Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Impact of artificial intelligence pilot policy on household medical burden: A quasi-experimental study based on the China Family Panel Studies

Yilong Wang
Northwest Minzu University
Aihong He
hah631020@163.com
Lanzhou Bowen College of Science and Technology

Research Article

Keywords: Artificial intelligence policy, Household medical burden, Catastrophic health expenditure, Mechanism channels, Difference-in-differences, Health financial protection

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11490016/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Impact of artificial intelligence pilot policy on household medical burden: A quasi-experimental study based on the China Family Panel Studies

Yilong Wang[1], Aihong He[2,*]

1. School of Management, Northwest Minzu University, Lanzhou, Gansu, China

2. School of Economics and Management, Lanzhou Bowen College of Science and Technology, Lanzhou, Gansu, China

*Corresponding author: Aihong He, School of Economics and Management, Lanzhou Bowen College of Science and Technology, Lanzhou, Gansu, China. Email: hah631020@163.com

Abstract

Artificial intelligence policy is usually judged on innovation and output. What it does to the way households pay for medical care is a different question, and a far less examined one. This paper takes that question up. China's National New Generation Artificial Intelligence Innovation and Development Pilot Zones were designated in three approval cohorts covering 18 cities, and that staggered rollout, paired with four waves of the China Family Panel Studies (2016-2022), gives us a quasi-natural experiment that can be taken to a staggered difference-in-differences design. Designation cuts real household medical expenditure by roughly 37% in the fully adjusted specification ($\beta = -0.470$ on log(1 + expenditure), $\rho < 0.01$; -0.372 with fixed effects alone). Catastrophic health expenditure (CHE) becomes about 6.0 percentage points less likely ($\beta = -0.060$, $\rho < 0.01$; -0.055 with fixed effects alone). Both estimates survive. Alternative clustering, winsorised controls, restricted samples, and placebo assignment of treatment status leave them where they stand. Insurance enrolment rises, household health status improves, and Internet use increases; each carries part of the effect. Gains in financial protection are not confined to particular institutional settings; formal interaction tests reject equality across subsamples only for long-term care insurance pilot status, and the effect is smaller where other health policies are already in force. Industrial policy aimed at AI may therefore double as social protection. How much of it households keep depends on the institutions around them.

Keywords: Artificial intelligence policy; Household medical burden; Catastrophic health expenditure; Mechanism channels; Difference-in-differences; Health financial protection

JEL classification: I12; I13; I15; O33

1 Introduction

Household medical burden remains a prominent policy concern in China. Despite the near-universal coverage achieved by the three basic medical insurance schemes, out-of-pocket payments continue to absorb a sizeable share of total health spending, and illness remains one of the main routes by which families fall into poverty [1]. The incidence of catastrophic health expenditure illustrates the problem. Measured as out-of-pocket payments exceeding 40% of household non-food expenditure, it fell from

14.7% in 2010 to 8.7% in 2018, but it stayed considerably higher in rural than in urban households, at 10.9% compared with 6.6% [2]. Medical costs of this kind force households to cut back on other consumption, sell assets, or borrow, and the consequences can persist for years [3-4]. Reducing household medical burden has therefore become an explicit goal of health policy, and identifying instruments capable of delivering it remains an unfinished task.

A large body of research has examined what drives household medical burden. Studies on the supply side have focused on insurance coverage, system capacity, and provider payment rules [5]. Studies on the demand side have emphasized income, chronic disease, and household composition [2-3, 6-7]. Reforms that expand service capacity and contain costs have been evaluated at length, and their effects on household financial burden remain mixed [8-11]. Work on digital transformation tells an equally equivocal story. Gains from digital service delivery are unevenly captured across socioeconomic groups [12-14], and digital inclusion policies can backfire for disadvantaged populations [15]. So far, most of this literature has looked inside the health system, and comparatively little attention has been paid to policies that originate outside it.

Artificial intelligence has moved to the centre of China's industrial policy over the same period. Beginning in 2019, the central government designated National New Generation Artificial Intelligence Innovation and Development Pilot Zones in successive cohorts, eventually covering 18 cities and regions [16-18]. Health care is among the sectors the policy targets. Assisted diagnosis, chronic disease management, medical resource allocation, and insurance claim review are all named as priority applications [19]. Public investment in artificial intelligence has grown rapidly, and the pilot zones have become one of the main instruments through which the state promotes its diffusion [20-21]. The policy is therefore large enough in scale, and close enough to health care delivery, to plausibly reach the household.

Research on the economic consequences of artificial intelligence has concentrated on productivity, labour markets, and firm performance [22-25]. Studies of AI industrial policy in China have followed a similar pattern, tracing outcomes such as factor shares, green innovation efficiency, and enterprise ESG performance [18, 20-21]. A smaller literature has examined artificial intelligence in health care directly, though mostly at the level of service delivery and policy design rather than household outcomes [19, 26]. So far, however, there has been little discussion about whether AI policy affects the financial burden households carry for medical care. Therefore, this study set out to explore the impact of AI Pilot Zone designation on household medical burden in China.

This study uses four waves of the China Family Panel Studies (CFPS, 2016-2022) and a staggered difference-in-differences design to estimate what AI Pilot Zone designation does to household medical burden. Two outcomes carry the analysis, namely total household medical expenditure and whether the

household crosses the catastrophic threshold. In the fully adjusted specification, real medical expenditure falls by about 37%. Catastrophic health expenditure becomes roughly 6.0 percentage points less likely. Insurance enrolment, household health status, and Internet use all appear to carry part of the effect. The estimates survive alternative clustering, winsorised specifications, restricted samples, and placebo assignment of treatment status.

The main contributions of this study are presented in three aspects. First, this study contributes to the literature on the social determinants of health by treating technology-oriented industrial policy as a structural determinant of household financial protection in its own right [6, 27]. Second, this study extends research on the welfare consequences of artificial intelligence beyond firm-level outcomes, tracing a path from a regional industrial policy down to household financial resilience in health care [22-24]. Finally, this study documents how the effect varies across institutional settings, which adds evidence to the debate over whether technology-oriented industrial strategy in middle-income countries is inclusive [14, 28].

The remainder of this study is organized as follows. Section 2 develops the conceptual framework and sets out the research hypotheses. Section 3 describes the policy background, the data, and the empirical strategy. Section 4 reports the empirical results, covering baseline estimates, identification checks, robustness tests, mechanism analysis, and heterogeneity analysis. Section 5 discusses the findings, and Section 6 concludes with policy implications.

2 Conceptual framework and hypotheses

2.1 Direct effects of artificial intelligence on household medical burden

At bottom, the AI Pilot policy is an institutional attempt to put artificial intelligence to work in public services. Health care is one of them. The policy supports AI tools for assisted diagnosis, chronic disease management, medical resource allocation, and insurance claim review [19]. Reorganise care that way and the financial burden reaching households shifts with it.

Clinical services come first. Intelligent diagnostic systems and decision-support tools are meant to raise diagnostic efficiency, though they also cut repeat examinations and pull back unnecessary provision. Where they work as intended, evidence-based decisions become the default. Supplier-induced demand weakens as a result, and the cost of a single visit falls [29-30].

Access is the second route, and the costs that matter here are often indirect. Teleconsultation, intelligent triage, and online follow-up each push high-quality resources out from large urban hospitals toward primary-level and remote areas. Triage does the sorting work by assessing symptoms, stratifying risk, and steering patients to the right level of care, with travel, accommodation, and waiting costs to follow [31].

Prevention is the third route. Integrating health data, predicting risk, and managing chronic disease intelligently all push delivery away from passive treatment and toward early management. At the system level, AI-assisted surveillance should make it less likely that a preventable condition turns into a high-cost episode [13, 32].

That reasoning yields a single testable hypothesis.

H1: The AI Pilot policy significantly reduces household medical burden.

2.2 Indirect effects of artificial intelligence on household medical burden

Direct service delivery is only part of the story. Indirect effects are possible as well, shifting the economic and informational conditions under which households decide how to finance care. What follows takes up three indirect channels in turn.

The distinction matters because the two families of effects act at different points in the household's decision. Direct effects act on the care encounter itself, changing what a given illness costs to treat. Indirect effects act upstream, on the resource and information environment in which the household weighs that cost and decides how to meet it. A household that is enrolled, healthier, and better informed faces a different financing problem even when the underlying illness is identical. Three such channels are examined below, namely enrolment in social health insurance, health status, and internet use. Each is a plausible consequence of the pilot, and each is a recognised determinant of medical burden in its own right.

2.2.1 The medical insurance channel

Fiscal capacity and digital governance infrastructure both expand under the pilot. Enrolment in social health insurance may therefore become easier for households, since insurance is the main institutional buffer against medical impoverishment. Insurance cuts the out-of-pocket share of spending, and with it the expected level and the variance of household health expenditure [6, 33-34]. Higher enrolment should therefore translate fairly directly into better financial protection [35-36]. Enrolment changes behaviour as well. Insured households make more use of preventive and chronic disease services, which lowers the odds of a high-cost episode [37].

The administrative mechanism deserves emphasis. Enrolment in China's social health insurance is voluntary for many residents, and gaps in coverage arise less from ineligibility than from the transaction costs of registering, renewing, and porting coverage across localities [34]. Digital governance lowers those costs. Where registration, benefit verification, and reimbursement are handled through integrated digital platforms, the marginal household is more likely to be enrolled and less likely to let coverage lapse. Fiscal capacity matters for the same reason, since it underwrites the subsidies that keep participation affordable for low-income households [29]. The pilot therefore plausibly raises enrolment along both an affordability margin and a transaction-cost margin.

The composition of protection shifts as well. Insurance converts an uncertain out-of-pocket exposure into a predictable premium, and that matters for burden even where average spending is unchanged, because catastrophic health expenditure is driven by the upper tail of the spending distribution rather than by its mean [6, 38]. The size of the gain, however, is unlikely to be uniform. Consolidation of urban and rural schemes has improved financial protection unevenly across regions and income groups [7, 11, 39], and the moral-hazard margin runs in both directions. Insurance raises utilisation, but much of the additional care is efficient prevention and chronic management rather than low-value treatment [30].

H2a: AI Pilot Zone designation raises the share of households enrolled in medical insurance, and household medical burden declines as a consequence.

2.2.2 The health status channel

Health status can improve through two routes at once. Access to preventive services reduces the incidence of acute episodes; AI-enabled chronic disease management keeps long-term conditions under better control [26]. Whichever route does the work, better health mechanically reduces the quantity of care required, and with it total expenditure and the probability of crossing the catastrophic threshold [16, 40].

The reasoning is standard health capital reasoning. Health is a durable stock that raises utility directly and lowers the shadow cost of producing it, so an improvement in health reduces the derived demand for curative care [40]. Two consequences follow for burden. The first is arithmetic. Fewer and milder episodes mean fewer services purchased, hence lower spending at any given price [16]. The second is indirect and often overlooked. Healthier households earn more and are less exposed to the health shocks that deplete savings and force borrowing [4, 41]. Capacity to pay therefore rises at the same time as medical need falls, and both movements push in the same direction.

Two qualifications temper the channel. The first concerns how gains are achieved. The return to AI-enabled chronic disease management depends on continuity and adherence rather than on the technology alone, and evaluations of smart-healthcare policy show that effects are concentrated where delivery is well organised [26]. The second concerns measurement. Health status in survey data is largely self-reported, and self-assessed health may respond to the policy environment even where clinical status does not [42]. The mechanism test therefore reads as a test of perceived and realised health together, not of clinical endpoints in isolation.

H2b: AI Pilot Zone designation improves household health status, and household medical burden declines as a consequence.

2.2.3 The Internet use channel

The pilot may also raise household internet use, by speeding up digital infrastructure deployment and lowering access costs. Once online, households reach health information more easily, book appointments

and consultations, and take up the digital financial instruments that smooth health-related spending [43-44]. Informational and transactional capacity of this kind makes it easier to obtain timely care at a fair price [45]. Uniformity should not be assumed, though. Returns to digital access depend on digital literacy, which is unevenly distributed and tends lowest among rural residents and older adults [46-47].

Two margins are at work. The information margin concerns search. Online households compare providers, check prices and drug availability, and encounter prevention messages they would otherwise miss, all of which shortens the path to appropriate care [31, 43]. The transaction margin concerns finance. Digital payment and credit instruments let a household spread a lumpy medical outlay across time, which is exactly the problem catastrophic expenditure poses [44]. Teleconsultation adds a third margin by substituting a remote encounter for a hospital trip, removing travel and waiting costs that fall hardest on rural households [45].

The distributional caveat is not incidental. Digital returns are conditional on capability, and capability is unequally held, so the same increase in connectivity can widen rather than narrow gaps in health consumption between urban and rural, and between younger and older, households [13, 47]. Where the digital divide is binding, the pilot may deliver smaller burden reductions to precisely the households that carry the largest medical burden [28]. The channel should therefore be read as an average effect with a heterogeneous, and possibly regressive, distribution underneath.

H2c: AI Pilot Zone designation raises household internet use, and household medical burden declines as a consequence.

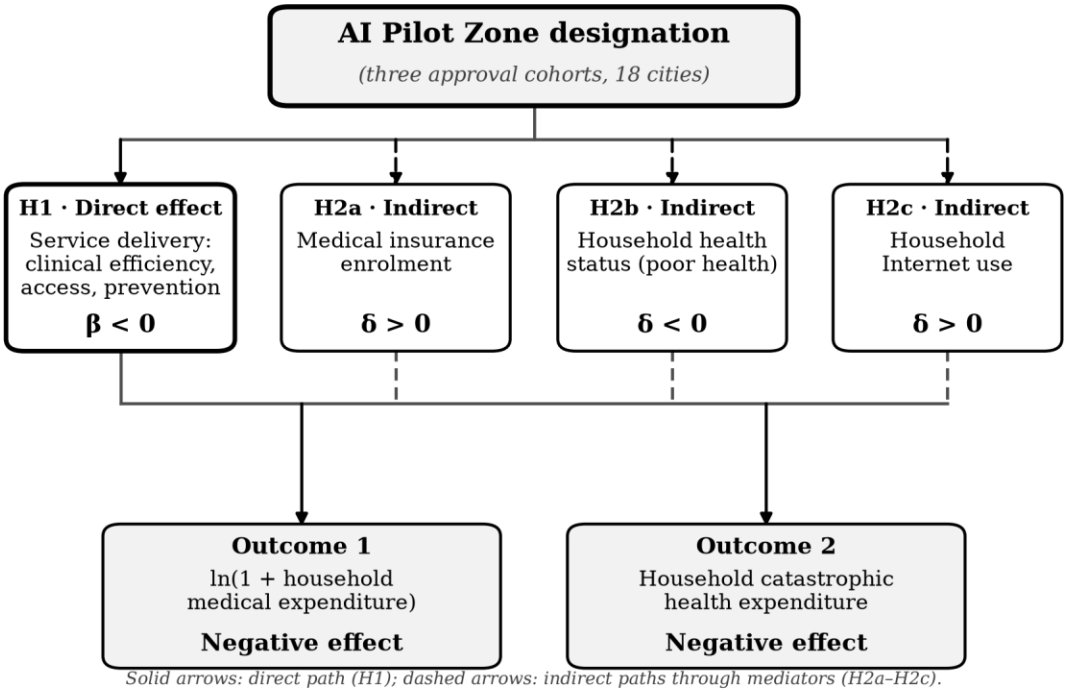


Fig. 1 Research framework

Notes: H1 covers the direct service-delivery effect of AI Pilot Zone designation on household medical burden; H2a-H2c are the three indirect channels tested in the mechanism analysis. Expected signs refer to coefficients on the AI pilot indicator in Equations (1) and (2).

3 Methods

3.1 Study setting and policy background

China began designating National New Generation Artificial Intelligence Innovation and Development Pilot Zones in 2019 [16-18]. Cohorts accumulated across successive years. Two goals drove the programme. One was to accelerate AI innovation; the other, to fold artificial intelligence into economic and social governance. Zones were organised around scenario-based applications, health among them, alongside public services and urban governance. Such a staggered rollout supplies the variation this paper needs. Without cities entering at different times, there would be no way to ask whether regional AI policy reaches household medical burden.

Approvals arrived in three cohorts. Seven zones entered in 2019 (Beijing, Shanghai, Tianjin, Shenzhen, Hangzhou, Hefei, and Deqing County), six in 2020 (Chongqing, Chengdu, Xi'an, Jinan, Guangzhou, and Wuhan), and five in 2021 (Suzhou, Changsha, Zhengzhou, Shenyang, and Harbin). Eighteen in total. Deqing County is a county-level unit, but it was designated under the same national framework and is treated here like the prefecture-level pilot cities. A city counts as treated from its year of

approval onward, and stays treated. CFPS cities never designated during the study period form the control group. Temporal and spatial variation of exactly this kind is what the identification strategy rests on.

Designation is plausibly exogenous to pre-existing household medical burden trajectories, and the selection criteria are the reason. The Ministry of Science and Technology assessed candidate cities mainly on AI industry foundations, technological innovation capacity, data infrastructure, and the availability of application scenarios in manufacturing and urban governance. None of these is conceptually or empirically the same thing as health expenditure outcomes or household financial vulnerability. Regional health system performance, medical expenditure levels, and catastrophic health expenditure incidence played no explicit role. Pilot selection was therefore driven by technology and industrial objectives, not by health system characteristics. Still, designated cities could differ systematically from undesignated ones in ways correlated with medical burden trends, higher income and more developed public service infrastructure being the obvious candidates. Household fixed effects absorb time-invariant heterogeneity, and a PSM-DID check establishes comparability between treated and control households on observable pre-treatment characteristics.

3.2 Data and sample

The data come from the China Family Panel Studies (CFPS), a nationally representative longitudinal survey run by Peking University. Individuals, households, and communities are all tracked, recording demographics, socioeconomic status, health status, and medical expenditure [34, 39], a combination that fits the question better than most alternatives. Four waves are used, namely 2016, 2018, 2020, and 2022. Together they straddle policy implementation. Policy information is matched to households by city of residence and survey year. Observations with missing values on key variables, or with unmatched city information, are dropped; what remains is the final sample of 25294 observations used throughout. Analysis at household level is the right choice here. Medical expenditure is typically a family decision, and medical burden is better captured by the economic pressure a household absorbs as a whole than by anything measured at the individual level.

3.3 Empirical strategy

Estimation uses a multi-period difference-in-differences model [48-49], which compares medical burden before and after pilot approval between treated and untreated cities:

$$Y_{ict} = \alpha + \beta DID_{ct} + \gamma X_{ict} + \mu_i + \lambda_t + \varepsilon_{ict} \quad (1)$$

Y_{ict} is household medical burden for household i in city c in year t, DID_{ct} takes the value 1 if city c has been approved as a pilot zone by year t and 0 otherwise, and X_{ict} collects individual and householdlevel controls; μ_i and λ_t are household and year fixed effects. Here β is the average treatment effect. Read the sign before the size. A negative β means designation reduces household medical burden. Parallel trends are checked with an event-study analysis of pre-policy coefficients [50]. Because the

catastrophic health expenditure outcome is binary, it is estimated as a linear probability model, which allows household and year fixed effects to be included directly; standard errors are clustered at the city level throughout.

3.4 Mechanism specification

The channels are tested with a parallel specification:

$$M_{ict} = \alpha + \delta DID_{ct} + \gamma X_{ict} + \mu_i + \lambda_t + \varepsilon_{ict} \quad (2)$$

M_{ict} is the mechanism variable, and the coefficient δ captures what designation does to the proposed mediator. H2a-H2c give rise to three mediators. Household enrolment in medical insurance, household health status, household Internet use. Where δ is significant and carries the predicted sign, the corresponding channel holds up.

3.5 Measures

3.5.1 Dependent variables

Two variables capture household medical burden, and they do different jobs. First comes the logarithm of one plus real household medical expenditure, lnmed, deflated to constant prices with the health care consumer price index and shifted by one so that households reporting zero expenditure stay in the sample. It picks up the intensive margin of routine spending. Then there is a catastrophic health expenditure indicator, che40, equal to 1 when out-of-pocket health spending reaches or exceeds 40% of household capacity to pay, or of total household consumption depending on the definition adopted [6], and 0 otherwise. lnmed tells us what a household typically spends. che40 tells us how likely it is that spending pushes the household into financial distress. For financial protection, the second is the outcome of primary interest.

3.5.2 Core explanatory variable

DID_{ct} equals 1 for households in pilot cities from the year of approval onward, and 0 otherwise. Note what it measures. City-level policy exposure to AI development, not household-level technology adoption. Staggered approvals allow comparisons across treated, not-yet-treated, and untreated cities.

3.5.3 Mechanism variables

One mechanism variable is built for each channel. Medical insurance is an indicator equal to 1 if any household member is enrolled in a social health insurance scheme and 0 otherwise, standing for the institutional buffer against medical expenditure risk. Health status is a household-level indicator of poor self-reported health, capturing underlying morbidity pressure. Internet use is an indicator equal to 1 if any household member uses the Internet and 0 otherwise, proxying the household's informational and transactional capacity to reach health services.

3.5.4 Control variables

The control set follows the household medical burden literature [7, 14], collecting household head and household characteristics that plausibly affect medical spending on their own, including age and its square, years of education, marital status, gender, rural hukou status, household size, urban residence, housing loan status, outstanding housing debt, non-housing debt status, and self-employment status. Monetary variables are all deflated to constant prices.

3.6 Descriptive statistics

Table 1 reports the descriptive statistics. Mean ln(1 + real medical expenditure) is 6.998 with a standard deviation of 2.821, and the median, at 7.503, sits above the mean. The ordering is a familiar feature of medical spending data: a sizeable share of households report no medical expenditure in a given wave, and the log of one plus expenditure maps those zeros to the bottom of the support, which pulls the mean below the median without distorting the upper part of the distribution. The observed range runs from zero to 13.998, and it is that right tail which the logarithmic transform compresses. Catastrophic health expenditure is recorded in 10.6% of household-wave observations, so the outcome that speaks directly to financial protection is a minority event; its mean is an incidence rate, not a level, and the linear probability model in Table 2 is estimated on that interpretation.

Three features of the covariates bear on how the estimates should be read. First, the treatment indicator has a mean of 0.054, so treated observations make up about one household-wave in twenty and the comparison rests on a comparatively small treated group against a large untreated one. That is what a staggered design supplies, and it is also why standard errors are clustered at the city level and why the wild cluster bootstrap is reported alongside the baseline. Second, medical insurance coverage already stands at 93.7%, close enough to the ceiling that the room for further expansion along that margin is narrow. Internet use, at 50.5%, and poor household health, at 17.1%, both lie well away from their bounds and therefore carry the variation a channel requires. Third, household heads average 52.1 years of age and 7.6 years of schooling, household size averages 3.65, and 76.4% of households hold a rural hukou while 48.8% live in an urban area. The gap between the last two reflects rural-to-urban migration, in which registration does not follow residence, and it means the two indicators are not mirror images of each other.

Table 1 Descriptive statistics of the main variables

Variable	N	Mean	S.D.	Min	Median	Max
ln(1 + real medical expenditure)	25294	6.998	2.821	0.000	7.503	13.998
Catastrophic health expenditure	25294	0.106	0.308	0.000	0.000	1.000
AI pilot (DID)	25294	0.054	0.225	0.000	0.000	1.000
Age of household head	25294	52.105	13.391	16.000	52.000	95.000
Age squared / 100	25294	2894.248	1416.596	256.000	2704.000	9025.000
Years of education	25294	7.595	4.752	0.000	9.000	22.000
Married (=1)	25294	0.859	0.348	0.000	1.000	1.000
Male (=1)	25294	0.553	0.497	0.000	1.000	1.000
Rural hukou (=1)	25294	0.764	0.425	0.000	1.000	1.000
Housing loan (=1)	25294	0.202	0.401	0.000	0.000	1.000
Household size	25294	3.653	1.822	1.000	3.000	19.000
Urban residence (=1)	25294	0.488	0.500	0.000	0.000	1.000
Housing debt (yuan)	25294	4.04e+04	1.94e+05	0.000	0.000	1.01e+07
Non-housing debt (=1)	25294	0.179	0.384	0.000	0.000	1.000
Self-employed (=1)	25294	0.093	0.290	0.000	0.000	1.000
Medical insurance (=1)	25294	0.937	0.243	0.000	1.000	1.000
Poor household health (=1)	25294	0.171	0.377	0.000	0.000	1.000
Internet use (=1)	25294	0.505	0.500	0.000	0.000	1.000

Notes: The sample matches the baseline estimation sample (CFPS panel, four waves 2016-2022, N = 25294) and excludes singleton observations. lnmed is the log of real medical expenditure deflated by the health care consumer price index; che40 indicates catastrophic health expenditure.

4 Results

4.1 Baseline results

The baseline estimates appear in Table 2. Across all four specifications, for both outcomes, with and without controls, the coefficient on the AI pilot indicator is negative and significant at the 1% level. Designation reduces household medical burden, as H1 predicts. Magnitudes matter here. With household and year fixed effects alone, the estimate is −0.372 for ln(1 + real medical expenditure) (column 1) and −0.055 for the catastrophic health expenditure indicator (column 3). Adding the household head and household controls gives −0.470 (column 2) and −0.060 (column 4); these fully adjusted estimates are the preferred ones and are what the abstract reports. Adding observables raises the expenditure coefficient by about 27%, from an implied 31.0% reduction to 37.5%, while the catastrophic expenditure coefficient rises by about 8%. Both remain negative, significant, and of the same order of magnitude, so the conclusion does not hinge on the control set, though the expenditure estimate is the more control-sensitive of the two and gains precision with controls, its standard error falling from 0.141 to 0.104.

The coefficients need translating before they mean much. Because the expenditure outcome is a logarithm of one plus expenditure, −0.470 corresponds to a proportional reduction of $\exp(-0.470) - 1 \approx -37.5\%$ in (1 + expenditure), and therefore to a reduction of roughly 37% in the expenditure level. At the pre-policy mean of treated households (7095 yuan in constant prices), that is about 2660 yuan per household per year. The catastrophic health expenditure coefficient needs no transformation, since it comes from a linear probability model, so −0.060 is a reduction of about 6.0 percentage points in incidence, against a pre-policy incidence of 9.8% among treated households. Both figures are average

effects over the observed post-policy window, which for the earliest-adopting cohort spans two survey waves.

Controls behave as expected. Education, marital status, household size, indebtedness, and rural hukou status are all positively associated with medical expenditure in the fully specified model, while self-employment is indistinguishable from zero. These patterns line up with earlier evidence on the socioeconomic correlates of household medical spending [7, 51].

Table 2 Baseline regression results

Variable	(1) lnmed	(2) lnmed	(3) CHE	(4) CHE
DID	-0.372***	-0.470***	-0.055***	-0.060***
	(0.141)	(0.104)	(0.013)	(0.012)
Age of household head	–	-0.071	–	-0.004
		(0.077)		(0.007)
Age squared / 100	–	0.001***	–	0.000
		(0.0003)		(0.000)
Years of education	–	0.035	–	-0.002
		(0.022)		(0.002)
Married (=1)	–	0.523***	–	0.032*
		(0.143)		(0.016)
Male (=1)	–	0.334	–	0.031
		(0.686)		(0.151)
Rural hukou (=1)	–	0.289***	–	0.016
		(0.083)		(0.010)
Housing loan (=1)	–	0.279***	–	-0.003
		(0.053)		(0.007)
Household size	–	0.174***	–	-0.001
		(0.022)		(0.003)
Urban residence (=1)	–	0.209	–	-0.013
		(0.134)		(0.015)
Housing debt (yuan)	–	0.000	–	0.000
		(0.000)		(0.000)
Non-housing debt (=1)	–	0.473***	–	0.044***
		(0.062)		(0.008)
Self-employed (=1)	–	-0.004	–	-0.016
		(0.096)		(0.010)
Household fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
R^2	0.625	0.631	0.537	0.539
Obs	25294	25294	25294	25294

Notes: Standard errors are cluster-robust at the city level and appear in parentheses. Columns (1) and (2) use ln(1 + real medical expenditure); columns (3) and (4) use the catastrophic health expenditure indicator. Household and year fixed effects only in columns (1) and (3), with the household head and household controls added in columns (2) and (4). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4.2 Parallel trend test

4.2.1 ln(1 + household medical expenditure)

The staggered design requires that, absent the policy, treated and control households would have followed parallel trajectories in medical burden, and the event-study estimates in Fig. 2 speak to that requirement for ln(1 + household medical expenditure). Pre-treatment coefficients sit close to zero. None

is statistically distinguishable from it at conventional levels, which supports the parallel trend assumption. After designation the estimates turn negative and become significant, pointing to a decline in household medical expenditure once the policy is in force.

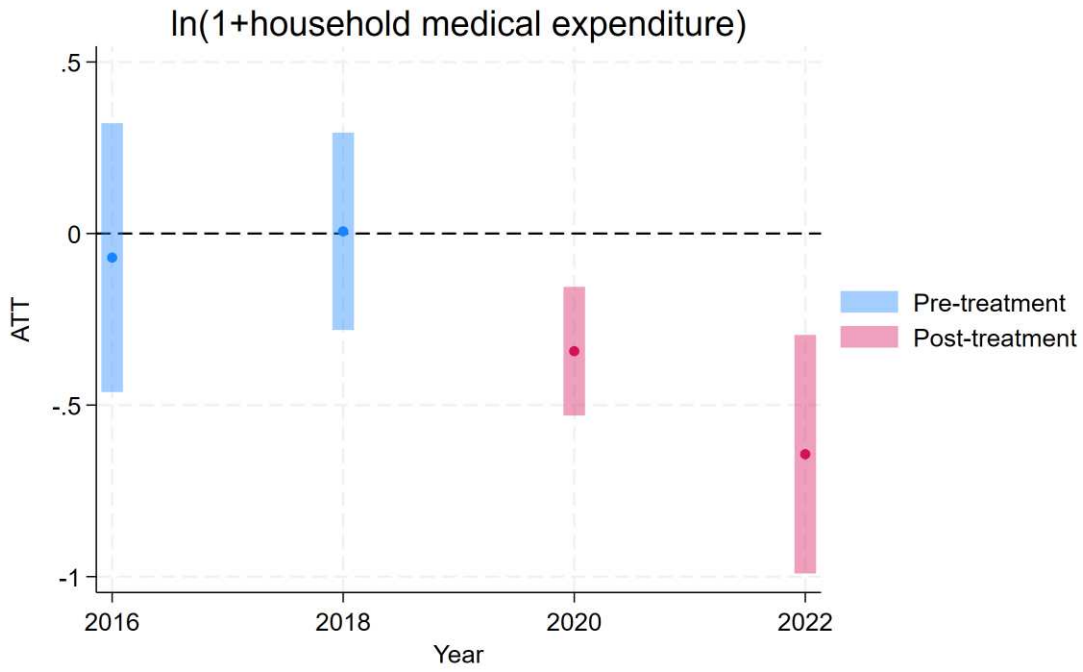


Fig. 2 Parallel trend test for ln(1 + household medical expenditure)

4.2.2 Household catastrophic health expenditure

Fig. 3 gives the corresponding estimates for catastrophic health expenditure. Here the pattern repeats. Pre-treatment coefficients are insignificant, and once the policy takes effect they turn negative and significant. Neither outcome shows a discernible pre-trend, which strengthens the causal reading.

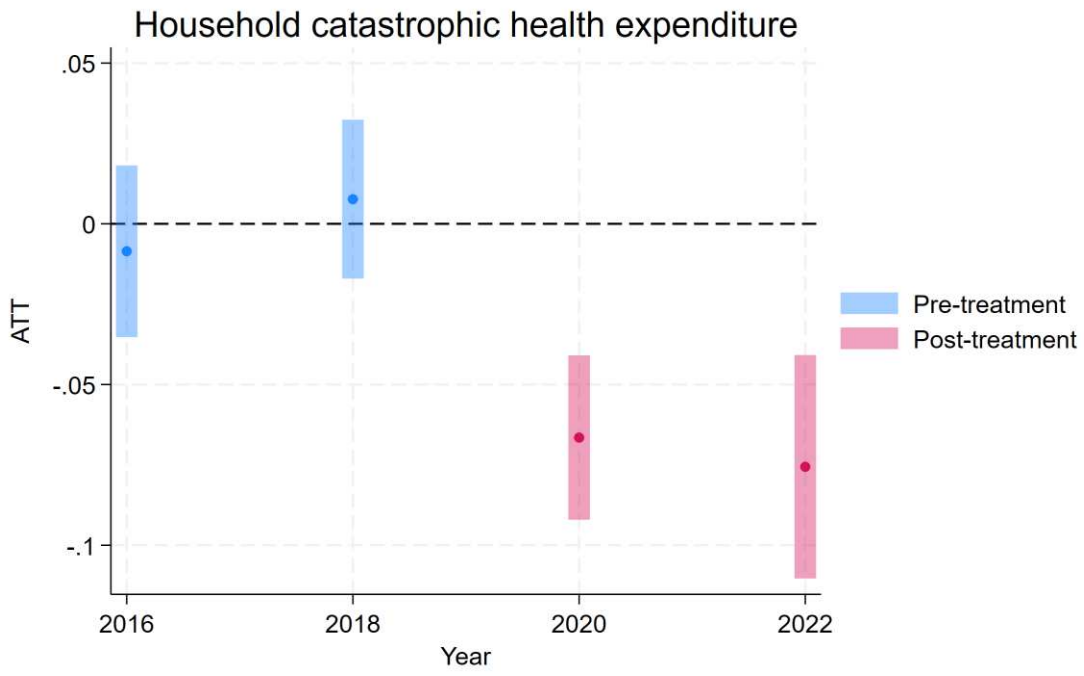


Fig. 3 Parallel trend test for household catastrophic health expenditure

4.3 Placebo test

4.3.1 ln(1 + household medical expenditure)

Could the estimates simply reflect spurious correlation? A placebo test addresses that. Pseudo pilot status is assigned to cities at random, 500 times over, and the baseline specification re-estimated each time. Fig. 4 shows the resulting distribution for ln(1 + household medical expenditure). Estimates centre on zero, and the actual estimate falls in the far left tail, well away from anything random variation would produce. The empirical p-value, meaning the share of the 500 simulated estimates at least as large in absolute value as the observed one, is 0.002.

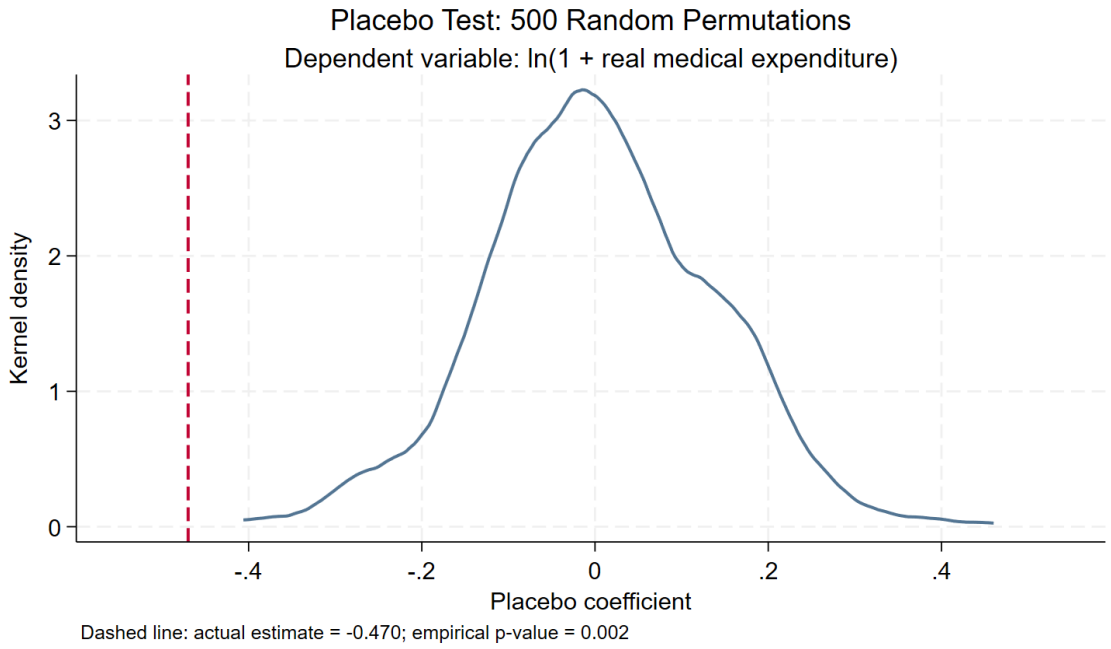


Fig. 4 Placebo test for ln(1 + household medical expenditure)

4.3.2 Household catastrophic health expenditure

Fig. 5 repeats the exercise for catastrophic health expenditure. Again the placebo estimates cluster around zero, and the actual estimate lies outside the bulk of the simulated distribution. The empirical p-value is 0.002. Nothing here unsettles it.

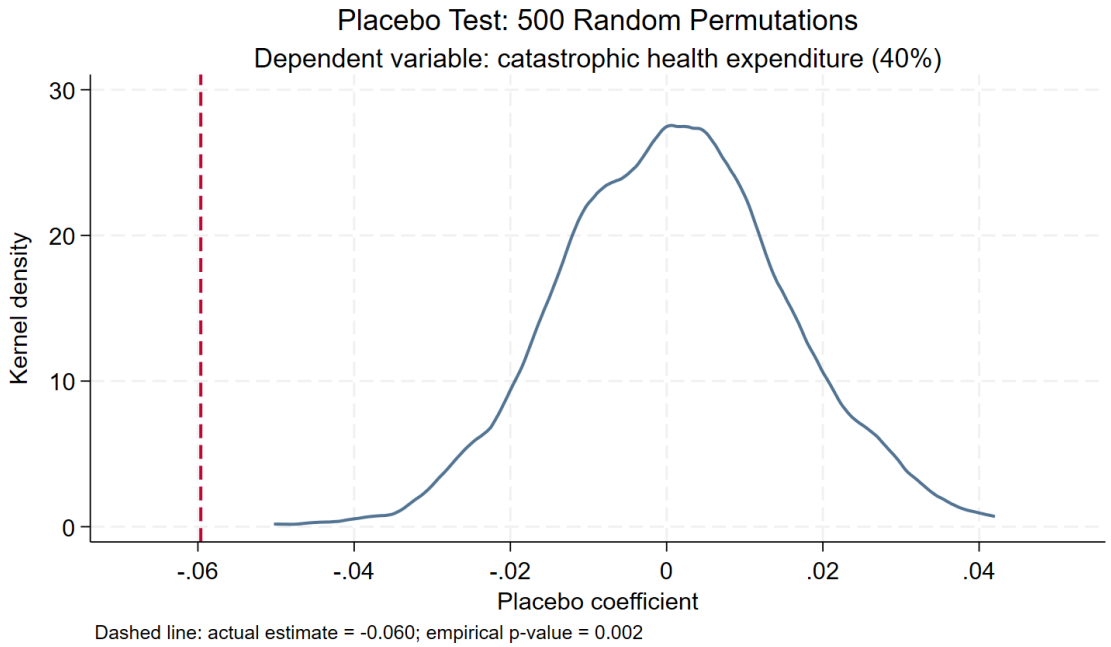


Fig. 5 Placebo test for household catastrophic health expenditure

4.4 Robustness checks

4.4.1 ln(1 + household medical expenditure)

Table 3 puts ln(1 + household medical expenditure) through five robustness checks. Standard errors are clustered at the province rather than the city level in column (1), the controls are winsorised at the 1st and 99th percentiles in column (2), the directly administered municipalities are dropped in column (3), the 2020 wave is dropped in column (4), and column (5) keeps only the 2020 pilot cohort. Across all five, the coefficient on the AI pilot indicator stays negative and significant at the 1% level, ranging from −0.451 to −0.663.

Table 3 Robustness checks: ln(1 + household medical expenditure)

Variable	(1) Province clustering	(2) Winsorised controls	(3) Excl. municipalities	(4) Excl. 2020 wave	(5) 2020 cohort only
DID	-0.470***	-0.468***	-0.600***	-0.663***	-0.451***
	(0.111)	(0.103)	(0.097)	(0.108)	(0.120)
Controls	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
R^2	0.631	0.631	0.632	0.658	0.630
Obs	25294	25294	23231	17742	23902

*Notes: The dependent variable is ln(1 + real medical expenditure). Parentheses hold cluster-robust standard errors, clustered at the province level in column (1) and at the city level elsewhere. Column (2) applies 1%/99% winsorisation to the controls. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

4.4.2 Household catastrophic health expenditure

Table 4 runs the same five specifications on catastrophic health expenditure. Estimates are negative and significant at the 1% level throughout, between −0.060 and −0.065. Both sets of results hold together under different clustering schemes, outlier treatments, sample restrictions, and cohort definitions, which is what makes them worth reporting. Baseline conclusions do not hinge on modelling choices.

Table 4 Robustness checks: household catastrophic health expenditure

Variable	(1) Province clustering	(2) Winsorised controls	(3) Excl. municipalities	(4) Excl. 2020 wave	(5) 2020 cohort only
DID	-0.060***	-0.060***	-0.063***	-0.065***	-0.062***
	(0.013)	(0.012)	(0.012)	(0.014)	(0.017)
Controls	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
R^2	0.539	0.539	0.539	0.587	0.539
Obs	25294	25294	23231	17742	23902

*Notes: The dependent variable is the catastrophic health expenditure indicator; column definitions match Table 3. Parentheses carry cluster-robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

4.5 Mechanism analysis

Table 5 carries estimates of Equation (2) for the three mechanism variables proposed in Section 2.2. Designation raises the probability of household enrolment in medical insurance, as column (1) shows, with an estimated effect of 0.077 significant at the 1% level. H2a is consistent with that. Insurance enrolment is the principal institutional buffer against medical expenditure risk [6, 34], so the channel plausibly accounts for part of the reduction in medical burden.

Poor health status falls, and column (2) puts the effect well inside conventional significance, consistent with H2b. Coding runs the other way here, and higher values mean worse health, which makes the negative coefficient a sign of improved household health; that in turn lowers routine expenditure and the probability of crossing the catastrophic threshold. Column (3) turns the other way, with a significant positive effect on Internet use, consistent with H2c. Better access to health information and to digital transaction channels makes care-seeking more timely and less costly [43-44]. None of the three channels competes with the others. Taken together, they point to the AI Pilot policy reducing household medical burden both through the supply-side service improvements discussed in Section 2.1 and by building up household-level insurance coverage, health endowments, and informational capacity.

Table 5 Mechanism analysis

Variable	(1) Medical insurance	(2) Health status	(3) Internet use
DID	0.077***	-0.103***	0.127***
	(0.012)	(0.010)	(0.028)
Controls	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.517	0.675	0.758
Obs	25294	25294	25294

*Notes: Each column reports the estimated effect of AI Pilot Zone designation on the indicated mechanism variable, with city-level cluster-robust standard errors in parentheses. Throughout, the full control set and household and year fixed effects are included. Note the coding of the health status indicator. Higher values denote poorer health. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

4.6 Heterogeneity analysis

Does the protective effect hold uniformly across institutional contexts? Its answer matters for policy design, and three dimensions speak to it. Each marks a concurrent policy that might complement the AI pilot, or substitute for it, namely long-term care insurance (LTCI) pilot status, participation in the centralised drug procurement ("4+7") pilot, and designation under the Western Development strategy.

4.6.1 ln(1 + household medical expenditure)

Table 6 reports subsample estimates for ln(1 + household medical expenditure). Every subsample returns a negative and significant coefficient on the AI pilot indicator, so the expenditure-reducing effect does not depend on institutional context. Point estimates run largest where no LTCI pilot operates concurrently (−0.607) and where procurement pilot status is absent (−0.658). They run smallest where those interventions are already in place (−0.530 and −0.199). Formal tests of cross-subsample differences qualify all of this. Only the LTCI dimension yields a significant interaction (p = 0.046); the rest cannot be distinguished from sampling variation. Heterogeneity in this table is best read as suggestive.

Table 6 Heterogeneity analysis: ln(1 + household medical expenditure)

Variable	(1) Non-LTCI	(2) LTCI pilot	(3) Non-procurement	(4) "4+7" procurement	(5) Non-western	(6) Western
DID	-0.607***	-0.530**	-0.658***	-0.199*	-0.436***	-0.687***
	(0.089)	(0.175)	(0.175)	(0.100)	(0.113)	(0.179)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.632	0.617	0.632	0.624	0.637	0.612
Obs	22770	2524	21911	3383	17933	7361

*Notes: Standard errors are cluster-robust with clustering at the city level, reported in parentheses, and each column is an independent subsample regression carrying the full control set and household and year fixed effects. Cross-subsample differences are best assessed from the interaction-term p-values, and only the LTCI dimension yields a significant difference (p = 0.046). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

4.6.2 Household catastrophic health expenditure

Table 7 reports the matching estimates for catastrophic health expenditure. Much the same shape appears here, and coefficients are negative throughout, and significant in every subsample except the smallest, the LTCI pilot cities (n = 2524), where the estimate stays negative but is not distinguishable from zero. Point estimates run somewhat larger in non-pilot subsamples. Here too the formal interaction tests fail to support significant heterogeneity in most dimensions. Read alongside Table 7, the results suggest a financial protection effect that is widely shared rather than concentrated in particular institutional settings, though concurrent health policies may still moderate its size.

Table 7 Heterogeneity analysis: household catastrophic health expenditure

Variable	(1) Non-LTCI	(2) LTCI pilot	(3) Non-procurement	(4) "4+7" procurement	(5) Non-western	(6) Western
DID	-0.070***	-0.048	-0.075***	-0.070***	-0.062***	-0.063**
	(0.016)	(0.025)	(0.023)	(0.021)	(0.013)	(0.027)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.539	0.520	0.537	0.549	0.542	0.532
Obs	22770	2524	21911	3383	17933	7361

*Notes: As for Table 6. Dependent variable: catastrophic health expenditure indicator. Cluster-robust standard errors, city-level clustering, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

5 Discussion

The question this study set out to answer was whether a state-led AI industrial policy produces measurable welfare effects at the household level, in the specific domain of health-related financial risk. Three answers emerge. In the fully adjusted specification, real household medical expenditure falls by about 37% and the probability of catastrophic health expenditure by about 6 percentage points, and the estimates survive alternative clustering, winsorised specifications, restricted samples, and placebo assignment. The effect travels through three detectable channels, namely expanded medical insurance enrolment, improved household health status, and greater Internet use. And it appears across every institutional subsample examined, though point estimates are smaller where concurrent health policies are already operating.

This first result speaks to a growing literature on what AI-oriented industrial policy does economically. Existing work covers firm-level outcomes such as factor shares and green innovation efficiency [21, 24], corporate sustainability [20], regional carbon and environmental outcomes [18], and the well-being of older adults [52]. Our results push the evidence one step further. Consequences of AI policy reach past the firm and the region, into household financial protection in health care. That matters. Financial protection is a core objective of health systems everywhere, and a central marker of progress toward universal health coverage [1, 6]. A lever of industrial policy that improves it is worth knowing about.

Transmission is the second result. Insurance (H2a) fits evidence that coverage is the primary institutional safeguard against medical impoverishment [7, 34, 38]. Health status (H2b) fits the health capital framework, in which better health lowers both the demand for and the cost of care [40]. Internet use (H2c) echoes findings that digital engagement improves access to health information and to services, while also enabling the digital financial instruments that smooth health spending [43-44]. These three are not mutually exclusive; they appear to run together. This is one reason the aggregate effect is detectable at all in a household survey, where any single channel might account for only a modest share of variation. Evidence on complementary coverage layers points the same way, showing lower out-of-pocket ratios

and less catastrophic expenditure [53-55]. Long-term care insurance, similarly, has been found to reduce out-of-pocket burdens among older adults [56]. More broadly, the joint operation of insurance, service utilisation, and infrastructure echoes recent work arguing that health equity emerges from how these components interact, not from any one of them alone [57-58].

Care is needed with the third result. Point estimates vary across subsamples, but formal interaction tests reject equality only for long-term care insurance. Read prudently, the financial protection effect is broadly shared, not concentrated in particular groups or places. Such a reading sits awkwardly with studies emphasising distributional skew in the benefits of digital transformation [13-14]. Mechanism may explain the difference. AI pilot designation works mainly through area-level infrastructure and fiscal capacity, not through individual technology adoption. Where realised benefit does depend on individual digital skills, as with Internet-mediated service access, macro-level policy gains may still be captured unevenly. That possibility deserves direct investigation. Several strands reinforce it. Urban-rural inequality of opportunity in Chinese health care persists [59]; the income-health gradient remains pronounced [41]; and medical expenditure trajectories diverge substantially across regions even within a single health system [42, 60]. Under the inverse care law, the populations with the greatest need tend to receive the least effective care [61]. What ultimately matters is not the availability of technology but whether households can convert it into valued health outcomes. Digital platforms can generate health benefits and new spending exposure at the same time [62], and the digital economy reshapes rural household resilience and income uncertainty in ways that bear directly on the capacity to absorb a health shock [4, 63-64].

A few limitations deserve mention. Selection on unobserved city-level trajectories cannot be fully excluded, even though designation followed technology and industrial criteria rather than health system characteristics; the parallel trend and placebo tests mitigate the concern without eliminating it. The post-policy window is short. AI diffusion and health system adaptation both unfold gradually, so longer-run effects may differ in size. CFPS measures medical expenditure by self-report, which is subject to recall error, though this is common to survey-based health expenditure research. The mechanism analysis documents effects on candidate mediators rather than estimating a fully mediated structural model, so the relative contribution of each channel cannot be decomposed. Finally, the findings come from a single middle-income country with a distinctive model of state-led technology policy, and generalisation beyond it should be cautious.

6 Conclusions

Artificial intelligence industrial policy can improve household financial protection in health care. That is the causal claim this study supports. Exploiting the staggered rollout of China's National New Generation Artificial Intelligence Innovation and Development Pilot Zone policy alongside four waves of

nationally representative household panel data, we find that designation reduces real household medical expenditure by roughly 37% and the probability of catastrophic health expenditure by roughly 6 percentage points, both in the fully adjusted specification. The estimates hold under alternative clustering, winsorised specifications, restricted samples, and placebo testing. The path from a regional technology policy to household financial resilience runs through medical insurance enrolment, household health status, and Internet use.

Three policy implications follow. Technology-oriented industrial policy should be judged on its welfare consequences for households, not only on innovation and productivity returns; the results here suggest AI policy can serve as a complementary instrument of social protection, a dimension largely missing from current evaluation frameworks. Because the channels run partly through insurance enrolment and health status, coordination between technology policy and health policy is likely to shape what households actually get; governments seeking financial protection gains should pair AI infrastructure investment with measures that expand coverage and strengthen primary care. And because the aggregate effect appears broadly shared, AI pilot policies need not be presumed inherently regressive. Complementary investment in digital literacy and rural primary care capacity remains warranted all the same, so that households can turn area-level policy inputs into realised benefits.

Several extensions suggest themselves. Longer post-policy horizons can be examined as further survey waves appear, and the relative contribution of each transmission channel can be decomposed. Whether comparable effects arise from AI policies in other national contexts also remains to be tested. Work along these lines would clarify when technological change translates into better financial protection for households.

Funding

This work was supported by the Soft Science Special Project of Gansu Basic Research Plan (Grant No. 26JRZA025), the Fundamental Research Funds for the Central Universities of Northwest Minzu University (Grant No. 31920250127), the Postgraduate Education and Teaching Reform Funds of Northwest Minzu University (Grant No. 2026JGYB13), and the Western Project of the National Social Science Fund of China (Grant No. 25XJY002).

Competing interests

The authors declare no competing interests. No commercial or financial relationship could be construed as a potential conflict of interest.

Ethics approval

This analysis relies on de-identified, publicly available survey data, collected under a CFPS protocol approved by the Institutional Review Board of Peking University, with informed consent obtained from every participant. No further approval was needed here.

Data availability

The CFPS data used here are publicly available. They come from the Institute of Social Science Survey at Peking University (https://www.isss.pku.edu.cn/cfps/). AI Pilot Zone designation dates were compiled by the authors from official announcements issued by the Ministry of Science and Technology of the People's Republic of China.

Author contributions

Yilong Wang: Conceptualization, Data curation, Formal analysis, Methodology, Writing – original draft. Aihong He: Supervision, Validation, Writing – review and editing. Both authors read and approved the final manuscript.

References

1. Liang, W., Tang, K., Zhu, J., Li, Z., Yip, W., Lin, V., Atun, R., Chan, M.: Tsinghua–Lancet commission on health and poverty alleviation in China. Lancet 403, 886–887 (2024). https://doi.org/10.1016/s0140-6736(23)01856-1

2. Liu, C., Liu, Z. M., Nicholas, S., Wang, J.: Trends and determinants of catastrophic health expenditure in China 2010–2018: A national panel data analysis. BMC Health Serv. Res. 21, 526 (2021). https://doi.org/10.1186/s12913-021-06533-x

3. Feng, Q., Che, Y., Yi, S., Wang, Y., Chen, W., Xia, X.: Analysis of medical impoverishment and its influencing factors among China's rural near-poor, 2016–2020. Front. Public Health 12, 1412536 (2024). https://doi.org/10.3389/fpubh.2024.1412536

4. Qiu, Y., Zhang, F.: Impact of health shocks on household consumption structure. Front. Public Health 12, 1431833 (2024). https://doi.org/10.3389/fpubh.2024.1431833

5. Zhao, N., Chen, K.: Equity and efficiency of medical and health service system in China. BMC Health Serv. Res. 23, 33 (2023). https://doi.org/10.1186/s12913-023-09025-2

6. Wagstaff, A., Flores, G., Hsu, J., Smitz, M. F., Chepynoga, K., Buisman, L. R., van Wilgenburg, K., Eozenou, P.: Progress on catastrophic health spending in 133 countries: A retrospective observational study. Lancet Glob. Health 6, e169–e179 (2018). https://doi.org/10.1016/S2214-109X(17)30429-1

7. Zhou, Q., He, Q., Eggleston, K., Liu, G. G.: Urban-rural health insurance integration in China: Impact on health care utilization, financial risk protection, and health status. Appl. Econ. 54, 2491–2509 (2022). https://doi.org/10.1080/00036846.2021.1998323

8. Chai, Y., Liu, Y., Li, S., Zhang, L., Chen, D., Yin, W., Guo, L.: Can public hospital reform reduce medical resource mismatches? Evidence from China. Health Econ. Rev. 14, 88 (2024). https://doi.org/10.1186/s13561-024-00567-7

9. Fu, X., Ren, X., Chen, Q.: Can China's national comprehensive medical reform increase medical resources and reduce healthcare burden: Evidence from Chinese provinces. Front. Public Health 12, 1444840 (2024). https://doi.org/10.3389/fpubh.2024.1444840

10. Jiang, L., He, Y., Hu, C.: Public hospital reform, family health consumption and health inequality: Evidence from China family panel studies. Front. Public Health 12, 1352417 (2024). https://doi.org/10.3389/fpubh.2024.1352417

11. Yang, D., Acharya, Y., Liu, X.: Social health insurance consolidation and urban-rural inequality in utilization and financial risk protection in China. Soc. Sci. Med. 308, 115200 (2022). https://doi.org/10.1016/j.socscimed.2022.115200

12. Li, H., Li, Y.: The impact of digital economy development on public health: Evidence from Chinese cities. Front. Public Health 12, 1347572 (2024). https://doi.org/10.3389/fpubh.2024.1347572

13. Yang, J., Luo, B., Hu, L.: The impact of the digital divide on residents' healthcare consumption inequality: Evidence from CFPS in China. Front. Public Health 13, 1699498 (2026). https://doi.org/10.3389/fpubh.2025.1699498

14. Zhao, Q., Li, Y.: The impact of the digital economy on health inequality: Micro evidence from China. Int. J. Health Plan. Manag. 40, 961–974 (2025). https://doi.org/10.1002/hpm.3931

15. Zhang, C., Wu, S., Dong, Y., Jiang, M.: Government-led digital governance and the digital divide among cities: Implications for sustainable digital transformation in China. Sustainability 17, 10700 (2025). https://doi.org/10.3390/su172310700

16. Feng, X., Ding, D.: Did the healthy cities pilot policy improve household health poverty? An empirical study from China. Cities 167, 106381 (2025). https://doi.org/10.1016/j.cities.2025.106381

17. Guo, B., Feng, W., Cai, H., Lin, J.: Influence of public hospital reform on public health: Evidence from a quasi-natural experiment in China. Front. Public Health 11, 1104328 (2023). https://doi.org/10.3389/fpubh.2023.1104328

18. Wang, L., Zhao, Z., Xu, X., Wang, X., Wang, Y.: How does the construction of new generation of national AI innovative development pilot zones affect carbon emissions intensity? Empirical evidence from China. Sustainability 17, 6858 (2025). https://doi.org/10.3390/su17156858

19. Zhou, Z., Xiang, Y., Liu, C., Huang, X., Fu, Y.: Quantitative evaluation and optimization of AI policy and regulatory texts for smart healthcare. Front. Public Health 13, 1731013 (2025). https://doi.org/10.3389/fpubh.2025.1731013

20. Huang, Y., Liu, S., Gan, J., Liu, B., Wu, Y.: How does the construction of new generation of national AI innovative development pilot zones drive enterprise ESG development? Empirical evidence from China. Energy Econ. 140, 108011 (2024). https://doi.org/10.1016/j.eneco.2024.108011

21. Zhang, Z., Zhang, M., Xu, X.: Artificial intelligence, green innovation efficiency, and public health benefits: Evidence from China's pilot zones with a DID approach. Front. Public Health 13, 1713295 (2025). https://doi.org/10.3389/fpubh.2025.1713295

22. Acemoglu, D., Restrepo, P.: Automation and new tasks: How technology displaces and reinstates labor. J. Econ. Perspect. 33, 3–30 (2019). https://doi.org/10.1257/jep.33.2.3

23. Acemoglu, D., Restrepo, P.: Tasks, automation, and the rise in U.S. wage inequality. Econometrica 90, 1973–2016 (2022). https://doi.org/10.3982/ECTA19815

24. Fan, X., Wu, Y., Zhou, Y., Wu, S.: How does artificial intelligence shock affect labor income distribution? Evidence from China. Pac.-Basin Finance J. 90, 102691 (2025). https://doi.org/10.1016/j.pacfin.2025.102691

25. Feng, P., Peng, Z., Lu, Y., Bi, X., Liu, C.: Artificial intelligence and employment structure: Evidence from China. Appl. Econ. 1–17 (2026). https://doi.org/10.1080/00036846.2026.2642882

26. Yang, H., Wang, X., Zhang, T., Zhao, S., Cai, R., Wang, D.: Empowering chronic disease management with smart healthcare in China: A policy effectiveness evaluation by PMC index model. Front. Public Health 14, 1729450 (2026). https://doi.org/10.3389/fpubh.2026.1729450

27. Kickbusch, I., Piselli, D., Agrawal, A., Balicer, R., Banner, O., Adelhardt, M., Capobianco, E., Fabian, C., Singh Gill, A., Lupton, D., Medhora, R. P., Ndili, N., Ryś, A., Sambuli, N., Settle, D., Swaminathan, S., Morales, J. V., Wolpert, M., Wyckoff, A. W., ..., Wong, B. L. H.: The Lancet and Financial Times commission on governing health futures 2030: Growing up in a digital world. Lancet 398, 1727–1776 (2021). https://doi.org/10.1016/S0140-6736(21)01824-9

28. Peng, Z., Dan, T.: Digital dividend or digital divide? Digital economy and urban-rural income inequality in China. Telecommun. Policy 47, 102616 (2023). https://doi.org/10.1016/j.telpol.2023.102616

29. Cui, K., Li, B., Wang, H.: Quantitative analysis of health insurance reform in China: Pure consolidation or universal health insurance?. Econ. Model. 101, 105550 (2021). https://doi.org/10.1016/j.econmod.2021.105550

30. Li, Y., Li, L., Liu, J.: The efficient moral hazard effect of health insurance: Evidence from the consolidation of urban and rural resident health insurance in China. Soc. Sci. Med. 324, 115884 (2023). https://doi.org/10.1016/j.socscimed.2023.115884

31. Yang, F., Shu, H., Zhang, X.: Understanding “Internet Plus Healthcare” in China: Policy text analysis. J. Med. Internet Res. 23, e23779 (2021). https://doi.org/10.2196/23779

32. Zhang, Z., Tang, D., Wang, X., Feng, Q.: The impact of digital technology on health inequality: Evidence from China. BMC Health Serv. Res. 24, 1531 (2024). https://doi.org/10.1186/s12913-024-12022-8

33. Li, Z., Chen, Y., Ding, J.: Impact of health insurance equity on poverty vulnerability: Evidence from urban–rural health insurance integration in rural China. Front. Public Health 11, 1328265 (2023). https://doi.org/10.3389/fpubh.2023.1328265

34. Xu, M., Yang, W.: Who will drop out of voluntary social health insurance? Evidence from the new cooperative medical scheme in China. Health Policy Plan. 36, 1013–1022 (2021). https://doi.org/10.1093/heapol/czab017

35. Hao, Z., Bian, R., Wang, Y.: The impact of health insurance on informal care: Evidence from China's urban-rural basic health insurance consolidation. China Econ. Rev. 94, 102514 (2025). https://doi.org/10.1016/j.chieco.2025.102514

36. Liu, W., Huang, G.: Health insurance policy enforcement and catastrophic health expenditure: A case study in Sichuan province, China. Front. Public Health 13, 1596377 (2025). https://doi.org/10.3389/fpubh.2025.1596377

37. Liu, W., Kang, L., Yan, W., Yan, C., Wang, T., Zhang, N., Cui, V. Y., Liu, S., Tao, L., Liu, M., Liu, J., Liang, W.: Health insurance schemes and chronic disease management behaviours in China: Evidence from women

with non-communicable diseases. Risk Manag. Healthc. Policy 19, 1–19 (2026). https://doi.org/10.2147/rmhp.s600628

38. Liu, K., Liu, W., Frank, R., Lu, C.: Assessing the long-term effects of basic medical insurance on catastrophic health spending in China. Health Policy Plan. 37, 747–759 (2022). https://doi.org/10.1093/heapol/czac020

39. Shi, Z., He, P., Zhu, D., Lu, F., Meng, Q.: Changes in health care utilization and financial protection after integration of the rural and urban social health insurance schemes in Beijing, China. BMC Health Serv. Res. 22, 1226 (2022). https://doi.org/10.1186/s12913-022-08602-1

40. Grossman, M.: On the concept of health capital and the demand for health. J. Polit. Econ. 80, 223–255 (1972). https://doi.org/10.1086/259880

41. Nie, P., Li, Q., Cohen, A. A., Sousa-Poza, A.: In search of China's income-health gradient: A biomarker-based analysis. Appl. Econ. 53, 5599–5618 (2021). https://doi.org/10.1080/00036846.2021.1927962

42. Kim, G. H., Song, J. S., Nam, J. W., Lee, W. R., Yoo, K. B.: Trajectory of medical expenditure and regional disparities in hypertensive patients in South Korea. Front. Public Health 12, 1294045 (2024). https://doi.org/10.3389/fpubh.2024.1294045

43. Lin, L., He, P.: The effect of Internet use on depressive symptoms in older adults: Based on the chain mediating role of physical exercise and self-rated health. Front. Public Health 12, 1472445 (2024). https://doi.org/10.3389/fpubh.2024.1472445

44. Zhang, Y., Li, K., Pang, Y., Coyte, P. C.: The role of digital financial inclusion in China on urban–rural disparities in healthcare expenditures. Front. Public Health 12, 1397560 (2024). https://doi.org/10.3389/fpubh.2024.1397560

45. Jin, X., Yuan, Y., Chang, C., Wu, X., Tan, X., Liu, Z.: Telemedicine in China: Effective indicators of telemedicine platforms for promoting health and well-being among healthcare consumers. Digit. Health 11, 20552076251341163 (2025). https://doi.org/10.1177/20552076251341163

46. Chen, G., He, J.: From access to capability: Digital literacy and health inequality in contemporary China. Front. Public Health 14, 1843334 (2026). https://doi.org/10.3389/fpubh.2026.1843334

47. Cui, Y., He, Y., Xu, X., Zhou, L., Nutakor, J. A.: Can cultural capital, cognitive ability, and economic capacity help rural older adults bridge the digital divide? Evidence from an empirical study. Front. Public Health 12, 1393651 (2024). https://doi.org/10.3389/fpubh.2024.1393651

48. Callaway, B., Sant'Anna, P. H. C.: Difference-in-differences with multiple time periods. J. Econom. 225, 200–230 (2021). https://doi.org/10.1016/j.jeconom.2020.12.001

49. Sun, L., Abraham, S.: Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. J. Econom. 225, 175–199 (2021). https://doi.org/10.1016/j.jeconom.2020.09.006

50. Kahn-Lang, A., Lang, K.: The promise and pitfalls of differences-in-differences: Reflections on 16 and Pregnant and other applications. J. Bus. Econ. Stat. 38, 613–620 (2020). https://doi.org/10.1080/07350015.2018.1546591

51. Liu, X., Burge, G. S.: The effect of the new cooperative medical scheme on rural labor supply in China. China Econ. Rev. 85, 102159 (2024). https://doi.org/10.1016/j.chieco.2024.102159

52. Zhang, X.: Financial pension security and older adults' well-being: The moderating role of artificial intelligence pilot zones—Evidence from a natural experiment in China. Econ. Anal. Policy 93, 1333–1344 (2026). https://doi.org/10.1016/j.eap.2026.08.032

53. Jiang, Y., Shao, X.: Association between commercial health insurance and out-of-pocket health expenditure: Evidence from provincial panel data in China, 2011–2022. BMC Health Serv. Res. 26, 367 (2026). https://doi.org/10.1186/s12913-026-14079-z

54. Li, Z., Yuan, X., Zhang, Z., Chu, Y.: The heterogeneous relationship between public health expenditures and household medical expenditures: Evidence from relative poverty households in China. Front. Public Health 11, 1195223 (2023). https://doi.org/10.3389/fpubh.2023.1195223

55. Song, S., Li, C., Li, Y.: Catastrophic health expenditure and commercial health insurance purchase: Evidence from China. BMC Health Serv. Res. 26, 437 (2026). https://doi.org/10.1186/s12913-026-14181-2

56. Liu, Z., Shi, R., Jiang, L., Tang, F., Qiu, Z., Yao, L.: Impact of long-term care insurance on health out-of-pocket expenditure ratios for older adults. Front. Public Health 13, 1616761 (2025). https://doi.org/10.3389/fpubh.2025.1616761

57. Chen, X., Chen, Y., Qin, B., He, Q.: Enhancing health equity in China: The interplay of public health infrastructure, service utilization, and health insurance. Sustainability 17, 4785 (2025). https://doi.org/10.3390/su17114785

58. Guan, X., Xu, J., Huang, X.: Digital economy and the medical and health service supply in China. Front. Public Health 12, 1441513 (2024). https://doi.org/10.3389/fpubh.2024.1441513

59. Ma, C., Song, Z., Zong, Q.: Urban-rural inequality of opportunity in health care: Evidence from China. Int. J. Environ. Res. Public Health 18, 7792 (2021). https://doi.org/10.3390/ijerph18157792

60. Wang, X., Guo, Y., Qin, Y., Nicholas, S., Maitland, E., Liu, C.: Regional catastrophic health expenditure and health inequality in China. Front. Public Health 11, 1193945 (2023). https://doi.org/10.3389/fpubh.2023.1193945

61. Hart, J. T.: The inverse care law. Lancet 297, 405–412 (1971). https://doi.org/10.1016/S0140-6736(71)92410-X

62. Tian, L., Guo, Q., Shen, J., Liu, Y.: The cost of digital convenience: Food delivery platform entry, health, and medical spending. Soc. Sci. Med. 403, 119393 (2026). https://doi.org/10.1016/j.socscimed.2026.119393

63. Xu, J., Wang, C., Yin, X., Wang, W.: Digital economy and rural household resilience: Evidence from China. Agric. