment variable—per capita exemption amount—affect key outcome variables such as average age and number of farm households. This function enables the identification of the policy's marginal effect across varying treatment intensities, beyond simple linear assumptions.

The DRF can also be applied in simulation analyses for policy prediction. By inputting a hypothetical level of tax exemption into the function, one can forecast its corresponding effect on county-level outcomes, assuming a constant structural framework. For instance, given a projected exemption amount, the DRF can estimate how much change in average age or farm household count would be induced, enabling evidence-based policy design and ex-ante evaluation of potential reforms.

Figure 1. Dose-Response Function
between Outcome Variables and Exemption Amount

Outcome Variable 1 (Average Age of Farm Household Heads)	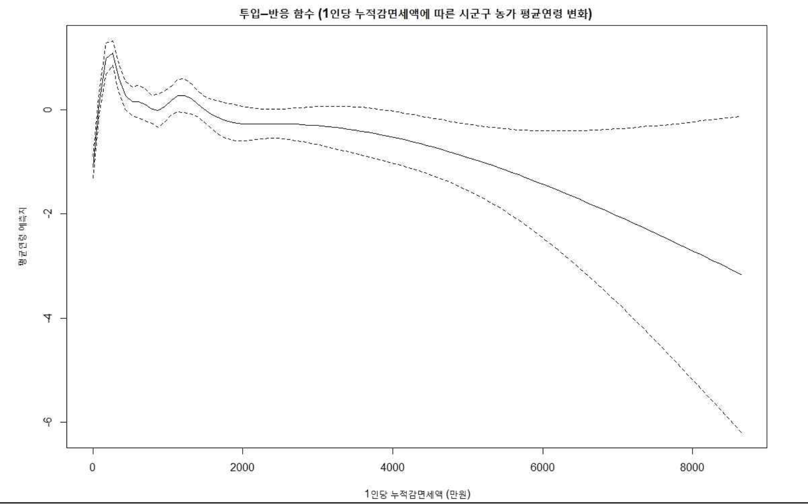

Outcome Variable 2 (Number of Farm Households)	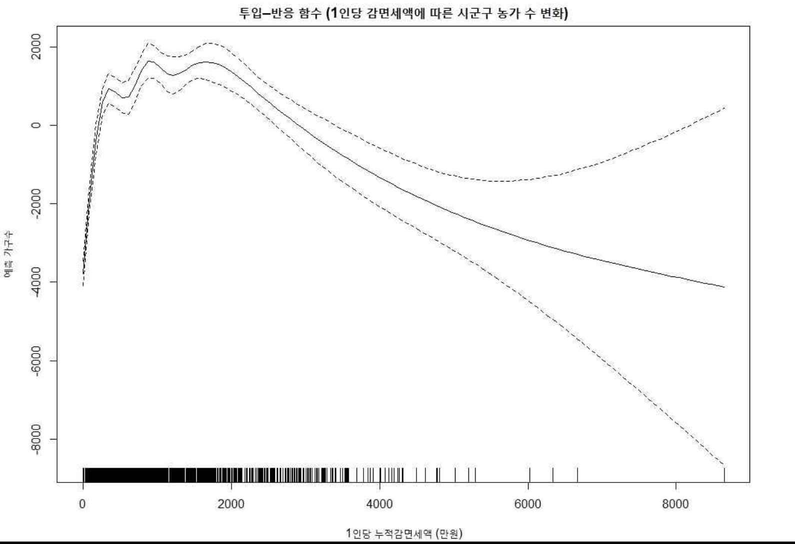

To enable a more precise estimation of the policy's dynamic effects, the original period-based dataset was interpolated into annual units. In addition, administrative data on gift tax exemption amounts from 2021 to 2023 were incorporated. This extended panel allows for the analysis of how the cumulative per capita exemption amount—measured annually—affects relative changes in average age and the number of farm households at the county level. More importantly, it facilitates a forward-looking assessment of the extent to which continued increases in exemption amounts might induce demographic or structural transitions in the agricultural sector.

The figure below presents the dose-response function estimated for the outcome variable of average age. The horizontal axis represents the annual per capita cumulative exemption amount (in 10,000 KRW units), while the vertical axis indicates the relative change in the average age of farm household heads, compared to the baseline where the exemption amount is zero. The curve reflects the estimated trajectory of policy impact, and the dashed lines denote the 95% confidence interval.

This visualization reveals how increases in exemption levels are associated with relative declines in average age, particularly in the lower to moderate range of the exemption spectrum. The slope of the curve flattens as exemption levels rise, suggesting diminishing marginal effects at higher exemption thresholds. These findings provide additional insight into the nonlinear and possibly saturating nature of the policy's demographic effects.

The estimated dose-response function for average age reveals a distinctly nonlinear pattern. In the initial segment—from 0 to approximately 20 million KRW—the curve exhibits a positive slope, indicating that the average age of farm heads actually increases with modest levels of exemption. Specifically, counties where the per capita exemption amount rises from zero to around 20 million KRW experience an average increase of about 1.1 years in farm head age compared to the baseline. This suggests that in low-exemption areas, the policy benefits may have primarily accrued to existing older-generation farmers, rather than facilitating the entry of younger successor farmers. In these cases, the tax benefits likely supported inter vivos transfers among elderly landholders, resulting in a short-term rise in average age.

