

Smart Health and Elderly-Care Services and Labor-Market Outcomes: Evidence from a Multi-period Difference-in-Differences Analysis

Yue Pan

*School of Economics, Xiamen University, Xiamen, China
shellflicker@stu.xmu.edu.cn*

Abstract. China is aging at an unprecedented pace: by 2020, 264 million residents were aged 60 or above, and the traditional family-based model of old-age care is under mounting strain. Between 2017 and 2019, three central ministries jointly designated 33 Smart Health and Elderly-Care Application Demonstration Bases in three batches, covering 20 prefecture-level cities. This paper exploits the staggered designation as a quasi-natural experiment. Combining the China Family Panel Studies (CFPS) from 2014 to 2020 with prefecture-level statistics, a sample of 32,176 working-age adults in 126 cities is constructed, and multi-period difference-in-differences (DID) models are estimated. The pilot is found to raise the annual wage income of working-age residents by 11.4 percent on average, without reducing employment or working hours. The effect emerges only after designation, grows over time, and survives province-by-year fixed effects, the exclusion of long-term-care-insurance pilot cities, propensity-score-matched samples, and a 500-repetition placebo test. The wage gains are concentrated among women (16.8 percent) and rural-hukou workers (15.2 percent), the groups that typically shoulder family caregiving, suggesting that technology-enabled formal care relaxes the family-care constraint on labor supply and job quality.

Keywords: smart elderly care, population aging, difference-in-differences, wages, family caregiving

1. Introduction

Population aging is one of the defining demographic facts of twenty-first-century China [1]. The Seventh Population Census recorded 264 million residents aged 60 and above in 2020, or 18.7 percent of the population [2], and the share continues to rise by roughly one percentage point every two years. Because formal long-term-care institutions remain underdeveloped and publicly provided care is scarce, the burden of supporting the elderly still falls overwhelmingly on families, and in particular on working-age adult children. Understanding how the supply of formal elderly-care services interacts with the labor market is therefore not only a question of social policy but also of labor economics: if caring for aging parents absorbs time, energy and geographic mobility, policies that expand affordable formal care may relax a binding constraint on the labor supply and earnings of the middle generation.

A large literature documents that informal caregiving is costly for caregivers. Carmichael and Charles [3] show that informal carers in Britain face reduced employment opportunities; Lilly, Laporte and Coyte [4] reach similar conclusions for Canada; Van Houtven, Coe and Skira [5] estimate that informal care modestly reduces wages and hours in the United States; and Schmitz and Westphal [6] find persistent negative long-run effects in Germany. Much less is known about the reverse experiment: what happens to labor-market outcomes when the supply of formal, affordable elderly care expands? Existing evidence mostly concerns long-term-care insurance (LTCI), studied for China by Lei, Bai, Hong and Liu [7], among others. Direct evidence on service-side supply shocks, especially in a setting where family care is the default, remains scarce.

This paper provides such evidence by evaluating China's Smart Health and Elderly-Care Application Demonstration program. Between 2017 and 2019, three central ministries jointly selected 33 demonstration bases, located in 20 prefecture-level cities, in three annual batches. The program promotes technology-enabled elderly care: internet-of-things health monitoring, tele-medicine, integrated home- and community-based care platforms, and smart assistive devices. Four waves (2014-2020) of the China Family Panel Studies (CFPS) are combined with prefecture-level statistics. Twelve of the twenty pilot cities appear in the CFPS sample, giving a treated group of five cities designated in 2017, two in 2018 and five in 2019, and a control group of 114 never-treated cities. The baseline sample comprises 32,176 individuals aged 18 to 60. Two-way fixed-effects DID models are estimated with city and year fixed effects, a rich set of individual and city-level controls, and standard errors clustered at the city level.

