Expanding Free Influenza Vaccination to Primary School Children: Evidence from Korea*

Inhyuk Hwang†

April 30, 2026

Abstract

School-aged children play a central role in seasonal influenza transmission, with high attack rates and frequent contacts that amplify community spread. Publicly funded vaccination programs for this age group, however, remain uncommon internationally. This paper evaluates Korea's 2018 expansion of the National Immunization Program, which extended free vaccination to all children ages 5–12. Using nationally representative survey data, I find that difference-in-differences estimates show a sharp and persistent increase in vaccination coverage among newly eligible children. To evaluate the expansion's effect on influenza-related healthcare utilization, I turn to administrative claims data. Because pre-expansion utilization trends differ across the two age groups, I adopt a triple-difference design that exploits variation in vaccine match quality across time, and find a meaningful reduction in influenza-related claims during high-match months. A back-of-the-envelope calculation indicates that direct healthcare-cost savings offset a substantial share of the program's budgetary cost.

JEL Codes: I12, I18, J13
Keywords: Influenza; Vaccination; Healthcare utilization; Children

---

*I am grateful to my advisors, Dae-il Kim and Sok Chul Hong, for their invaluable guidance and support throughout this project. I also thank Chulhee Lee, Wankyo Chung, and Kyeongbae Kim for their thoughtful comments and suggestions, as well as seminar participants at Seoul National University and the AASLE 2024 Conference. I gratefully acknowledge financial support from the National Research Foundation of Korea (NRF-2020S1A3A2A02104190), funded by the Ministry of Education of the Republic of Korea. This study uses the National Health Information Database (NHIS-2024-11-1-025) provided by the National Health Insurance Service (NHIS) under IRB exemption No. E2404/004-005. The author declares no conflict of interest with NHIS. All remaining errors are my own.

†Associate Research Fellow, National Assembly Futures Institute. Email: inhyuk@nafi.re.kr.

1 Introduction

Seasonal influenza is an acute respiratory illness that poses a major global public health challenge. Although influenza is mild for most healthy individuals, it can progress to severe complications such as pneumonia and bronchitis in vulnerable populations, particularly young children and the elderly (World Health Organization, 2025). Each year, influenza is estimated to infect roughly one billion people globally, resulting in three to five million cases of severe illness and 290,000 to 650,000 respiratory deaths (World Health Organization, 2025). Early-life exposure to influenza has also been linked to lasting socioeconomic costs, including reduced height, educational attainment, income, and employment (Almond, 2006; Lin and Liu, 2014; Kelly, 2011).

Despite this epidemiological rationale, publicly funded vaccination programs targeting school-aged children remain uncommon internationally (Goldin et al., 2024; Ortiz et al., 2016). Most countries limit free provision to infants, the elderly, and high-risk adults, leaving primary-school-aged children dependent on parental willingness to pay and on access to private clinics. Two open questions follow from this gap: whether extending free eligibility to school-aged children raises vaccination coverage, and whether any resulting coverage gains translate into meaningful reductions in influenza-related healthcare utilization. Both questions have direct implications for program design.

This paper answers these questions by evaluating a sharp expansion of Korea's National Immunization Program (NIP). The NIP expanded its pediatric influenza coverage in stages. Free vaccination was first offered to infants aged 6–12 months in the 2016–2017 season, extended to children under age 5 in 2017–2018, and then extended to all children ages 5 through 12 beginning in 2018–2019 (hereafter, the 2018 expansion). After the 2018 expansion, eligible children could receive free vaccination at public health centers and through a large network of contracted private clinics. Because the 2018 expansion was announced and implemented as a single national policy, it produced a discrete shift in eligibility at age 12

and left adolescents aged 13–18 as an ineligible comparison group.

I draw on two datasets, one for each outcome of interest. Vaccination status is measured using the Korea National Health and Nutrition Examination Survey (KNHANES), a nationally representative survey. Influenza-related healthcare utilization is measured using the National Health Information Database, maintained by the National Health Insurance Service (NHIS). Because NHIS operates Korea's single-payer system, its claims data cover virtually the entire population.

My identification strategy centers on the policy change but adapts to threats to identification that differ across the two outcomes. For vaccination coverage, I estimate a difference-in-differences specification that compares the newly eligible treatment group (ages 5–12) to the ineligible control group (ages 13–18) around the 2018 expansion. Age and flu-year fixed effects absorb time-invariant age differences and age-invariant year shocks. For influenza-related healthcare utilization, pre-expansion differences in utilization trends between younger and older children, driven by age-specific incidence and contact patterns, preclude a simple difference-in-differences design. I therefore adopt a triple-difference (DDD) strategy that exploits across-time variation in the influenza vaccine match rate as a third source of identifying variation. Because the health benefits of vaccination are larger when the circulating strain matches the vaccine composition more closely, the expansion's effect is identified from the interaction of three differences: eligible versus ineligible ages, post- versus pre-expansion periods, and high-match versus low-match periods. This DDD design controls for age-group-specific time trends nonparametrically and identifies the expansion's health effect from differential vaccine effectiveness across high- and low-match months (Ward, 2014; White, 2021).

I find that the 2018 expansion produced a sharp and persistent increase in pediatric influenza vaccination. The preferred difference-in-differences specification indicates that vaccination coverage among children ages 5–12 rose by 11.7 percentage points relative to the ages 13–18 comparison group—roughly a 20 percent increase over the pre-expansion treatment-

group mean of 57.4 percent. Dynamic DiD estimates show no differential pre-trends and a discrete jump in the 2018 expansion year that persists into the following flu-year; the coverage effect is also stable across robustness checks that address interview timing, age-based contamination, within-household spillovers, and the early COVID-19 period. Heterogeneity analysis indicates that the estimated gains are larger among higher-income children than among lower-income children, though the difference is not statistically significant.

For healthcare utilization, the DDD estimate implies a reduction of 13.3 influenza-related claims per 1,000 among newly eligible children during high-match months. I interpret this estimate as a lower bound on the expansion's intent-to-treat (ITT) effect on this group—an estimand that combines direct protection with within-group indirect protection. Two forces attenuate the estimate: the vaccine retains partial protection in low-match months, narrowing the high-vs-low contrast, and any within-household or within-school spillovers to the ineligible comparison group would further compress the difference. The DDD design therefore captures only part of the program's community benefit. A back-of-the-envelope calculation suggests that direct healthcare-cost savings offset roughly three-quarters of the program's budgetary cost. Once standard indirect benefits—such as averted caregiver work loss and within-school externalities—are included, the net return would likely turn positive.

This paper contributes to two strands of the literature. First, I add to a growing body of work on how public policy shapes vaccine uptake. Prior research has examined school-entry mandates (Abrevaya and Mulligan, 2011; Lawler, 2017; Carpenter and Lawler, 2019), free-provision programs and price reductions (Bouckaert, Gielen, and Van Ourti, 2020; Humlum, Morthorst, and Thingholm, 2024), and monetary incentives (Campos-Mercade et al., 2021; Schneider, 2023). Much of this evidence, however, concerns non-influenza vaccines such as varicella, pertussis, HPV, and COVID-19, or addresses adults and the elderly. Causal evidence on price interventions for flu vaccination among school-aged children remains thin, despite this group's central role in community transmission. The present study addresses this gap by evaluating a national program that extended free influenza vaccination to all

primary-school children in South Korea.

Second, I contribute to the literature on the population-level health effects of influenza vaccination. A large epidemiological literature estimates individual-level vaccine efficacy and effectiveness from randomized trials and test-negative designs (Osterholm et al., 2012; Choi et al., 2024). These designs recover direct protection well but are poorly suited to measuring indirect protection within a population, because randomization at the individual level removes the spillovers that drive community-level gains. Ward (2014) addresses this limitation with a vaccine-match-rate triple-difference design applied to Ontario's universal adult expansion; White (2021) builds on this approach to quantify the mortality and labor-supply externalities of influenza vaccination and shows that most of the total social value is external rather than private. I adopt the match-rate identification strategy in a pediatric expansion targeted at ages 5 through 12 rather than at adults. My DDD estimates deliver an intent-to-treat effect on influenza-related utilization among the eligible age group, reflecting the combined direct and within-group indirect protection.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of Korea's National Immunization Program. Section 3 introduces the data sources and sample construction. Section 4 presents the empirical strategy and results for vaccination coverage. Section 5 presents the strategy and results for influenza-related healthcare utilization. Section 6 concludes.

2 Institutional Background

Before the 2016 flu season, Korea's National Immunization Program covered only infants aged 6 to 12 months. In response to the high disease burden among children and their role in community transmission, the government expanded eligibility in stages. In the 2017 flu season, coverage extended to children up to 59 months. A larger expansion followed in 2018, covering all primary school students up to age 12. Eligible children can receive the influenza

vaccine free of charge at public health centers and designated medical institutions.[1]

To support the expanded program, the Korea Disease Control and Prevention Agency (KDCA), which administers the National Immunization Program, launched a nationwide campaign to promote pediatric influenza vaccination in the 2018 flu season. The KDCA designated an intensive vaccination week in early October to coincide with the start of the vaccination period and to encourage early, widespread uptake. During this week, selected public health centers offered weekend vaccination services.[2] The KDCA also collaborated with the Ministry of Education to disseminate vaccination information through schools and kindergartens, distributing official notices to parents and caregivers. In the 2019 flu season, the collaboration expanded to include text message reminders about vaccination schedules, sent to approximately 520,000 upper-grade primary school students (ages 10–12) before November.

This study evaluates the effect of the 2018 expansion, together with the accompanying information campaign, on vaccine uptake and influenza-related healthcare utilization among children aged 5 to 12.

3 Data

3.1 Korea National Health and Nutrition Examination Survey

To evaluate the effect of the 2018 expansion on influenza vaccination, I use the academic-research microdata release of KNHANES covering 2014 to 2020. I reorganize the data into flu-years using the interview date. In Korea, influenza vaccination typically begins in

1. The number of eligible children increased from 316,746 in the 2016–2017 flu season to approximately 2.14 million in 2017–2018 and 5.74 million in 2018–2019 (Korea Centers for Disease Control and Prevention, 2017; Korea Centers for Disease Control and Prevention, 2019). Over the same period, the number of designated medical institutions increased from 6,701 in 2017–2018 to 8,879 in 2018–2019 and 9,264 in 2019–2020 (Korea Centers for Disease Control and Prevention, 2020).

2. Weekend vaccination services were provided primarily in selected underserved areas rather than nationwide.

September and the epidemic season runs from October to April of the following year. I therefore define a flu-year as the period from September of calendar year t through August of year $t+1$.

I measure influenza vaccination status using the survey question "Have you been vaccinated against influenza within the past year?" This retrospective measure introduces a reference-window issue. Because influenza vaccinations are typically administered in the fall and winter, respondents interviewed shortly after the expansion may report vaccination behavior that occurred before the expansion yet still be coded as treated. In particular, interviews conducted in September and October 2018 may attenuate the estimated effect. I address this concern in the robustness checks.

Beyond the individual vaccination measure, KNHANES offers a feature central to my heterogeneity analysis: it surveys all household members aged one year and older. I therefore observe not only the children directly affected by the expansion but also their mothers, who make most pediatric healthcare decisions in Korea. I exploit this household linkage to examine heterogeneity in the expansion's effect.

Figure 1 plots influenza vaccination rates before and after the expansion for the treatment group (ages 5–12) and the control group (ages 13–18). During the pre-expansion period (2014–2017), vaccination rates in both groups follow a similar upward trend; although the levels differ, the gap between the two groups remains stable. After the expansion in 2018, vaccination rates in the treatment group increase sharply, while the control group continues its gradual upward trend. These patterns provide visual support for a difference-in-differences approach. As shown in Appendix Table A3, however, the repeated cross-sectional nature of the data produces changes in sample composition before and after the expansion. I address this issue in Section 4.1.