After reaching a local maximum at approximately 20 million KRW, the average age begins to decline steadily as the exemption amount increases further. By around 60 million KRW, the average age has dropped by approximately 2 years compared to the baseline. This turning point marks a transition wherein the policy may begin to take effect through the entry of younger successor farmers or the retirement of aging farm heads. At even higher levels of exemption—exceeding 80 million KRW—the decline in average age becomes more pronounced, with reductions reaching nearly 4 years. This indicates that large-scale exemptions exert stronger demographic impacts and may promote structural renewal in the agricultural labor force.

These results highlight a nonlinear, U-shaped effect of the tax exemption policy on average age, with an initial unintended consequence of aging, followed by an inflection point and eventual rejuvenation at higher exemption levels. The presence of such reversal effects underscores the importance of calibrating eligibility thresholds and targeting mechanisms within the policy. It also suggests that moderate-to-high levels of exemption are necessary to realize the intended generational turnover effects.

The lower panel of the figure presents the corresponding dose-response function for the number of farm households. The vertical axis reflects relative change in household count compared to the baseline where exemption is zero. In the low-exemption range (up to approximately 20 million KRW), the number of households increases significantly—by as many as 1,000 additional households—suggesting that even modest fiscal incentives may help retain existing farms or attract new entrants. This supports the view that small-scale exemptions act as policy signals, helping sustain farming communities by offsetting exit trends.

Beyond the 20 million KRW threshold, however, the number of farm households begins to decline sharply. By the time exemptions reach around 60 million KRW, household numbers fall by up to 4,000. This decline likely reflects two processes: (1) policy concentration on a small number of large farms receiving high exemptions, and (2) complete generational transfer accompanied by the retirement of the parent generation, which reduces the total count of independently registered farms. Where these dynamics coincide with broader trends of population decline, the reduction in farm household numbers may be especially pronounced.

Notably, the decline plateaus beyond 60 million KRW, stabilizing at approximately 4,200 fewer households. This suggests a saturation point where additional exemptions no longer induce further structural shifts—indicating diminishing marginal effects of the policy. In short, while modest exemptions are associated with farm retention and possibly entry, high-value exemptions correspond more closely to consolidation and farm exit dynamics.

These findings confirm that the response of farm household counts to tax exemptions is nonlinear and asymmetric. Small exemptions yield positive retention effects, while large exemptions may induce restructuring or reduction in farm numbers. This has direct implications for the design and targeting of the policy: to balance sustainability and equity, exemption caps and eligibility rules should be aligned with regional demographic and structural conditions.

6. Conclusion

This study evaluated the effectiveness of South Korea's gift tax exemption policy for successor farmers by employing a three-stage analytical framework that integrates continuous treatment intensity with repeated panel observations. First, covariate balance was achieved using the Covariate Balancing Generalized Propensity Score (CBGPS), minimizing potential endogeneity bias. Second, a multiple-period Difference-in-Differences (DID) approach was applied to the three observation points (2010, 2015, and 2020) to capture the policy's cumulative and dynamic effects. Third, a doubly robust re-estimation of the DID model using GPS-based weights was conducted to assess treatment heterogeneity and distributional equity.

The analysis found that the marginal effect of the per capita exemption amount on the average age of farm household heads weakened over time, indicating a diminishing impact of the policy across the study periods. In contrast, the effect of the exemption on the number of farm households remained consistently positive, though its magnitude also declined across successive periods. These patterns suggest that the marginal returns to additional exemption amounts may taper off as the policy matures—a dynamic characteristic that holds significant implications for future policy adjustments.

It is important to note that this study relied on aggregated data at the county (Si/Gun/Gu) level due to limitations in data availability. As a result, individual-level heterogeneity among farms could not be fully accounted for, and the analysis was necessarily centered on representative averages rather than micro-level variation. Despite this inherent limitation, the findings contribute to a more nuanced understanding of how tax incentives shape demographic and structural transitions in agriculture, and offer empirical guidance for the refinement of successor-focused agricultural policy.

References

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of Educational Psychology, 66(5), 688–701. https://doi.org/10.1037/h0037350

Bruhn & McKenzie (2009). In Pursuit of Balance: Randomization in Practice in Development Field Experiments, American Economic Journal: Applied Economics, 1(4), 200–232.

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. https://doi.org/10.1093/biomet/70.1.41

Hirano, K., & Imbens, G. W. (2004). The propensity score with continuous treatments. Retrieved from https://www.math.mcgill.ca/dstephens/PSMMA/Articles/HIrano-Imbens-2004.pdf

Fong, C., Hazlett, C., & Imai, K. (2018). Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. Annals of Applied Statistics, 12(1), 156–177. https://doi.org/10.1214/17-AOAS1101

Egami, N., & Yamauchi, S. (2022). Using multiple pretreatment periods to improve difference-in-differences and staggered adoption designs. Political Analysis, 31(1), 195–212. https://doi.org/10.1017/pan.2022.8

Cowell, F. (2011).「Measuring inequality」, Oxford University Press.