Three findings emerge. First, the pilot raises the annual wage income of working-age residents by approximately 11.4 percent on average, while employment and weekly working hours are essentially unchanged, indicating that the policy operates on the quality and intensity of paid work rather than on labor-market entry or exit. Second, the effect is genuinely dynamic: event-study estimates are flat before designation, then rise from about 8 percent in the first two post-treatment years to about 20 percent three or more years after designation, consistent with the gradual rollout of services and platforms. Third, the gains are strongly concentrated among women (16.8 percent) and rural-hukou workers (15.2 percent), precisely the demographic groups that shoulder the bulk of family caregiving in China. A battery of robustness checks, including province-by-year fixed effects, exclusion of LTCI pilot cities, dropping the last treatment cohort, propensity-score-matched city samples and a 500-repetition placebo test, supports a causal interpretation.

This paper makes three contributions. First, it provides, to the best of the authors' knowledge, the first labor-market evaluation of China's smart elderly-care pilots using nationally representative micro data, complementing recent work by Xie and Yang [8] on the program's family outcomes. Second, it contributes to the international literature on caregiving and labor supply by studying a plausibly exogenous expansion of formal care supply in a setting where family care is the norm; most existing evidence comes from developed countries with mature care markets. Third, the results carry a policy message that extends beyond elderly care itself: investment in care infrastructure functions, in part, as labor-market policy, and its beneficiaries are disproportionately women and rural workers.

The rest of the paper proceeds as follows: Section 2 gives the institutional background, Section 3 the data, Section 4 the empirical strategy, Section 5 the results, Section 6 heterogeneity and mechanisms, and Section 7 concludes.

2. Institutional background

2.1. Aging and the family-care model in China

China's old-age dependency ratio has risen steadily for two decades, while average household size fell to 2.62 persons in 2020 [2]. The traditional arrangement, under which adult children, and especially daughters-in-law, provide personal care, financial support and health management for aging parents, is increasingly strained by internal migration, smaller sibling networks and women's rising labor-force attachment. Surveys consistently show that a majority of disabled elderly are cared for exclusively by family members; for working-age adults, caring for a parent with functional limitations can absorb the equivalent of a part-time job. At the same time, the formal long-term-care sector remains thin: beds in eldercare institutions cover only a small fraction of the elderly population, home-visit nursing is rare outside major cities, and the LTCI system launched in 2016 still covers only a minority of prefectures [9].

2.2. The smart health and elderly-care application demonstration program

Against this backdrop, the Ministry of Industry and Information Technology (MIIT), the Ministry of Civil Affairs (MCA) and the National Health Commission (NHC) jointly launched the Smart Health and Elderly-Care Application Demonstration program. The program designates demonstration enterprises, demonstration streets (townships) and demonstration bases that integrate information technology into elderly-care services. Typical components include wearable health-monitoring devices linked to community clinics, tele-consultation and chronic-disease management platforms, one-call emergency response systems, and digital platforms that dispatch home-care and rehabilitation services.

Three batches of demonstration bases were announced: 9 bases in 2017 (Document No. 588 of 2017), 7 bases in 2018 (Document No. 482 of 2018), and 17 bases in 2019 (Document No. 415 of 2019), 33 bases in total. Because some bases are districts of the same prefecture city, the bases map into 20 prefecture-level cities. Designation brings policy support, funding channels and, most importantly, a mandate to scale up smart elderly-care services locally; services typically take one to three years to roll out, which motivates the focus on dynamic treatment effects.

Two features of the program make it suitable for a DID design. First, the staggered timing generates both never-treated and not-yet-treated comparison groups. Second, designation decisions were made by a joint expert panel at the central level on the basis of application materials, rather than in response to local labor-market conditions; Section 5 shows that pre-designation wage trends in treated and control cities are statistically indistinguishable. Twelve of the twenty pilot cities are covered by the CFPS: Dalian, Chengdu, Shenyang, Jinan and Zibo (2017 batch); Xuchang and Ganzhou (2018 batch); Lanzhou, Guangzhou, Hangzhou, Wuhan and Xi'an (2019 batch).