Figure 1: Trends in Influenza Vaccination Rates by Age Group

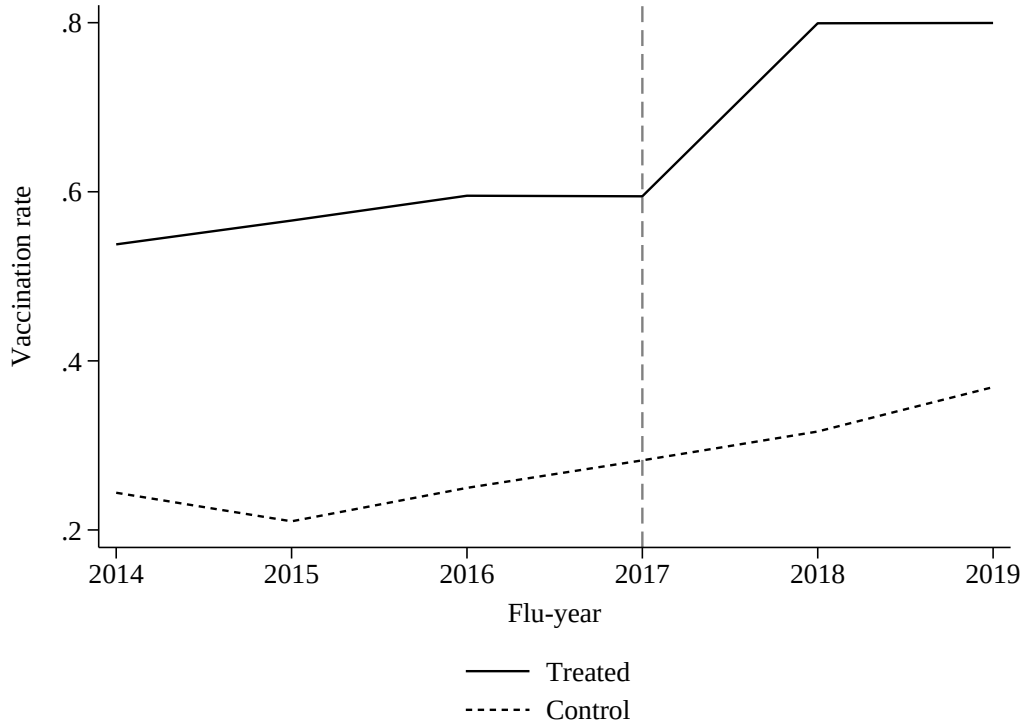

Notes: Figure 1 plots annual influenza vaccination rates by flu-year for the treatment and control groups. The vertical dashed line marks the flu-year immediately preceding the 2018 expansion. The treatment group consists of children aged 5–12 who became eligible under the 2018 expansion, and the control group includes adolescents aged 13–18 who were not affected by the eligibility change. A flu-year is defined as September of year t through August of year $t + 1$. Vaccination rates are calculated using KNHANES survey weights.

3.2 National Health Information Database

To examine influenza-related healthcare utilization, I use the National Health Information Database, maintained by NHIS. As Korea's single-payer insurer, NHIS holds eligibility and claims records that cover virtually the entire population. The database allows researchers to specify a target population, sampling design, and study period, and it returns the corresponding eligibility and claim-level records.

My extract draws from all NHIS enrollees aged 23 or younger as of 2017. Using sex- and age-stratified random sampling, I obtain approximately 27 percent of this target population, or roughly 3.4 million individuals. The extract contains two linked files: the eligibility file records each enrollee's sex, age, insurance type, and monthly insurance premium; the claims file records, at the individual claim level, ICD-10 diagnosis codes, service details, and healthcare expenditures. From this extract I construct an analysis sample comprising children and adolescents aged 5 to 18, observed over flu-years 2014 through 2018.

Observations are indexed by age, flu-year, and calendar month. This structure follows from my identification strategy, which uses month-to-month variation in the influenza vaccine match rate within flu-years. I construct two primary outcomes. The first, influenza-related claims, is the count of claims recording ICD-10 diagnosis codes J09, J10, or J11 (including all subcodes), standardized as a rate per 1,000 population; the denominator is the number of individuals in the corresponding age by flu-year cell of the analysis sample. The second outcome, influenza-related healthcare expenditures, sums NHIS benefit payments and patient out-of-pocket costs on claims identified by the same diagnostic criteria, also standardized per 1,000 population.[3]

3. Claims data record influenza episodes only when a patient visits a medical institution and receives a clinical diagnosis. The diagnosis is typically supported by rapid antigen testing, the most widely used method in Korean outpatient practice; RT-PCR and rapid molecular assays are used less frequently (Lee, 2023). Because mild influenza cases often resolve without medical contact, claims-based measures understate the true incidence of infection. Under-reporting of levels is not, on its own, a threat to my research design: identification relies on relative changes between the treatment and control groups, so a common multiplicative under-reporting factor differences out.

4 Effects of the Coverage Expansion on Influenza Vaccine Uptake

4.1 Empirical Approach

To estimate the effect of the 2018 expansion, I use a standard difference-in-differences framework:

$$Y_i = \beta_0 + \beta_1 Treat_i + \beta_2 Post_t + \beta_3 (Treat_i \times Post_t) + \gamma' X_i + \epsilon_i \tag{1}$$

where Y_i equals one if individual i received an influenza vaccination during the corresponding flu-year and zero otherwise. $Treat_i$ equals one if the individual belongs to the treated age group. The treatment group consists of children aged 5–12 who became newly eligible for free vaccination under the 2018 expansion; the control group consists of adolescents aged 13–18, whom the eligibility change did not directly affect. $Post_t$ equals one if the individual was surveyed in the post-expansion period, which I define as flu-years 2018/19 and 2019/20. X_i is a vector of child, household, and maternal characteristics: age, sex, and subjective health status for the child; household size, income, type of health insurance, and private health insurance coverage for the household; and age, college education, influenza vaccination status, and employment status for the mother. Detailed variable definitions appear in Appendix Table A2. ϵ_i is the error term, and standard errors account for the complex sampling design of KNHANES.

The coefficient of interest, β_3, measures the differential change in vaccination probability attributable to the expansion under two identifying assumptions. First, the parallel-trends assumption requires that, absent the expansion, vaccination rates in the treatment and control groups would have evolved similarly over time. Second, because the analysis uses repeated cross-sectional data, identification also requires no systematic compositional change at the time of the expansion (Hong, 2013): the distribution of observable and unobservable

characteristics within each age group should not shift in a way that differentially influences vaccination behavior across groups.

To assess the plausibility of parallel trends, I adopt two complementary approaches. First, I include group-specific linear flu-year trends, which yield a more conservative estimate of the expansion effect. Second, I estimate a dynamic difference-in-differences (dynamic DiD) specification:

$$Y_i = \beta_0 + \beta_1 Treat_i + \sum_{k \in K} \beta_2^k (Treat_i \times \mathbf{1}[Flu\text{-}year_t = k]) + \gamma' X_i + \delta_t + \epsilon_i \tag{2}$$

where $K = \{2014, 2015, 2016, 2018, 2019\}$, and the omitted category is the 2017 flu-year immediately preceding the expansion. δ_t denotes flu-year fixed effects, and the control variables match those in equation (1).

On compositional stability, the increase in vaccination rates for the treatment group after the expansion is abrupt rather than gradual, which makes it unlikely that the estimated effect reflects slow-moving compositional shifts within the age group. The summary statistics do show some changes in sample composition across periods, however (Appendix Table A3). To mitigate this concern, I control for the rich set of child, household, and maternal characteristics described above. Although this does not rule out changes in unobservables, it reduces bias from observable compositional differences.

4.2 Baseline Results

Table 1 reports estimates from equation (1). Column (1) shows the baseline without additional controls: the expansion increased vaccination rates among children aged 5–12 by 12.7 percentage points relative to the control group. Columns (2) through (4) sequentially add controls for child, household, and maternal characteristics. The estimated treatment effect remains stable across specifications, indicating that the results are not driven by observable compositional differences. Column (5) further includes group-specific linear time

trends to account for differential pre-existing trends across age groups. The estimated coefficient remains at 11.7 percentage points. Taking Column (4) as the preferred specification, the estimated effect of 11.7 percentage points corresponds to approximately a 20.4 percent increase over the pre-expansion vaccination rate of 57.4 percent among treated children.

Table 1: Effect of the 2018 Expansion on Influenza Vaccine Uptake

	(1)	(2)	(3)	(4)	(5)
$Treat \times Post$	0.127***	0.127***	0.130***	0.117***	0.117**
	(0.030)	(0.029)	(0.029)	(0.026)	(0.046)
Controls:					
Child		Y	Y	Y	Y
Household			Y	Y	Y
Maternal				Y	Y
Group-specific time trend					Y
Adjusted R-squared	0.166	0.196	0.204	0.376	0.377
Observation	6699	6699	6699	6699	6699

Notes: Table 1 reports estimates from equation (1), which evaluates the effect of the 2018 expansion on influenza vaccination among children. The dependent variable is an indicator equal to 1 if the child received an influenza vaccination in a given flu-year and 0 otherwise. $Treat \times Post$, the key regressor, equals 1 for children aged 5–12 in the post-expansion period (flu-years 2018/19 and 2019/20) and 0 otherwise; its coefficient is the parameter of interest. The pre-expansion vaccination rate among the treatment group is 57.4 percent. Columns (2)–(4) sequentially add child, household, and maternal characteristics, and Column (5) further includes group-specific linear flu-year trends. Detailed variable definitions are reported in Appendix Table A2. All specifications apply KNHANES survey weights, and standard errors account for the survey's complex sampling design. A single asterisk denotes statistical significance at the 10 percent level, double at the 5 percent level, and triple at the 1 percent level.

4.3 Robustness Checks

Figure 2 plots the dynamic DiD coefficients β_2^k from equation (2). In the pre-expansion years (flu-years 2014–2016), the estimated coefficients are small and statistically indistinguishable from zero across specifications, supporting the parallel trends assumption. In the 2018 flu-year the treatment effect jumps sharply and is both statistically significant and economically sizable, and the effect persists in the 2019 flu-year, indicating that the increase in vaccination uptake among newly eligible children was not transitory.

Figure 2: Dynamic DiD Estimates of the 2018 Expansion's Effect on Influenza Vaccine Uptake

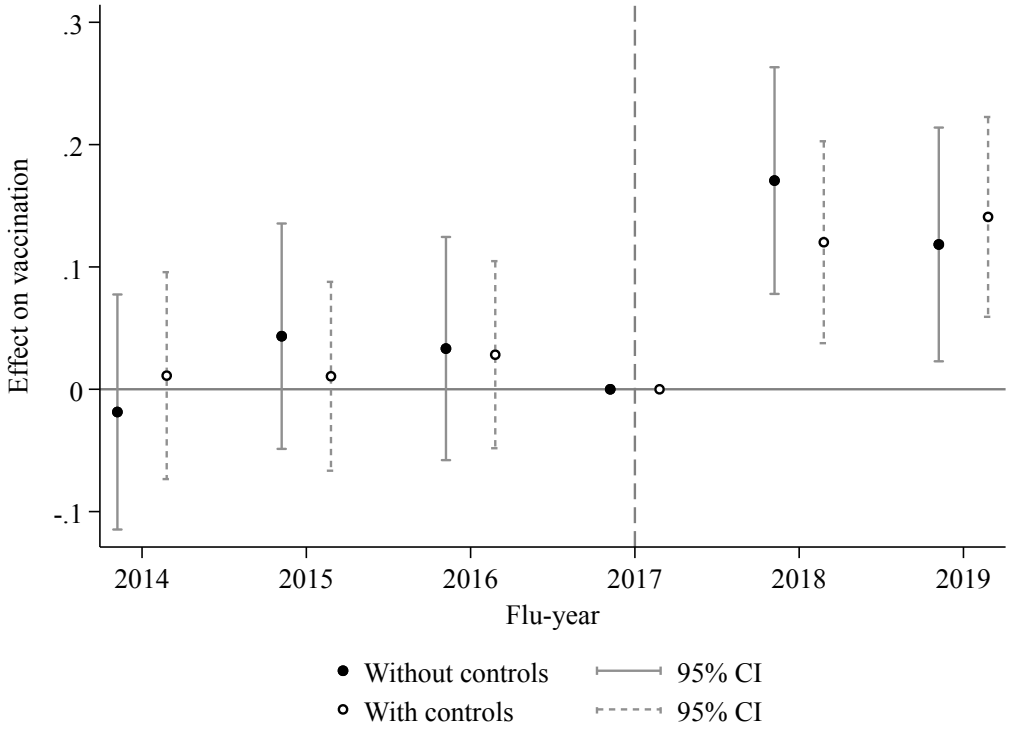

Notes: Figure 2 reports dynamic DiD estimates based on equation (2), assessing the effect of the 2018 expansion on influenza vaccination among children. The dependent variable is an indicator equal to 1 if the child received an influenza vaccination in a given flu-year and 0 otherwise. The coefficients shown are the interaction terms $Treat_i \times \mathbf{1}[Flu\text{-}year_t = k]$ between the treatment indicator (children aged 5–12) and flu-year dummies, with the 2017 flu-year as the omitted reference period. Filled circles denote estimates without control variables, and hollow circles denote estimates including the full set of controls used in Column (4) of Table 1. Vertical lines represent 95 percent confidence intervals. All specifications apply KNHANES survey weights, and standard errors account for the survey's complex sampling design.

Although the dynamic DiD estimates support the identifying assumption, three threats to the baseline estimate remain: interview timing, contamination from cohort transitions and household spillovers, and the onset of the COVID-19 pandemic. Table 2 reports robustness checks targeting each threat in turn.

First, the outcome variable comes from the survey question "Have you received an influenza vaccination within the past year?" Because influenza vaccinations are primarily administered between September and December, differences in interview timing may introduce attenuation bias (Bitler and Carpenter, 2016; White, 2021). In particular, individuals interviewed between September and December 2018 may report vaccination behavior from before the expansion yet still be classified as treated. To address this reference-window concern, Column (2) excludes observations interviewed during this period. Excluding these observations raises the estimated coefficient from 0.117 in Column (1) to 0.123, consistent with modest attenuation in the baseline. Column (3) further partitions the post-expansion period into two subperiods: $Post1$ indicates interviews conducted between September and December 2018, and $Post2$ indicates interviews conducted from January 2019 through August 2020. The estimated coefficient on $Treat \times Post1$ is 0.085 and is marginally significant, while the coefficient on $Treat \times Post2$ is 0.123 and is significant at the 1 percent level. The pattern is consistent with attenuation concentrated in the months immediately after the expansion. Across these specifications, the qualitative conclusion is unchanged and the magnitudes remain close to the baseline.

Second, Columns (4)–(6) address potential contamination. Some children younger than age five in 2017 were already covered by the program and enter the treatment group as they age; excluding children aged 5 and 6 raises the estimated coefficient slightly to 0.127 in Column (4). Column (5) considers the reverse transition, in which children initially in the treatment group age into the control group over time; excluding age 13 yields a similar estimate of 0.127, modestly larger than the baseline. Column (6) examines within-household spillover effects. Children in the control group may have siblings in the treatment group, and

such exposure could bias the estimates.[4] Excluding control-group children with siblings in the treatment group slightly increases the estimated coefficient, though the change is small. Contamination from cohort transitions or household spillovers does not materially affect the results.

Column (7) examines whether the onset of the COVID-19 pandemic affects the estimates. The 2019 flu-year overlaps with the early phase of the pandemic. Influenza vaccinations, however, are primarily administered between September and December, before the peak epidemic season, so substantial pandemic-related distortion is unlikely. Excluding individuals surveyed between February and August 2020 nevertheless increases the estimated coefficient. This pattern reflects differences in vaccination timing across age groups: older children in the control group tend to receive vaccinations later in the flu-year, so excluding late-survey observations disproportionately removes vaccinated individuals from the control group and mechanically raises the estimated treatment effect.[5]

4.4 Additional Results

4.4.1 Heterogeneous Effects

Monetary cost is a well-documented barrier to influenza vaccination across populations, including the general public, children, and healthcare workers (Schmid et al., 2017). In the Korean context, survey evidence indicates that 64.8 percent of parents perceived influenza vaccination as expensive, with lower-income households more likely to report cost concerns

4. The direction of this bias is ambiguous. If parents view vaccination as unnecessary for non-eligible children, the baseline estimate may be overstated (Bouckaert, Gielen, and Van Ourti, 2020). Conversely, if exposure to newly eligible siblings raises perceived necessity, vaccination among control children may rise, leading to underestimation (Ma et al., 2006; Yoo et al., 2010).

5. Survey responses indicate that 94.32 percent of influenza vaccinations occurred between September and December overall. When disaggregated by group, 96.51 percent of the treatment group reported vaccination during this period, compared with 86.28 percent of the control group.

Table 2: Robustness of the Estimated Effect on Influenza Vaccine Uptake

	Baseline	Reference window		Contamination			Excluding COVID-19 period
		Excluding 2018.9–12	Alternative post	Excluding age 5,6	Excluding age 13	Within household	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Treat \times Post$	0.117***	0.123***		0.127***	0.127***	0.136***	0.134***
	(0.026)	(0.028)		(0.027)	(0.027)	(0.028)	(0.030)
$Treat \times Post1$			0.085*				
			(0.051)				
$Treat \times Post2$			0.123***				
			(0.028)				
Adjusted R-squared	0.376	0.373	0.376	0.368	0.379	0.380	0.377
Observation	6699	6364	6699	6479	6560	6482	6134

Notes: Table 2 reports robustness checks of the baseline specification in Table 1. The dependent variable, sample, and controls follow Column (4) of Table 1 unless noted otherwise. Column (2) excludes interviews conducted between September and December 2018. Column (3) divides the post-expansion period into two subperiods, $Post1$ (interviews conducted September–December 2018) and $Post2$ (interviews conducted January 2019–August 2020). Column (4) excludes children aged 5 and 6 to address potential contamination from earlier-stage NIP eligibility, and Column (5) excludes age 13 to address contamination from cohort transitions out of the treatment group. Column (6) excludes control-group children who have a sibling in the treatment group, addressing potential within-household spillovers. Column (7) excludes observations surveyed between February and August 2020 to address potential disruption from the early COVID-19 period. All specifications apply KNHANES survey weights, and standard errors account for the survey's complex sampling design. A single asterisk denotes statistical significance at the 10 percent level, double at the 5 percent level, and triple at the 1 percent level.

than higher-income households (Hwang et al., 2017).[6] After the expansion, vaccination became free for eligible children. I therefore estimate the following model to examine whether removing out-of-pocket costs disproportionately increased vaccination uptake in low-income households:

$$Y_i = \beta_0 + \beta_1(Treat_i \times Post_t) + \beta_2(z_i \times Treat_i \times Post_t) + \text{Other interaction terms} + \gamma' X_i + \epsilon_i \quad (3)$$

where all baseline variables are defined as in equation (1), and z_i denotes the individual characteristic used to assess heterogeneity. In addition to the triple interaction term, all corresponding lower-order interaction terms among z_i, $Treat_i$, and $Post_t$ are included.

Table 3 presents heterogeneity in the expansion effect by household income. Column (1) reproduces the baseline estimate. Column (2) interacts $Treat \times Post$ with an above-median household-income indicator. The estimated increase in vaccination probability is 8.5 percentage points larger for children above the median, although the difference between the two income groups is not statistically significant. Column (3) replaces the binary indicator with household-income quartile indicators; the estimated effect rises monotonically with household income, but the across-quartile differences remain statistically insignificant. These results should be interpreted with caution, but they provide no evidence that vaccination uptake increased more among households facing greater financial constraints.

The larger estimated effect among higher-income groups does not imply that financial cost is unimportant. The expansion may have interacted with other constraints that differ systematically by income. Appendix Table A4 shows systematic differences across income levels in characteristics associated with childhood influenza vaccination. Lower-income mothers are less likely to hold a college degree and more likely to work irregular or night shifts. Local healthcare supply, measured by the number of vaccination providers per 1,000 eligible

6. According to KNHANES responses on unmet medical needs, the share citing financial reasons is substantially higher in the below-median income group (12.20 percent) than in the above-median group (2.34 percent). The cell sizes are small, however (128 above-median and 82 below-median observations), so this comparison should be interpreted with caution.

Table 3: Heterogeneous Effects on Influenza Vaccine Uptake by Household Income

	Baseline	By household income	
		Median income	Income quantile
	(1)	(2)	(3)
$Treat \times Post$	0.117***	0.063	0.030
	(0.026)	(0.047)	(0.098)
$Treat \times Post \times Above\ median\ income$		0.085	
		(0.057)	
$Treat \times Post \times Income\ quantile = 2$			0.048
			(0.113)
$Treat \times Post \times Income\ quantile = 3$			0.102
			(0.108)
$Treat \times Post \times Income\ quantile = 4$			0.134
			(0.109)
Adjusted R-squared	0.376	0.377	0.377
Observation	6699	6699	6699

Notes: Table 3 reports heterogeneity in the expansion effect by household income. The dependent variable, sample, and controls follow Column (4) of Table 1. Column (1) reproduces the baseline estimate. Columns (2) and (3) estimate equation (3), allowing the treatment effect to vary by household income. Column (2) splits the sample at the median household income, and Column (3) uses household income quartiles. All specifications apply KNHANES survey weights, and standard errors account for the survey's complex sampling design. A single asterisk denotes statistical significance at the 10 percent level, double at the 5 percent level, and triple at the 1 percent level.

children, also differs.

I next estimate equation (3) using the characteristics identified in Appendix Table A4. Figure 3 stacks the heterogeneity coefficients across four dimensions: maternal college education, employment, work schedule, and local healthcare access. Children whose mothers hold at least a bachelor's degree show a slightly larger increase in vaccination probability, though the difference is not statistically significant. The expansion effect is significantly larger among children with employed mothers than among those whose mothers are not employed. Within the employed subsample, the effect is concentrated among children whose mothers work standard daytime hours rather than non-daytime (night or irregular) shifts. The expansion effect is also larger in regions where the number of vaccination providers per 1,000 eligible children is above the sample mean.

The heterogeneity results indicate that the expansion was more effective among children with working mothers, which helps explain the income gradient documented earlier. During the vaccination campaign period, selected public health centers extended weekend operations in areas lacking referral providers. Administrative statistics, however, show that vaccinations at public health centers accounted for only a small fraction of total vaccinations after the expansion.[7] Given this limited scale, the temporary weekend measure alone is unlikely to have reduced mothers' time costs. Changes in the availability or scheduling of vaccination services at designated medical institutions may have reduced time costs, but the data do not allow a direct test of this channel.

4.4.2 Effect of Text Message Reminders on Influenza Vaccine Uptake

Beyond covering the cost of vaccination, the KDCA ran an information campaign through educational institutions to raise uptake. During the 2019 flu-year, the agency sent reminder messages to all upper-grade primary-school students aged 10–12 to promote vaccination. I

7. According to Korea Centers for Disease Control and Prevention (2019), 68,056 out of 5,744,731 eligible children (approximately 1.5 percent) were vaccinated at public health centers in the 2018–2019 flu-year, while the remaining 98.5 percent received vaccinations at contracted medical institutions.

Figure 3: Heterogeneous Effects on Influenza Vaccine Uptake by Maternal Characteristics and Healthcare Access

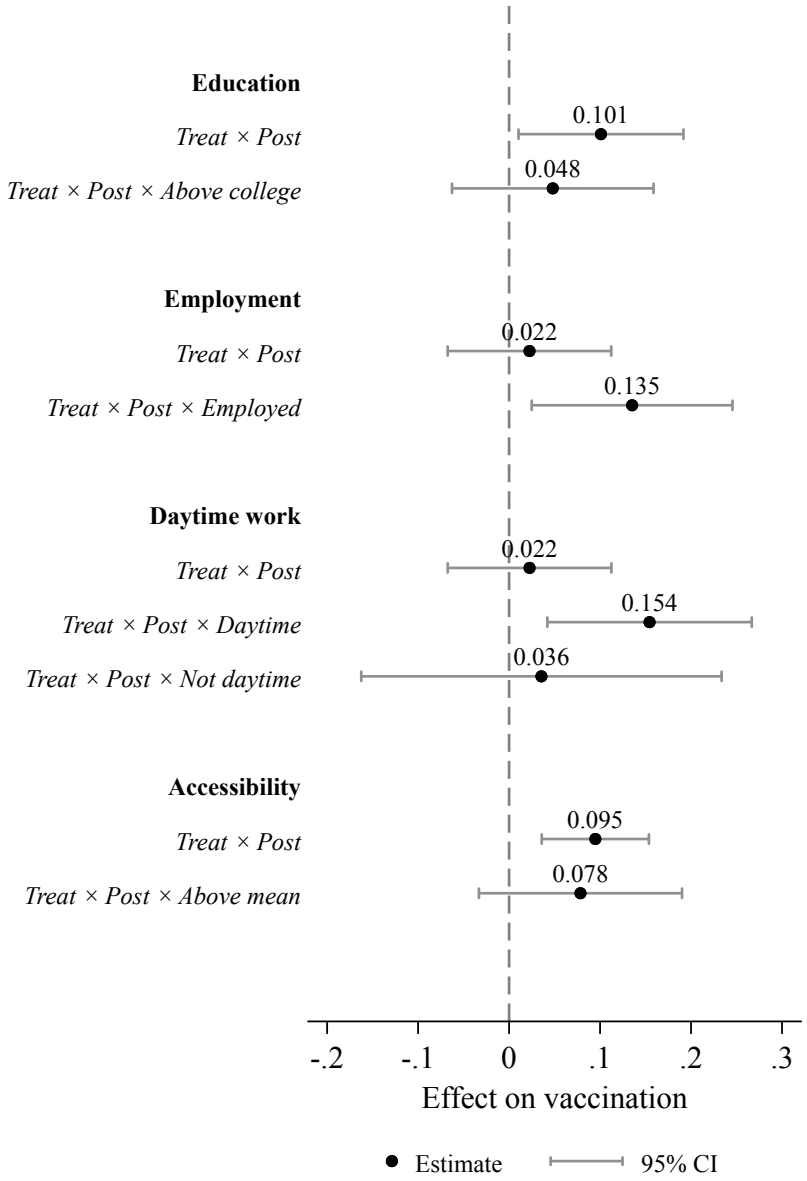

Notes: Figure 3 reports heterogeneity in the expansion effect by maternal characteristics and regional healthcare access, estimated from equation (3). The dependent variable, sample, and controls follow Column (4) of Table 1. The figure stacks four sets of coefficients: (i) maternal college education (bachelor's degree or higher), (ii) maternal employment, (iii) maternal work schedule (daytime, non-daytime, and not working), and (iv) regional healthcare access, defined by whether the number of vaccination providers per 1,000 eligible children in the municipality of residence is at or above the sample mean. Points represent estimated coefficients on the corresponding triple interaction, and horizontal lines denote 95 percent confidence intervals. All specifications apply KNHANES survey weights, and standard errors account for the survey's complex sampling design.

examine the effect of these reminders using the following model:

$$Y_i = \alpha + \sum_{k \in K} \beta_1^k (Treat1_i \times \mathbf{1}[Flu\text{-}year_t = k]) + \sum_{k \in K} \beta_2^k (Treat2_i \times \mathbf{1}[Flu\text{-}year_t = k]) + \text{Other terms} \tag{4}$$

where $Treat1_i$ equals 1 for children aged 5–9, who did not receive a reminder, and 0 otherwise; $Treat2_i$ equals 1 for children aged 10–12, who did receive a reminder, and 0 otherwise. All other control variables match those in equation (1).

Figure 4 presents the estimates of β_1^k and β_2^k from equation (4). If the reminder messages raised vaccination uptake, the coefficient on the ages 10–12 group should rise in the 2019 flu-year. However, no such increase is observed.

Figure 4: Effect of Reminder Messages on Influenza Vaccine Uptake

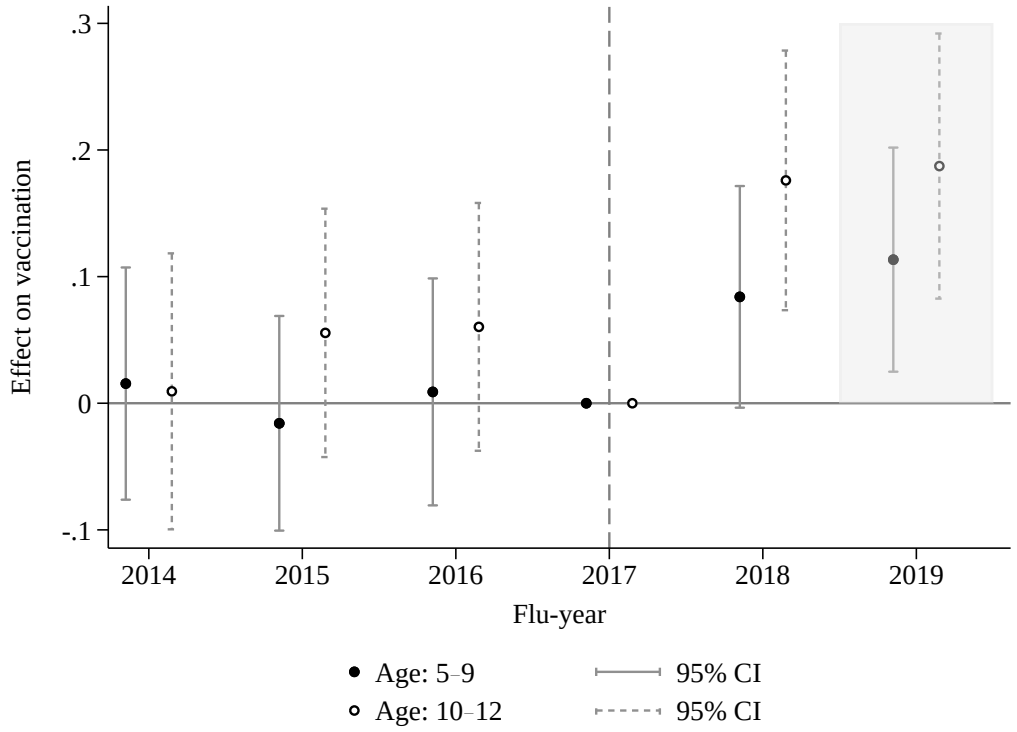

Notes: Figure 4 reports dynamic DiD estimates of β_1^k and β_2^k from equation (4), assessing whether vaccination reminder messages raised uptake. The dependent variable, sample, and controls follow Column (4) of Table 1. $Treat1_i$ equals 1 for children aged 5–9 (who did not receive a reminder), and $Treat2_i$ equals 1 for children aged 10–12 (who did receive a reminder); β_1^k and β_2^k trace each group's vaccination probability relative to the omitted 2017 flu-year. The shaded region marks the 2019 flu-year, when reminders were sent. Vertical lines represent 95 percent confidence intervals. All specifications apply KNHANES survey weights, and standard errors account for the survey's complex sampling design.

5 Effects of the Coverage Expansion on Influenza-Related Healthcare Utilization

Section 4 estimates the effect of the 2018 expansion on vaccination uptake. This section examines whether the resulting rise in vaccination translated into improvements in health outcomes.

5.1 Empirical Approach

5.1.1 Violation of Parallel Trends

Mirroring Section 4, I estimate the following dynamic DiD specification to examine whether a difference-in-differences approach is valid for influenza-related healthcare utilization:

$$Y_{imt} = \beta_0 + \beta_1 Treat_i + \sum_{k \in K} \beta_3^k (Treat_i \times \mathbf{1}[Flu\text{-}year_t = k]) + \gamma' X_{imt} + \delta_t + \epsilon_{imt} \qquad (5)$$

where Y_{imt} is the number of influenza-related claims per 1,000 children in age group i during calendar month m of flu-year t; $Treat_i$ equals 1 for the ages 5–12 group and 0 for the ages 13–18 group. The 2019 flu-year is excluded because of COVID-19, so the post-expansion sample consists only of the 2018 flu-year ($K = \{2014, 2015, 2016, 2018\}$), with the 2017 flu-year as the omitted reference. X_{imt} includes $Treat_i$ interacted with the quadratic of average monthly precipitation, the quadratic of average monthly temperature, and a linear flu-year trend. ϵ_{imt} is the error term, and standard errors are clustered at the age-by-flu-year level.

Figure 5 presents the dynamic DiD estimates based on equation (5). In the 2014 flu-year, the treatment group has significantly fewer influenza-related claims than the reference year (2017), indicating that the two age groups followed different pre-expansion dynamics. The estimated coefficients also vary substantially with the inclusion of control variables. This sensitivity indicates that age-specific responses to flu-year shocks, such as weather conditions

or underlying seasonal trends, play an important role in shaping influenza-related healthcare utilization. If these age-specific time-varying factors correlate with the expansion timing, a standard difference-in-differences approach may confound expansion effects with differential exposure to influenza.

Figure 5: Dynamic DiD Estimates of the 2018 Expansion's Effect on Influenza-Related Healthcare Utilization

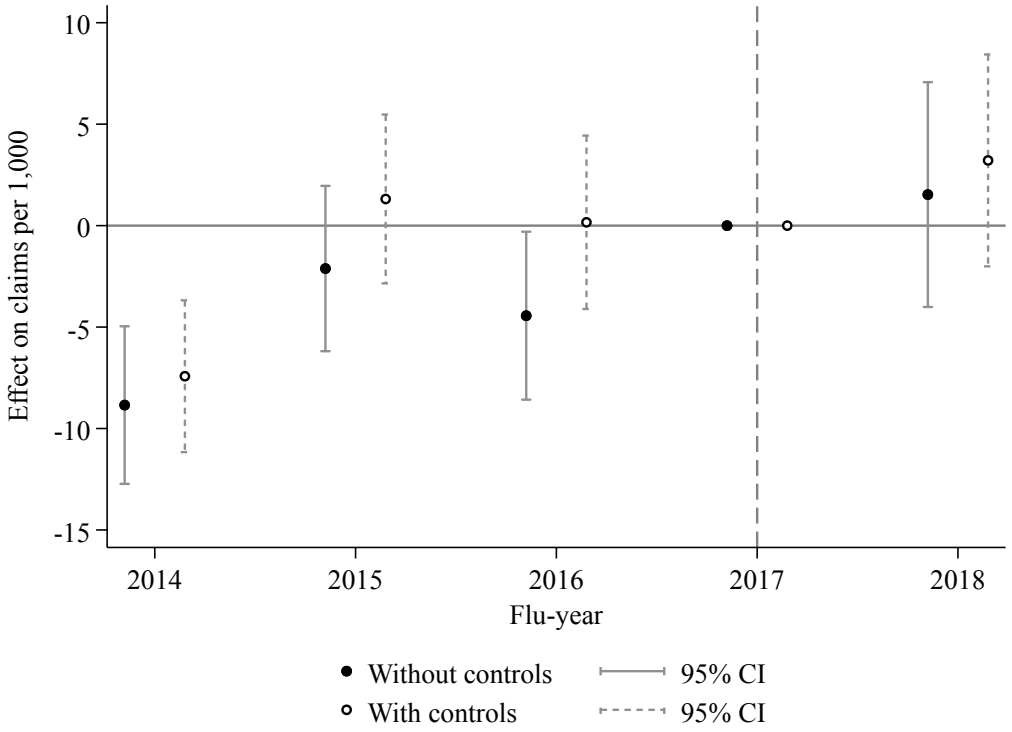

Notes: Figure 5 reports dynamic DiD estimates based on equation (5). The dependent variable is the number of influenza-related claims per 1,000 children by age, flu-year, and calendar month, computed over flu-years 2014–2018. The coefficients shown are the interaction terms $Treat_i \times \mathbf{1}[Flu\text{-}year_t = k]$ between the treatment indicator (age group 5–12) and flu-year dummies, with the 2017 flu-year as the omitted reference period. The 2019 flu-year is excluded due to COVID-19, so the post-expansion period consists only of the 2018 flu-year. Filled circles denote estimates without control variables, and hollow circles denote estimates including the full set of controls described in equation (5). Vertical lines represent 95 percent confidence intervals. Standard errors are clustered at the age-by-flu-year level.

5.1.2 Triple-Difference Approach

Using the vaccine match rate as an additional dimension requires that the effectiveness of influenza vaccination actually vary with match conditions. I verify this relationship in

Appendix Section B, which shows that pediatric influenza-related healthcare utilization falls more sharply with vaccination rates under high-match conditions than under low-match conditions. This supports using the match rate to identify the effects of expansion. To address the age-specific time-varying confounding documented in Section 5.1.1, I estimate the following triple-difference specification:

$$Y_{imt} = \beta_0 + \beta_1(Treat_i \times Post_t \times Good\ match_{mt}) + \text{Other interaction terms} + \gamma' X_{it} + \epsilon_{imt} \quad (6)$$

where $Good\ match_{mt}$ equals one if the vaccine match rate in month m of flu-year t is at or above the average match rate across all periods in the sample, and zero otherwise. All lower-order interactions among $Treat_i$, $Post_t$, and $Good\ match_{mt}$ are included. The control variables match those in equation (5).

The identifying assumption is that, in the absence of the expansion, the difference-in-differences in healthcare utilization between the treatment and control groups would have evolved similarly across high- and low-match environments.

Even when this assumption holds, two features of the DDD coefficient deserve emphasis before turning to the estimates. First, low-match conditions do not imply that the influenza vaccine is ineffective. Even in seasons with relatively poor strain match, the vaccine provides partial protection (Tricco et al., 2013; Belongia et al., 2016). The estimated coefficient in equation (6) should therefore be interpreted as the differential effect in high-match environments relative to low-match environments rather than as an effect relative to a zero-effect benchmark. As a consequence, the DDD coefficient understates the expansion's total health benefit by the amount of protection that accrues during low-match periods. Second, equation (6) identifies an intent-to-treat effect on the eligible age group, combining direct protection received by vaccinated children with indirect protection among unvaccinated peers in the same age band. Any indirect protection that the expansion produces among the ineligible adolescent control group—through within-school or within-household contact with

newly vaccinated younger siblings—biases the DDD coefficient toward zero. The coefficient should therefore be read as a lower bound rather than as a measure of the full community benefit of the expansion.

5.2 Results

5.2.1 Validity of Identifying Assumption

Figure 6 evaluates the identifying assumption of the triple-difference design directly by testing whether the difference between the treatment and control groups was stable across high- and low-match periods before the expansion. Figure 6a plots dynamic DiD coefficients separately for high- and low-match months. In the pre-expansion years, the magnitudes of the estimated coefficients differ somewhat across match levels. In the 2018 flu-year, by contrast, healthcare utilization rises sharply during the low-match period while the corresponding estimate remains close to zero during the high-match period. What matters for the triple-difference design is the pre-expansion stability of the difference-in-differences between high- and low-match periods, which Figure 6b examines: as the figure shows, the pre-expansion DDD coefficients are all small and statistically insignificant (Olden and Møen, 2022). These findings suggest that low-match periods provide a plausible counterfactual for high-match periods.

5.2.2 Main Results

Table 4 presents the estimated effects of the 2018 expansion on influenza-related healthcare utilization. Columns (1)–(4) report standard difference-in-differences specifications. Column (1), which excludes additional controls, suggests an increase in influenza-related healthcare utilization in the treatment group after the policy. Motivated by the pattern in Figure 5, Column (2) allows for group-specific linear flu-year trends; the coefficient changes sign and becomes negative, indicating that failing to account for differential pre-existing trends may

Figure 6: Dynamic DiD Estimates by Vaccine Match Quality: DiD and DDD

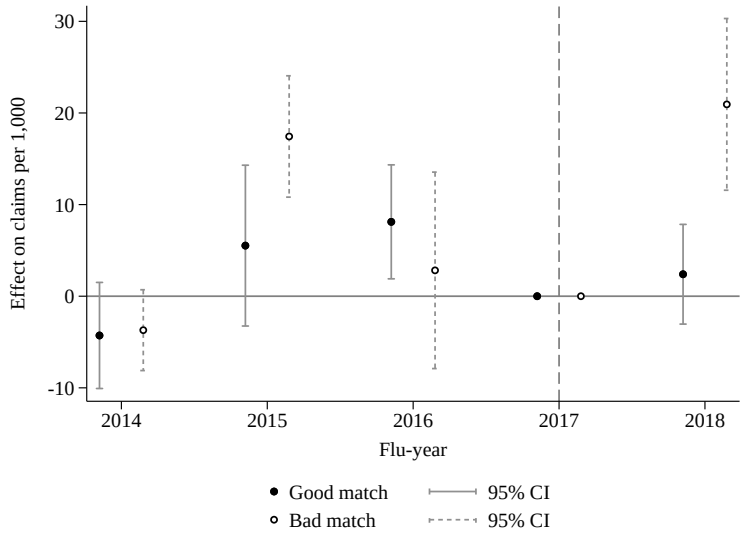

(a) Dynamic DiD by match level

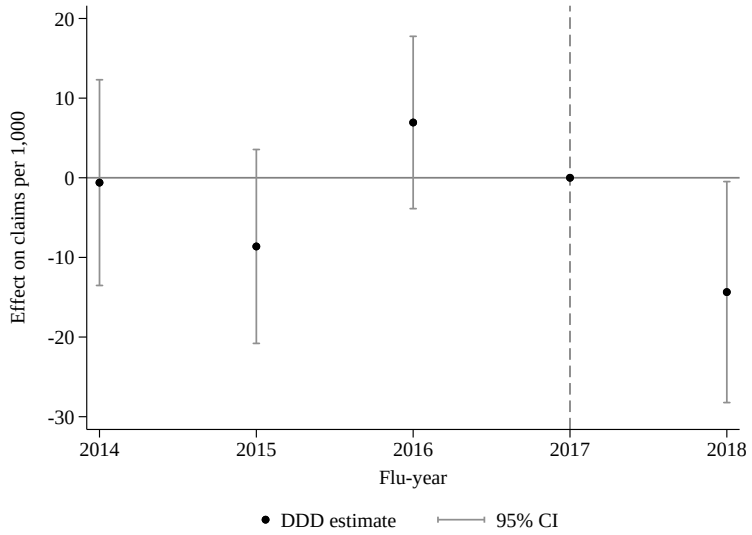

(b) Triple-difference dynamic DiD

Notes: Figure 6 reports dynamic DiD estimates assessing the validity of the identifying assumption underlying the triple-difference specification. The dependent variable is the number of influenza-related claims per 1,000 children at the age-by-year-month level, computed over flu-years 2014–2018. The treatment group is age group 5–12 and the control group is age group 13–18. A high-match month is defined as a month in which the vaccine match rate (constructed as described in Appendix Section A) is at or above the sample-period average; otherwise it is a low-match month. Figure 6a plots dynamic DiD coefficients on $Treat_i \times \mathbf{1}[Flu\text{-}year_t = k]$ separately for high- and low-match months, based on equation (5). Figure 6b plots triple-difference dynamic DiD coefficients on $Treat_i \times \mathbf{1}[Flu\text{-}year_t = k] \times Good\ match_{mt}$, the dynamic DiD analogue of equation (6). In both figures, the 2017 flu-year is the omitted reference period; the 2019 flu-year is excluded due to COVID-19. Vertical lines represent 95 percent confidence intervals. Standard errors are clustered at the age-by-flu-year level.

bias the baseline estimate. Columns (3) and (4) add climate controls, and the estimated effect remains close to zero at −0.475 in Column (4).

Columns (5) and (6) split the DiD estimates by high- and low-match periods. The results indicate that the increase in healthcare utilization in the treatment group is concentrated in low-match periods, whereas utilization declines during high-match periods, although the estimates are not statistically significant. This heterogeneity supports the validity of the match-rate-based DDD design: standard DiD estimates aggregate two opposing effects across high- and low-match months, while the DDD specification isolates the high-match component.

Column (7) presents the triple-difference estimates from equation (6), using an indicator for high-match periods. The estimated triple-interaction coefficient is statistically significant and implies a reduction of 13.3 cases per 1,000 in influenza-related healthcare utilization in the treatment group during high-match periods after the policy. Column (8) replaces the binary match indicator with the continuous vaccine match rate and shows a significant decline in healthcare utilization as the match rate rises, corresponding to a reduction of 10.56 cases per 1,000 for a one-standard-deviation increase in the match rate.[8] Column (9) reports a difference-in-differences estimate on the same restricted sample as Columns (7) and (8), addressing the concern that missing match-rate observations might alter sample composition; the results remain consistent, indicating that sample selection is unlikely to drive the findings in the triple-difference specifications.

6 Conclusion

This study examines the effects of Korea's 2018 expansion of the National Immunization Program, which extended free influenza-vaccine eligibility to all primary-school children aged 5 to 12. For vaccination uptake, the difference-in-differences estimates indicate that the expansion raised influenza vaccination rates among newly eligible children by approximately

8. The vaccine match rate has a mean of 0.6325 and a standard deviation of 0.3342.

Table 4: Effect of the 2018 Expansion on Influenza-Related Healthcare Utilization

	DiD				DiD – subsample		DDD		DiD with the sample used in (7), (8)
					Good match	Bad match	Binary	Continuous	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treat × *Post*	5.380*	-0.673	4.802*	-0.475	-4.154	14.022***	7.738	21.383***	-0.024
	(2.703)	(3.071)	(2.739)	(3.120)	(3.407)	(4.847)	(4.672)	(6.166)	(2.951)
Treat × *Post* × *Good match*							-13.300***		
							(3.804)		
Treat × *Post* × *Match rate*								-31.600***	
								(6.326)	
Controls:									
Group-specific flu-year trend		Y		Y	Y	Y	Y	Y	Y
Weather conditions			Y	Y	Y	Y	Y	Y	Y
Adjusted R-squared	0.806	0.807	0.843	0.844	0.870	0.809	0.848	0.846	0.844
Observation	840	840	840	840	420	364	784	784	784

Notes: Table 4 reports the estimated effects of the 2018 expansion on influenza-related healthcare utilization. The dependent variable is the number of influenza-related claims per 1,000 children at the age-by-year-month level. The treatment group is age group 5–12 and the control group is age group 13–18; the post-expansion period covers the 2018/19 flu-year, and the 2019/20 flu-year is excluded due to COVID-19. A high-match month is defined as in Figure 6 (constructed as described in Appendix Section A). Columns (1)–(4) report difference-in-differences (DiD) estimates under alternative specifications. Columns (5) and (6) report DiD estimates separately for high- and low-match months. Columns (7) and (8) report triple-difference (DDD) estimates based on equation (6), using a dummy variable for high match in Column (7) and the continuous vaccine match rate in Column (8). Column (9) reports a DiD estimate on the same restricted sample as Columns (7) and (8). Group-specific linear flu-year trends and weather controls are included where indicated. Standard errors are clustered at the age-by-flu-year level. A single asterisk denotes statistical significance at the 10 percent level, double at the 5 percent level, and triple at the 1 percent level.

11.7 percentage points—a roughly 20 percent increase over the pre-expansion mean. The heterogeneity analyses provide no evidence that the expansion disproportionately benefited lower-income households; the estimated effects are instead larger among children from higher-income families.

To evaluate the health effects of the expansion, I use a triple-difference approach that exploits variation in influenza vaccine match rates over time. A standard difference-in-differences approach is not appropriate here because the two age groups show different trends in influenza-related healthcare utilization. By comparing high- and low-match periods, the triple-difference approach isolates the policy-induced health benefit specific to high-match environments, where vaccine effectiveness is greater. The estimates indicate that the expansion reduced influenza-related healthcare utilization among children aged 5 to 12 by approximately 13.3 cases per 1,000 during high-match periods after the policy.

To translate these utilization gains into monetary terms, I conduct a simple cost-benefit comparison of the 2018 expansion. On the cost side, the government supported 2,114,414 vaccinations among children aged 5 to 12 in the 2018–2019 flu season at a unit cost of 27,650 won per vaccination, for a total program cost of approximately 58.5 billion won.[9] On the benefit side, I estimate the reduction in influenza-related healthcare expenditures using a continuous match-rate specification parallel to Column (8) of Table 4, with healthcare expenditure per 1,000 children as the outcome. Applying the estimated coefficient to the total population of children aged 5 to 12 as of December 2018 (3,743,725 children) and to the realized match rates over the post-expansion period yields an estimated reduction in healthcare expenditure of approximately 44.9 billion won.[10] The resulting benefit-to-cost ratio is approximately 0.77, indicating that direct healthcare-cost savings alone do not

9. Won amounts can be converted to USD using the average won–dollar exchange rate over the 2018–2019 flu season (approximately 1,124 won per dollar). Under this rate, the unit cost is approximately USD 25 and the total program cost is approximately USD 52 million. This figure captures only the direct vaccination costs and excludes administrative expenses associated with program implementation, for which data are unavailable; the true program cost is therefore higher.

10. The DDD coefficient on healthcare expenditure is −1,446,612 (standard error 340,413; significant at the 1 percent level).

fully offset the program cost. The indirect benefits of influenza vaccination are substantial, however.[11] I do not estimate these indirect components in this paper. The literature suggests, however, that indirect benefits often exceed direct healthcare-cost savings, so once these benefits are accounted for, the total benefits of the program likely exceed its cost.

Several limitations warrant mention. First, the analysis in Section 5 relies on a single post-expansion flu-year (2018–2019), which raises concerns about the generalizability of the estimates to seasons with different epidemic conditions. Second, the triple-difference coefficient identifies an ITT effect on the eligible age group, combining direct protection with within-group indirect protection among children ages 5–12. Because influenza vaccines retain partial effectiveness even in low-match environments, and because any spillover to the ineligible adolescent control group attenuates the differential, the estimated reduction represents a lower bound on the policy's ITT effect on newly eligible children. The design does not identify the full community benefit of the program.

11. For school-aged children, these include reductions in school absenteeism and, importantly, productivity losses incurred by caregivers who miss work to attend to sick children. The existing literature consistently finds that indirect costs constitute a large share of the total economic burden of influenza. Molinari et al. (2007) estimate that indirect costs—primarily lost earnings due to illness and loss of life—amount to $16.3 billion annually in the United States, compared with $10.4 billion in direct healthcare costs. Putri et al. (2018), updating these estimates with 2015 data, find that indirect costs ($8.0 billion) are approximately 2.5 times the direct healthcare costs ($3.2 billion). Studies focusing specifically on children report similar patterns: Ortega-Sanchez et al. (2012) find that caregivers of hospitalized children with influenza miss an average of 73 work hours per episode, and Li and Leader (2007) document that households with school-aged children experiencing influenza-like illness incur significantly higher medical expenses and work-loss days.

References

Abrevaya, Jason, and Karen Mulligan. 2011. "Effectiveness of State-Level Vaccination Mandates: Evidence from the Varicella Vaccine." *J. Health Econ.* 30 (5): 966–976. https://doi.org/10.1016/j.jhealeco.2011.06.003.

Almond, Douglas. 2006. "Is the 1918 Influenza Pandemic Over? Long-Term Effects of In Utero Influenza Exposure in the Post-1940 U.S. Population." *J. Polit. Econ.* 114.

Belongia, Edward A, Melissa D Simpson, Jennifer P King, Maria E Sundaram, Nicholas S Kelley, Michael T Osterholm, and Huong Q McLean. 2016. "Variable Influenza Vaccine Effectiveness by Subtype: A Systematic Review and Meta-Analysis of Test-Negative Design Studies." *The Lancet Infectious Diseases* 16 (8): 942–951.

Bitler, Marianne P, and Christopher S Carpenter. 2016. "Health Insurance Mandates, Mammography, and Breast Cancer Diagnoses." *American Economic Journal: Economic Policy* 8 (3): 39–68. https://doi.org/10.1257/pol.20120298.

Bouckaert, Nicolas, Anne C Gielen, and Tom Van Ourti. 2020. "It Runs in the Family – Influenza Vaccination and Spillover Effects." *J. Health Econ.* 74. https://doi.org/10.1016/j.jhealeco.2020.102386.

Campos-Mercade, Pol, Armando N Meier, Florian H Schneider, and Erik Wengström. 2021. "Prosociality Predicts Health Behaviors during the COVID-19 Pandemic." *J. Public Econ.* 195 (104367): 104367. https://doi.org/10.1016/j.jpubeco.2021.104367.

Carpenter, Christopher S, and Emily C Lawler. 2019. "Direct and Spillover Effects of Middle School Vaccination Requirements." *American Economic Journal: Economic Policy* 11 (1): 95–125. https://doi.org/10.1257/pol.20170067.

Choi, Yu Jung, Joon Young Song, Seong-Heon Wie, Won Suk Choi, Jacob Lee, Jin-Soo Lee, Young Keun Kim, Shin Woo Kim, Sun Hee Lee, Kyung-Hwa Park, Hye Won Jeong, Jin Gu Yoon, Hye Seong, Eliel Nham, Ji Yun Noh, Hee Jin Cheong, and Woo Joo Kim. 2024. "Real-World Effectiveness of Influenza Vaccine over a Decade during the 2011–2021 Seasons—Implications of Vaccine Mismatch." *Vaccine* 42 (26): 126381. https://doi.org/10.1016/j.vaccine.2024.126381.

Goldin, Shoshanna, Donald Brooks, Pernille Jorgensen, Pushpa Wijesinghe, Heeyoun Cho, Rania Attia, Reena Doshi, Francisco Nogareda, Belinda Herring, Laure Dumolard, et al. 2024. "Seasonal Influenza Vaccination: A Global Review of National Policies in 194 WHO Member States in 2022." *Vaccine* 42 (26): 126274.

Hong, Seung Hyun. 2013. "MEASURING THE EFFECT OF NAPSTER ON RECORDED MUSIC SALES: DIFFERENCE-IN-DIFFERENCES ESTIMATES UNDER COMPOSITIONAL CHANGES." *J. Appl. Econometrics* 28 (2): 297–324. https://doi.org/10.1002/jae.1269.

Humlum, Maria Knoth, Marius Opstrup Morthorst, and Peter Rønø Thingholm. 2024. "Sibling Spillovers and the Choice to Get Vaccinated: Evidence from a Regression Discontinuity Design." *J. Health Econ.* 94 (102843): 102843. https://doi.org/10.1016/j.jhealeco.2023.102843.

Hwang, Ji Hyen, Chang Hoon Lim, Dong Ho Kim, Byung Wook Eun, Dae Sun Jo, Young Hwan Song, and Yun Kyung Kim. 2017. "A Survey of Parental Perception and Pattern of Action in Response to Influenza-like Illness in Their Children: Including Healthcare Use and Vaccination in Korea." *J. Korean Med. Sci.* 32 (2): 204–211. https://doi.org/10.3346/jkms.2017.32.2.204.

Kelly, Elaine. 2011. "The Scourge of Asian Flu: In Utero Exposure to Pandemic Influenza and the Development of a Cohort of British Children." *J. Hum. Resour.* 46 (4): 669–694. https://doi.org/10.1353/jhr.2011.0004. JSTOR: 41304838.

Korea Centers for Disease Control and Prevention. 2017. *Guidelines for the 2017–2018 season influenza national immunization program (for public health centers).* Manual. Cheongju, Republic of Korea: Korea Centers for Disease Control and Prevention.

Korea Centers for Disease Control and Prevention. 2019. *Final Report on the 2018–2019 Season Influenza National Immunization Program.* Government Report. Cheongju, Republic of Korea: Korea Centers for Disease Control and Prevention.

Korea Centers for Disease Control and Prevention. 2020. *Guidelines for the 2020–2021 Season Influenza National Immunization Program (for Public Health Centers).* Program Manual. Cheongju, Republic of Korea: Korea Centers for Disease Control and Prevention.

Lawler, Emily C. 2017. "Effectiveness of Vaccination Recommendations versus Mandates: Evidence from the Hepatitis A Vaccine." *J. Health Econ.* 52:45–62. https://doi.org/10.1016/j.jhealeco.2017.01.002.

Lee, Jae-Tae. 2023. *Rapid Test for Influenza A and B Virus Antigen Detection [General Immunoassay].* Technical Report. Seoul, Republic of Korea: National Evidence-based Healthcare Collaborating Agency (NECA).

Li, Su, and Shelah Leader. 2007. "Economic Burden and Absenteeism from Influenza-like Illness in Healthy Households with Children (5–17 Years) in the US." *Respiratory Medicine* 101 (6): 1244–1250. https://doi.org/10.1016/j.rmed.2006.10.022.

Lin, Ming Jen, and Elaine M Liu. 2014. "Does in Utero Exposure to Illness Matter? The 1918 Influenza Epidemic in Taiwan as a Natural Experiment." *J. Health Econ.* 37 (1): 152–163. https://doi.org/10.1016/j.jhealeco.2014.05.004.

Ma, K K, W Schaffner, C Colmenares, J Howser, J Jones, and K A Poehling. 2006. "Influenza Vaccinations of Young Children Increased with Media Coverage in 2003." *Pediatrics* 117 (2): e157–63. https://doi.org/10.1542/peds.2005-1079.

Molinari, Noelle-Angelique M., Ismael R. Ortega-Sanchez, Mark L. Messonnier, William W. Thompson, Pascale M. Wortley, Eric Weintraub, and Carolyn B. Bridges. 2007. "The Annual Impact of Seasonal Influenza in the US: Measuring Disease Burden and Costs." *Vaccine* 25 (27): 5086–5096. https://doi.org/10.1016/j.vaccine.2007.03.046.

Olden, Andreas, and Jarle Møen. 2022. "The Triple Difference Estimator." *Econom. J.* 25 (3): 531–553. https://doi.org/10.1093/ectj/utac010.

Ortega-Sanchez, Ismael R., Noelle-Angelique M. Molinari, Gerry Fairbrother, Peter G. Szilagyi, Kathryn M. Edwards, Marie R. Griffin, Amy Cassedy, Katherine A. Poehling, Carolyn Bridges, and Mary Allen Staat. 2012. "Indirect, out-of-Pocket and Medical Costs from Influenza-Related Illness in Young Children." *Vaccine* 30 (28): 4175–4181. https://doi.org/10.1016/j.vaccine.2012.04.057.

Ortiz, Justin R, Marc Perut, Laure Dumolard, Pushpa Ranjan Wijesinghe, Pernille Jorgensen, Alba Maria Ropero, M Carolina Danovaro-Holliday, James D Heffelfinger, Carol Tevi-Benissan, Nadia A Teleb, et al. 2016. "A Global Review of National Influenza Immunization Policies: Analysis of the 2014 WHO/UNICEF Joint Reporting Form on Immunization." *Vaccine* 34 (45): 5400–5405.

Osterholm, Michael T., Nicholas S. Kelley, Alfred Sommer, and Edward A. Belongia. 2012. "Efficacy and Effectiveness of Influenza Vaccines: A Systematic Review and Meta-Analysis." *The Lancet Infectious Diseases* 12 (1): 36–44. https://doi.org/10.1016/S1473-3099(11)70295-X. PMID: 22032844.

Putri, Wayan C. W. S., David J. Muscatello, Melissa S. Stockwell, and Anthony T. Newall. 2018. "Economic Burden of Seasonal Influenza in the United States." *Vaccine* 36 (27): 3960–3966. https://doi.org/10.1016/j.vaccine.2018.05.057.

Schmid, Philipp, Dorothee Rauber, Cornelia Betsch, Gianni Lidolt, and Marie-Luisa Denker. 2017. "Barriers of Influenza Vaccination Intention and Behavior - A Systematic Review of Influenza Vaccine Hesitancy, 2005 - 2016." *PLoS One* 12 (1): e0170550. https://doi.org/10.1371/journal.pone.0170550.

Schneider, Eric B. 2023. "The Determinants of Child Stunting and Shifts in the Growth Pattern of Children: A Long-run, Global Review." *J. Econ. Surv.,* https://doi.org/10.1111/joes.12591.

Tricco, Andrea C, Ayman Chit, Charlene Soobiah, David Hallett, Genevieve Meier, Maggie H Chen, Mariam Tashkandi, Chris T Bauch, and Mark Loeb. 2013. "Comparing Influenza Vaccine Efficacy against Mismatched and Matched Strains: A Systematic Review and Meta-Analysis." *BMC medicine* 11 (1): 153.

Ward, Courtney J. 2014. "Influenza Vaccination Campaigns: Is an Ounce of Prevention Worth a Pound of Cure?" *Am. Econ. J. Appl. Econ.* 6 (1 A): 38–72. https://doi.org/10.1257/app.6.1.38.

White, Corey. 2021. "Measuring Social and Externality Benefits of Influenza Vaccination." *J. Hum. Resour.* 56 (3): 749–785. https://doi.org/10.3368/jhr.56.3.1118-9893R2.

World Health Organization. 2025. "Influenza (Seasonal)." Accessed April 29, 2026. https://www.who.int/news-room/fact-sheets/detail/influenza-(seasonal).

Yoo, Byung-Kwang, Margaret L Holland, Jay Bhattacharya, Charles E Phelps, and Peter G Szilagyi. 2010. "Effects of Mass Media Coverage on Timing and Annual Receipt of Influenza Vaccination among Medicare Elderly." *Health Serv. Res.* 45 (5 Pt 1): 1287–1309. https://doi.org/10.1111/j.1475-6773.2010.01127.x.

Appendix

A Vaccine Effectiveness and Match Rate

Following Weinberg and Szilagyi (2010), vaccine effectiveness refers to the reduction in disease observed under real-world conditions. Unlike vaccine efficacy measured in randomized controlled trials, vaccine effectiveness reflects how vaccination translates into population-level disease outcomes in routine practice. It reflects the biological properties of the vaccine together with behavioral responses, coverage rates, and contextual factors that shape vaccine uptake and disease transmission.

Influenza vaccine effectiveness varies substantially across seasons, depending largely on the degree of match between the vaccine strains and circulating influenza viruses. Under well-matched conditions, seasonal influenza vaccines achieve 70–90 percent effectiveness in preventing laboratory-confirmed influenza among healthy adults (Osterholm et al., 2012). When vaccine strains are poorly matched to circulating viruses, effectiveness declines considerably. Evidence from Korea shows that vaccine mismatch can substantially reduce real-world effectiveness across seasons (Choi et al., 2024). These findings motivate accounting for match conditions when evaluating policy impacts on influenza-related health outcomes.

Seasonal influenza vaccines protect against multiple circulating strains. The traditional trivalent vaccine includes three strains: two influenza A subtypes (H1N1 and H3N2) and one influenza B lineage. Influenza A viruses currently circulating in humans are primarily H1N1 and H3N2, although other subtypes have circulated historically; for example, H2N2 was the dominant human influenza A virus between 1957 and 1968 (Krammer et al., 2018). Influenza B viruses circulate as two distinct lineages, Victoria and Yamagata, which may co-circulate within a season. Because multiple strains and lineages co-circulate, vaccine effectiveness depends on how well the selected vaccine strains match the strains that predominate during the season. Mismatch may arise when the dominant circulating strain differs from the strain included in the vaccine, particularly for influenza B, where the choice of lineage can vary

across seasons.

Influenza vaccine match for month t is calculated as:

$$E_{mt} = M_t^{H1} \times I_{mt}^{H1} + M_t^{H3} \times I_{mt}^{H3} + M_t^{B} \times I_{mt}^{B} \tag{A1}$$

where E_{mt} is the overall vaccine match rate in month m of flu-year t. M_t^j is the seasonal match between the vaccine strain and the circulating strain for virus type $j \in \{H1N1, H3N2, B\}$, defined at the flu-year level. I_{mt}^j is the share of virus type j among all laboratory-confirmed influenza cases detected in month m of flu-year t.[12]

This formulation weights strain-specific match rates by the contemporaneous distribution of circulating strains and captures the effective match between the vaccine and the viral environment in a given month. Figure A1 plots the resulting seasonal match rate by flu-year for 2014–2018, showing substantial across-year variation that the triple-difference approach exploits.

Temporal variation in vaccine match arises from two main mechanisms: changes in viral composition and changes in the relative incidence of circulating strains. First, influenza viruses evolve continuously. Antigenic drift refers to small, gradual mutations that accumulate as the virus replicates. Vaccine composition is determined several months before the influenza season begins, typically in February or March; antigenic drift between strain selection and the subsequent epidemic season may reduce the match between the vaccine and circulating strains (Smith et al., 2004; Petrova and Russell, 2018). This mechanism differs from antigenic shift, a rare and abrupt major change in viral subtype such as the 2009 H1N1 pandemic. The mismatch considered in this study reflects antigenic drift rather than shift. Multiple influenza strains also circulate simultaneously, while seasonal vaccines include only three or four strains. Even without substantial mutation, mismatch may occur if a strain not included in the vaccine becomes dominant during the epidemic. Because pre-

12. I thank Corey White for sharing materials related to the construction of the match rate used in White (2021).

dicting dominant strains months in advance is inherently difficult, vaccine mismatch is often unpredictable.

Second, influenza A viruses typically peak in winter while influenza B viruses are often observed later, such as in spring. The relative timing and intensity of these peaks vary across seasons, however (Caini et al., 2015). Even after the overall seasonal match rate (M) is determined, within-season variation in the distribution of circulating strains may alter the effective match at a given point in time. The realized effectiveness of the vaccine in a particular month therefore depends on both the ex ante match rate and which strains are actively circulating during that period.

B Vaccine Match Rate as Moderating Channel

Using the vaccine match rate as an additional difference variable requires that vaccine effectiveness vary with match conditions. To explore this, I collected influenza vaccination coverage data by district (sigungu in Korean) and age group for the 2018 flu-year.[13] Using the vaccination rates and match rates, I estimate the following difference-in-differences model:

$$Y_{icm} = \beta_1(V_{ic} \times E_m) + \delta_{ic} + \delta_{sm} + \epsilon_{icm} \tag{A2}$$

where Y_{icm} is the number of influenza-related claims per 1,000 children in district c, age group i, and month m. The age groups are 5–6, 7–9, and 10–12 years. V_{ic} is the age-group–district influenza vaccination rate, and E_m is the monthly influenza vaccine match rate. The fixed effect δ_{ic} absorbs time-invariant age-group–district characteristics related to influenza vaccination, and δ_{sm} absorbs province–month time trends. Standard errors are clustered at the age-group–district level.

Table A1 presents the results of estimating equation (A2). My preferred specification, Column (4), controls for both age-group–district and province–month fixed effects and yields a negative and statistically significant interaction between vaccination rates and match rates. This coefficient provides direct evidence that vaccine match moderates the effect of vaccination on influenza-related healthcare utilization: higher vaccination rates reduce utilization more sharply when the strain match is better.

Columns (1)–(3) motivate why these fixed effects are essential. Column (1) regresses utilization on the vaccination rate alone and finds a positive association, reflecting confounding by local healthcare infrastructure, climate, and other time-invariant factors that drive both vaccination and care-seeking behavior. Column (2) adds the match rate and yields a positive association between match and utilization, because the high-match period in the 2018

13. KDCA only provides information on influenza vaccination coverage for those eligible for the National Vaccination Program.

flu-year coincided with the peak of a relatively virulent influenza A season. Column (3) adds the interaction between vaccination and match rates, with a coefficient that points in the expected negative direction but is not statistically significant; the estimate stabilizes only once age-group–district and province–month fixed effects are added in Column (4). Columns (5) and (6) re-estimate the specification of Column (4) separately for influenza-active and non-active months and confirm that the moderating effect is concentrated in active epidemic periods, consistent with a vaccine-effectiveness interpretation.

C Figures

Figure A1: Seasonal Vaccine Match Rate by Flu-Year

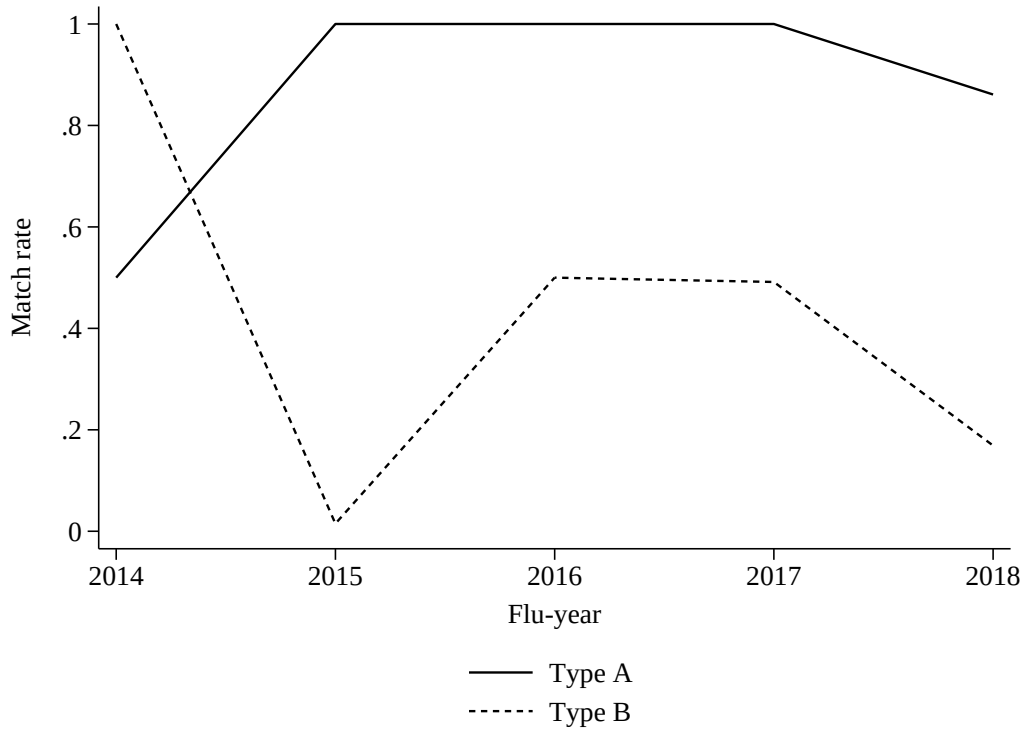

Notes: Figure A1 plots the seasonal influenza vaccine match rate by flu-year for 2014–2018. The match rate is computed as the strain-specific match between vaccine and circulating viruses, aggregated over the flu-year using the within-season distribution of laboratory-confirmed influenza cases as weights (see Appendix Section A for details). Higher values indicate closer alignment between vaccine composition and circulating strains.

D Tables

Table A1: Moderating Role of Vaccine Match Rate in the Vaccination–Utilization Relationship

	All months				Flu season	Non-flu season
	(1)	(2)	(3)	(4)	(5)	(6)
Vaccination rate	66.601***	66.601***	80.067***			
	(4.780)	(4.780)	(13.977)			
Match rate		13.850***	27.232**			
		(1.621)	(11.741)			
Vaccination rate × *Match rate*			-19.471	-54.315***	-43.841	-7.260**
			(17.430)	(15.206)	(26.634)	(3.103)
Controls:						
Age × District FE				Y	Y	Y
Province × Month FE				Y	Y	Y
Adjusted R-squared	0.008	0.010	0.010	0.801	0.736	0.389
Observation	9,000	9,000	9,000	9,000	4,500	4,500

Notes: Table A1 reports estimates of equation (A2), which examines how the influenza vaccine match rate moderates the association between vaccination rates and influenza-related healthcare utilization. The dependent variable is the number of influenza-related claims per 1,000 children at the age-group–district–month level. The sample covers age groups 5–6, 7–9, and 10–12 in the 2018 flu-year. V_{ic} is the district-level influenza vaccination rate by age group, and E_m is the monthly influenza vaccine match rate (constructed as described in Appendix Section A). Column (1) regresses utilization on V_{ic} alone; Column (2) adds E_m; Column (3) adds the interaction $V_{ic} \times E_m$; Column (4) further adds age-group–district and province–month fixed effects. Columns (5) and (6) re-estimate the Column (4) specification separately for influenza-active and non-active months. Standard errors are clustered at the age-group–district level. A single asterisk denotes statistical significance at the 10 percent level, double at the 5 percent level, and triple at the 1 percent level.

Table A2: Variable Definitions

Variable	Definition
Section 4: Effects of the Coverage Expansion on Influenza Vaccination	
Child characteristics	
Age	Child's age in years at survey
Female	=1 if child is female
Subjective health status	Categorical: Poor or very poor / Fair / Good or very good
Household characteristics	
Household size	Number of household members
Above median income	=1 if household income is above the sample median
Private health insurance	Categorical: Yes / No / No response
Health insurance type	Categorical: Self-employed insured / Employee insured / Medical Aid beneficiary / No response
Maternal characteristics	
Age	Mother's age in years (missing indicator included)
College education	Categorical: Less than college / College or above / Missing
Influenza vaccination status	Categorical: Unvaccinated / Vaccinated / Missing
Employment status	Categorical: Yes / No / Missing
Section 5: Effects of the Coverage Expansion on Influenza-Related Healthcare Utilization	
Weather conditions	
Mean temperature	Monthly average temperature (°C, quadratic term)
Precipitation	Monthly average precipitation (mm, quadratic term)

Table A3: Descriptive Statistics: Vaccination Sample

	Control				Treated			
	Pre		Post		Pre		Post	
	N	Mean	N	Mean	N	Mean	N	Mean
Vaccinated	1776	0.245	776	0.344	2805	0.574	1342	0.799
Child characteristics:								
Age	1776	15.652	776	15.687	2805	8.547	1342	8.587
Female	1776	0.471	776	0.459	2805	0.488	1342	0.490
Subjective health status								
Poor or very poor	1776	0.053	776	0.057	2805	0.022	1342	0.027
Fair	1776	0.366	776	0.357	2805	0.230	1342	0.199
Good or very good	1776	0.581	776	0.585	2805	0.747	1342	0.774
Household characteristics:								
Household size	1776	3.975	776	3.885	2805	4.172	1342	4.095
Above median income	1776	0.638	776	0.638	2805	0.653	1342	0.629
Private health insurance								
Yes	1776	0.886	776	0.930	2805	0.918	1342	0.932
No	1776	0.103	776	0.063	2805	0.073	1342	0.059
No response	1776	0.011	776	0.007	2805	0.009	1342	0.009
Health insurance type								
Self-employed insured	1776	0.253	776	0.228	2805	0.220	1342	0.186
Employee insured	1776	0.683	776	0.744	2805	0.749	1342	0.783
Medical Aid beneficiary	1776	0.052	776	0.028	2805	0.026	1342	0.030
No response	1776	0.012	776	0.000	2805	0.005	1342	0.001
Maternal characteristics:								
Age	1776	40.718	776	42.530	2805	37.218	1342	37.316
Age: missing	1776	0.092	776	0.070	2805	0.054	1342	0.064
College education								
Less than college	1776	0.455	776	0.411	2805	0.320	1342	0.276
College or above	1776	0.385	776	0.469	2805	0.557	1342	0.605
Missing	1776	0.160	776	0.120	2805	0.123	1342	0.119
Influenza vaccination status								
Unvaccinated	1776	0.660	776	0.616	2805	0.566	1342	0.479
Vaccinated	1776	0.180	776	0.264	2805	0.312	1342	0.402
Missing	1776	0.160	776	0.120	2805	0.122	1342	0.119
Employment status								
Yes	1776	0.645	776	0.684	2805	0.553	1342	0.579
No	1776	0.195	776	0.197	2805	0.324	1342	0.300
Missing	1776	0.160	776	0.120	2805	0.123	1342	0.121

Notes: Table A3 reports sample sizes and means of vaccination rates and key covariates for the analysis sample used in the baseline specification (Table 1, Column 4). Columns are organized by group (control, ages 13–18; treatment, ages 5–12) and period (pre- and post-expansion). Missing values for maternal characteristics are coded with a missing-category indicator and retained in the sample. Variable definitions are provided in Table A2. All statistics are calculated using survey weights.

Table A4: Maternal Characteristics and Healthcare Supply by Household Income

	Household income		
	Below median	Above median	Difference
Age	41.808	42.188	-0.381
College education:			
Less than college	0.485	0.304	0.180
College or above	0.318	0.597	-0.279
Missing	0.197	0.098	0.099
Employment status:			
Yes	0.532	0.646	-0.114
No	0.270	0.255	0.015
Missing	0.198	0.099	0.099
Daytime work:			
Yes	0.840	0.855	-0.015
No	0.160	0.145	0.015
Healthcare facilities (per 1,000)	2.629	2.884	-0.256

Notes: Table A4 reports means of key maternal characteristics and local healthcare supply by household income status, motivating the income-based heterogeneity analysis in Table 3. The sample corresponds to that used in Table 3. Columns (1) and (2) report means for households below and above the median income, respectively, and Column (3) reports the difference (below minus above). Mother's education is categorized as less than college, college or above, or missing. Employment status indicates whether the mother is currently employed, and Daytime work indicates whether the mother works primarily during standard daytime hours among those employed. Healthcare facilities is the number of healthcare facilities per 1,000 vaccination-eligible children in the municipality of residence. Missing values for maternal characteristics are coded with a missing-category indicator and retained in the sample. All statistics are calculated using survey weights.

References for Appendix

Caini, Saverio, Q. Sue Huang, Meral A. Ciblak, Gabriela Kusznierz, Rhonda Owen, Sonam Wangchuk, Claudio Maierovitch Pessanha Henriques, Richard Njouom, Rodrigo A. Fasce, Hongjie Yu, et al. 2015. "Epidemiological and Virological Characteristics of Influenza B: Results of the Global Influenza B Study." *Influenza and Other Respiratory Viruses* 9 (S1): 3–12. https://doi.org/10.1111/irv.12319.

Choi, Yu Jung, Joon Young Song, Seong-Heon Wie, Won Suk Choi, Jacob Lee, Jin-Soo Lee, Young Keun Kim, Shin Woo Kim, Sun Hee Lee, Kyung-Hwa Park, Hye Won Jeong, Jin Gu Yoon, Hye Seong, Eliel Nham, Ji Yun Noh, Hee Jin Cheong, and Woo Joo Kim. 2024. "Real-World Effectiveness of Influenza Vaccine over a Decade during the 2011–2021 Seasons—Implications of Vaccine Mismatch." *Vaccine* 42 (26): 126381. https://doi.org/10.1016/j.vaccine.2024.126381.

Krammer, Florian, Gavin J. D. Smith, Ron A. M. Fouchier, Malik Peiris, Katherine Kedzierska, Peter C. Doherty, Peter Palese, Megan L. Shaw, John Treanor, Robert G. Webster, and Adolfo Garcia-Sastre. 2018. "Influenza." *Nature reviews. Disease primers* (LONDON) 4.

Osterholm, Michael T., Nicholas S. Kelley, Alfred Sommer, and Edward A. Belongia. 2012. "Efficacy and Effectiveness of Influenza Vaccines: A Systematic Review and Meta-Analysis." *The Lancet Infectious Diseases* 12 (1): 36–44. https://doi.org/10.1016/S1473-3099(11)70295-X. PMID: 22032844.

Petrova, Velislava N., and Colin A. Russell. 2018. "The Evolution of Seasonal Influenza Viruses." *Nature Reviews Microbiology* 16 (1): 47–60. https://doi.org/10.1038/nrmicro.2017.118.

Smith, Derek J., Alan S. Lapedes, Jan C. de Jong, Theo M. Bestebroer, Guus F. Rimmelzwaan, Albert D. M. E. Osterhaus, and Ron A. M. Fouchier. 2004. "Mapping the Antigenic and Genetic Evolution of Influenza Virus." *Science* 305 (5682): 371–376. https://doi.org/10.1126/science.1097211.

Weinberg, Geoffrey A., and Peter G. Szilagyi. 2010. "Vaccine Epidemiology: Efficacy, Effectiveness, and the Translational Research Roadmap." *The Journal of Infectious Diseases* 201 (11): 1607–1610. https://doi.org/10.1086/652404.

White, Corey. 2021. "Measuring Social and Externality Benefits of Influenza Vaccination." *J. Hum. Resour.* 56 (3): 749–785. https://doi.org/10.3368/jhr.56.3.1118-9893R2.