2.3. Related literature

Three strands of research frame the analysis. The first documents the labor-market costs of informal caregiving. Studies for Europe, North America and Japan consistently find that intensive caregiving reduces labor-force participation, working hours and wages, and that the effects are larger for women and for caregivers who co-reside with the care recipient [3-6, 10]. A corollary that has received less attention is that expanding the supply of formal care should, symmetrically, improve caregivers' labor-market outcomes; this paper tests this corollary directly.

The second strand evaluates China's long-term-care institutions. Feng et al. [9] review the fragmented landscape of eldercare provision and financing, and evaluations of the 2016 LTCI pilots find effects on medical utilization, health and family transfers [7]. Relative to insurance, which relaxes the budget constraint for purchasing care, the smart elderly-care program acts on the service supply side, and its labor-market consequences have not previously been quantified. Xie and Yang [8] study the same program with a focus on family and health outcomes; the present paper complements their work by shifting the outcome space to employment, hours and earnings of the working-age population.

The third strand is methodological. The credibility of staggered-adoption DID has been reshaped by recent work on treatment-effect heterogeneity across cohorts [11-13]. These concerns are taken seriously here: the design uses only never-treated cities as controls in the matched and placebo exercises, cohort-robust estimates are reported in the replication package, and the results are shown not to be driven by any single treatment cohort.

3. Data and variables

3.1. Data sources

The micro data come from the China Family Panel Studies, a biennial, nationally representative longitudinal survey conducted by the Institute of Social Science Survey at Peking University, covering 25 provinces. The 2014, 2016, 2018 and 2020 waves are used, which provide individual-level information on employment, annual wage income, working hours, demographics and health, together with prefecture identifiers. The pilot-city list is compiled from the three ministerial designation documents. City-level controls (GDP per capita, the urbanization rate, fiscal expenditure over GDP, and the tertiary-industry share) are taken from a prefecture-level panel assembled from official statistical sources.

Three sample restrictions are imposed: individuals aged 18 to 60; survey years 2014-2020, so that every treated city contributes at least one pre-treatment wave; and non-missing prefecture identifiers. The resulting unbalanced panel contains 32,176 person-wave observations in 126 cities. Twelve cities are ever treated (3,448 person-wave observations), and 114 cities are never treated.

3.2. Variables

Four labor-market outcomes are studied. Employment is an indicator for engaging in either non-agricultural or agricultural work in the reference period. Annual wage income is total wage earnings over the past year, winsorized at the 99th percentile by survey year and expressed in logs. Weekly working hours are reported usual hours. Household per-capita income (log) serves as a broader measure of economic well-being. Individual controls include gender, age and its square, years of schooling, urban-hukou status, marital status, self-rated health, medical-insurance coverage and internet use; city controls are log GDP per capita, the urbanization rate, fiscal expenditure relative to GDP, and the tertiary-industry share.

Table 1 reports pre-treatment (2014-2016) descriptive statistics by pilot status. Treated cities are on average richer, more educated and more urbanized than control cities, as expected for centrally designated demonstration sites. These level differences are absorbed by city fixed effects and controls; the identifying assumption concerns trends rather than levels, which is examined directly in the event-study analysis. Importantly, employment, working hours, health and insurance coverage are very similar across the two groups before 2017.

Table 1. Pre-treatment descriptive statistics by pilot status (2014-2016)

Variable	Treated mean	(SD)	Control mean	(SD)	Difference
Employed	0.880	(0.325)	0.889	(0.314)	-0.009
Annual wage (yuan)	20,614	(62,668)	14,575	(25,465)	6,039***
Weekly work hours	46.33	(22.60)	45.71	(22.63)	0.61
Household income p.c. (yuan)	23,912	(57,116)	20,254	(58,956)	3,658***
Age	44.65	(10.32)	44.14	(10.19)	0.52**
Male	0.465	(0.499)	0.495	(0.500)	-0.031**
Years of schooling	9.01	(4.30)	8.04	(4.34)	0.98***
Urban hukou	0.374	(0.484)	0.251	(0.434)	0.123***
Married	0.865	(0.342)	0.889	(0.314)	-0.024***
Self-rated health	2.98	(1.18)	2.98	(1.21)	0.00
Internet user	0.416	(0.493)	0.352	(0.478)	0.065***
Has medical insurance	0.911	(0.284)	0.909	(0.287)	0.002
Elderly (65+) in household	0.157	(0.364)	0.189	(0.392)	-0.032***
Log city GDP p.c.	10.73	(0.28)	10.68	(0.40)	0.05***
Fiscal expenditure / GDP	0.210	(0.083)	0.234	(0.087)	-0.024***
Tertiary-industry share	1.06	(0.20)	1.18	(0.42)	-0.11***

Notes: Sample means with standard deviations in parentheses; the last column reports the difference in group means with significance from a two-sided t-test with unequal variances. *** p<0.01, ** p<0.05, * p<0.1.

4. Empirical strategy

The baseline specification is the standard two-way fixed-effects model for staggered-adoption designs:

$$Y_{ict} = \alpha + \beta DID_{ct} + X'_{it}\gamma + W'_{ct}\delta + \mu_c + \lambda_t + \varepsilon_{ict} \quad (1)$$

where Y_{ict} is an outcome of individual i in city c in wave t ; DID_{ct} equals one for observations in a pilot city from its designation year onward; X_{it} and W_{ct} are the individual and city-level controls listed above; μ_c and λ_t are city and year fixed effects. The coefficient of interest, β , is the average effect of the pilot on the treated cities' working-age residents. Standard errors are clustered at the city level, the level of treatment assignment.

To assess pre-trends and dynamics, the single DID term is replaced with event-time bins relative to the designation year $k = t - T(c)$: k at or below -4, k in [-3,-2], the baseline bin k in [-1,0] (omitted), k in [1,2], and k at or above 3. Because CFPS waves are biennial, the bins correspond roughly to four-or-more years before, two-to-four years before, the designation wave, two-to-three years after, and four-or-more years after designation.

Two identification concerns deserve emphasis. First, recent work shows that two-way fixed-effects estimates can be contaminated by negative weighting when treatment effects are

heterogeneous across cohorts [11, 12]. This concern is addressed by reporting cohort-robust Callaway and Sant'Anna [13] estimates in the replication package, which yield the same sign and significance, and by showing that the result survives dropping the 2019 cohort. Second, treated cities differ in levels from control cities (Table 1). The strategy therefore relies on parallel trends rather than level comparability; it is probed with the event study, a placebo test with 500 random treatment assignments, province-by-year fixed effects, city-specific linear trends, and propensity-score matching of cities on pre-treatment characteristics. Fifteen first-batch LTCI pilot cities (including Guangzhou and Chengdu) are also excluded, so the estimates are not driven by the concurrent LTCI reform.

5. Empirical results

5.1. Baseline estimates

Table 2 reports the baseline DID estimates. For each outcome, a specification without controls (odd columns) and one with the full set of individual and city controls (even columns) are shown; all regressions include city and year fixed effects. The headline result appears in column (4): designation of a smart elderly-care demonstration base raises the annual wage income of working-age residents by 0.114 log points, significant at the 1 percent level, which corresponds to roughly a 12 percent increase. Employment (columns 1-2), weekly working hours (columns 5-6) and log household per-capita income (columns 7-8) show no statistically significant response, and the point estimates are small.

Table 2. Baseline multi-period DID estimates

	Employed		Log wage		Weekly hours		Log HH income p.c.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
DID	0.0162 (0.0104)	-0.0027 (0.0150)	0.0802 (0.0528)	0.1138*** (0.0360)	-1.0539 (1.3006)	-1.6289 (1.1370)	0.0093 (0.0532)	0.0302 (0.0465)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
City FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,176	30,182	11,022	10,496	21,159	19,956	31,312	29,688

Notes: Standard errors clustered at the city level in parentheses. Individual controls: gender, age, age squared, years of schooling, urban hukou, marital status, self-rated health, medical insurance, internet use. City controls: log GDP per capita, urbanization rate, fiscal expenditure/GDP, tertiary-industry share. The wage sample is restricted to employed individuals with positive wage income. *** p<0.01, ** p<0.05, * p<0.1.

The magnitude is meaningful but not implausible: with a treated-sample mean annual wage of about 20,600 yuan in the pre-period, an 11-12 percent increase corresponds to roughly 2,300 yuan per working individual per year, similar in size to the wage penalties associated with intensive informal care [5]. The null results also matter: the policy induces neither labor-market exit among current caregivers nor a measurable reduction in aggregate hours, so the wage gains are unlikely to reflect composition shifts in who works.

5.2. Pre-trends and dynamic effects

Figure 1 and Table 3 present the event-study estimates. The two pre-designation coefficients are small and statistically insignificant (-0.098 and 0.043 log points), so the null hypothesis that treated and control cities shared a common wage trend before the program cannot be rejected. After designation the effect turns positive, reaching 0.083 log points in the first two post-treatment years (significant at the 10 percent level) and 0.202 log points three or more years after designation (significant at the 1 percent level). The monotonically increasing profile is consistent with the gradual rollout of smart elderly-care platforms and services after designation, and argues against the interpretation that the baseline result simply picks up a one-time level shift or anticipation. Employment shows no significant response at any horizon (Table 3, right panel).




Figure 1. Event-study estimates for log annual wage

Table 3. Event-study estimates

Event time	Log wage	Employed
k at or below -4	-0.0984 (0.0694)	-0.0149 (0.0245)
k in [-3,-2]	0.0430 (0.0738)	-0.0035 (0.0136)
k in [1,2]	0.0832* (0.0473)	-0.0128 (0.0120)
k at or above 3	0.2019*** (0.0641)	-0.0201 (0.0142)
Observations	10,496	30,182

Notes: All regressions include the full control set and city and year fixed effects; standard errors clustered at the city level in parentheses. The omitted baseline is k in [-1,0]. *** p<0.01, ** p<0.05, * p<0.1.

5.3. Robustness

Table 4 subjects the wage result to seven specification checks. Column (2) adds province-by-year fixed effects, so that identification comes only from within-province comparisons in the same year; the estimate rises to 0.149 and remains significant at the 1 percent level. Column (3) adds city-specific linear trends: the point estimate falls to 0.051 and becomes imprecise, although it remains positive; with only four post-treatment waves and 126 cities, a trend-adjusted specification is demanding, and when province-by-year effects and city trends are combined in column (4) the estimate recovers to 0.110 ($p = 0.098$). Column (5) excludes the 15 first-batch LTCI pilot cities, addressing the concern that the concurrent LTCI reform drives the result; the estimate is 0.094 and significant at the 5 percent level. Column (6) drops the 2019 cohort, which has the shortest post-treatment window; the estimate is essentially unchanged at 0.112. Column (7) restricts the sample to 34 propensity-score-matched cities (1:2 nearest-neighbor matching on pre-treatment city characteristics); the estimate is 0.088 and significant at the 5 percent level.

Table 4. Robustness checks (dependent variable: log annual wage)

	Baseline	Prov-by-year FE	City trends	Prov-by-yr + trends	Excl. LTCI	Drop 2019 batch	PSM sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DID	0.1138*** (0.0360)	0.1485*** (0.0401)	0.0505 (0.0668)	0.1098* (0.0664)	0.0941** (0.0370)	0.1121** (0.0482)	0.0883** (0.0397)
Observations	10,496	10,496	10,496	10,496	10,019	9,847	3,049

Notes: Each column reports the DID coefficient from a separate regression with the full control set and city and year fixed effects unless stated otherwise. Standard errors clustered at the city level in parentheses. Column (7) uses 1:2 nearest-neighbor propensity-score matching on pre-treatment (2014-2016) city characteristics. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




Figure 2. Placebo test with 500 random treatment assignments

Finally, a placebo test is conducted. In each of 500 replications, 12 of the sample cities are randomly designated as fake pilot cities, assigned fake designation years drawn from 2017-2019, and the baseline model is re-estimated. Figure 2 plots the distribution of the 500 placebo coefficients: it is centered close to zero (mean 0.005, standard deviation 0.070), while the true estimate of 0.114 lies far in the right tail; only 5.4 percent of the placebo coefficients exceed the true estimate. It is therefore unlikely that the result reflects a spurious correlation between arbitrary city-year shocks and wages.

6. Heterogeneity and mechanisms

6.1. Who benefits

If the wage effect of the pilot operates through the relaxation of family-care constraints, it should be largest for the groups that provide the most informal care. In China, as in most countries, caregiving is highly gendered: wives, daughters and daughters-in-law supply the bulk of personal care for the elderly. Rural-hukou workers face an additional constraint, because many migrate for work and must either leave aging parents behind with thin local services or withdraw from better-paying urban jobs. Table 5 splits the sample along these dimensions.

Table 5. Heterogeneity of the wage effect (dependent variable: log annual wage)

	Female	Male	Rural hukou	Urban hukou	Edu 9+ yrs	Edu <9 yrs	Elderly HH	No elderly HH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
DID	0.1679*** (0.0537)	0.0695 (0.0647)	0.1523*** (0.0437)	0.0373 (0.0582)	0.1007*** (0.0334)	0.0135 (0.1009)	0.1340 (0.1056)	0.1192*** (0.0427)
Observations	4,432	6,064	6,937	3,559	7,992	2,504	1,948	8,548

Notes: Each column reports the DID coefficient from a subsample regression with the full control set and city and year fixed effects; standard errors clustered at the city level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

The wage effect is strongly concentrated among women: 0.168 log points (significant at the 1 percent level), against 0.070 and insignificant for men. Similarly, rural-hukou workers gain 0.152 log points (1 percent level), while the estimate for urban-hukou workers is 0.037 and insignificant. Workers with at least nine years of schooling, who are more likely to hold formal wage jobs, gain 0.101 log points; the estimate for the less-educated group is close to zero. Splitting by whether the household contains an elderly member (65 or older) yields point estimates of 0.134 for elderly households and 0.119 for non-elderly households; the former is larger but imprecisely estimated because the subsample is small, so equality of the two coefficients cannot be rejected. A triple-difference specification that interacts the DID term with the elderly-household indicator likewise yields an insignificant differential. The household-elderly split is therefore interpreted as suggestive, and more weight is placed on the sharp gender and hukou contrasts.

6.2. Mechanisms

What behavioral adjustments generate the wage gains? Table 6 examines intermediate outcomes. For women, weekly working hours fall by about 2.8 hours on average (statistically insignificant), while

employment is unchanged for both genders. The pattern of higher wages with stable or slightly reduced hours is consistent with women substituting away from the combination of paid work and unpaid caregiving: when formal care services become available, caregivers can keep their jobs, avoid career interruptions and part-time downgrades, and bargain from a stronger position, all of which raise earnings without necessarily lengthening hours. The international literature documents precisely these margins in reverse: caregiving forces job separations and switches to more flexible, lower-paying work [5, 6], so relieving the care burden should raise the quality and pay of the jobs caregivers can keep.

Table 6. Mechanism checks

	Hours, female	Hours, male	Employ, female	Employ, male	Internet use	Life satisfaction	Log med. exp.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DID	-2.7965 (2.2633)	-0.8465 (1.1791)	0.0008 (0.0178)	-0.0060 (0.0134)	-0.0202 (0.0233)	0.0326 (0.0529)	-0.0185 (0.1045)
Observations	9,045	10,911	15,177	15,005	30,182	19,058	29,967

Notes: Each column reports the DID coefficient from a separate regression with the full control set and city and year fixed effects; standard errors clustered at the city level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Several alternative explanations receive no support in the data. The pilot does not measurably increase individual internet use (column 5), which speaks against the view that the estimates simply pick up a generic digitalization dividend. Life satisfaction (column 6) and household medical expenditure (column 7) are also unchanged, suggesting that the wage effect is not a by-product of large health shocks or of household budget relief through medical savings. A demand-side channel is also plausible: demonstration bases create care-sector jobs and stimulate local services, which can bid up wages for the female and rural-hukou workers who staff them. Disentangling the caregiver-relief channel from this demand channel is left for future work, as the CFPS does not measure informal caregiving hours directly.

7. Conclusion

This paper evaluates the labor-market consequences of China's Smart Health and Elderly-Care Application Demonstration program, exploiting the staggered designation of demonstration bases in 2017, 2018 and 2019 as a quasi-natural experiment. Using CFPS micro data from 2014 to 2020 matched to prefecture-level statistics, multi-period DID estimates show that the pilot raises the annual wage income of working-age residents by about 11 percent, with no detectable effect on employment or hours. The effect follows a clean dynamic profile, flat before designation and rising thereafter, and is robust to province-by-year fixed effects, exclusion of long-term-care-insurance pilot cities, dropping the 2019 cohort, propensity-score matching and a 500-repetition placebo test. Wage gains are concentrated among women and rural-hukou workers, the groups that bear the greatest family-care burden.

The policy implication is twofold. First, by relaxing the family-care constraint, investment in elderly-care infrastructure raises the earnings of working-age adults, especially women; elderly-care policy is, in part, labor-market and gender-equality policy. Second, the technology-enabled service

model appears capable of delivering these gains at scale in a middle-income setting, so continued expansion of the program is likely to yield economic benefits beyond its direct welfare objectives.

The study has limitations. The panel covers only four waves and twelve treated cities, and treatment intensity is unmeasured; the specification with city-specific trends is underpowered, so modest trend differences cannot be ruled out. Nor are informal caregiving hours observed directly, which limits the mechanism analysis. Future work using the 2022 CFPS wave, richer caregiving modules and cohort-robust estimators would help confirm and refine these findings.

Data and code availability

All estimation code and the analysis dataset are available in the accompanying replication package.

References

- [1] Beard, J. R., Officer, A., de Carvalho, I. A., Sadana, R., Pot, A. M., Michel, J.-P., et al. (2016). The world report on ageing and health: A policy framework for healthy ageing. *The Lancet*, 387(10033), 2145–2154.
- [2] National Bureau of Statistics of China. (2021). Main data of the seventh national population census. China Statistics Press (in Chinese).
- [3] Carmichael, F., & Charles, S. (2003). The opportunity costs of informal care: Does gender matter? *Journal of Health Economics*, 22(5), 781–803.
- [4] Lilly, M. B., Laporte, A., & Coyte, P. C. (2007). Labor market work and home care's unpaid caregivers: A systematic review of labor force participation rates, predictors of labor market withdrawal, and hours of work. *The Milbank Quarterly*, 85(4), 641–690.
- [5] Van Houtven, C. H., Coe, N. B., & Skira, M. M. (2013). The effect of informal care on work and wages. *Journal of Health Economics*, 32(1), 240–252.
- [6] Schmitz, H., & Westphal, M. (2017). Informal care and long-term labor market outcomes. *Journal of Health Economics*, 56, 1–18.
- [7] Lei, X., Bai, C., Hong, J., & Liu, H. (2022). Long-term care insurance and the well-being of older adults and their families: Evidence from China. *Social Science & Medicine*, 296, 114745.
- [8] Xie, E., & Yang, Y. (2025). The socio-economic effects of smart health and elderly-care services: Quasi-experimental evidence from the application demonstration pilots. *Management World* (in Chinese).
- [9] Feng, Z., Glinskaya, E., Chen, H., Gong, S., Qiu, Y., Xu, J., & Yip, W. (2020). Long-term care system for older adults in China: Policy landscape, challenges, and future prospects. *The Lancet*, 396(10259), 1362–1372.
- [10] Coe, N. B., & Van Houtven, C. H. (2009). Caring for mom and neglecting yourself? The health effects of caring for an elderly parent. *Health Economics*, 18(9), 991–1010.
- [11] de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
- [12] Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- [13] Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.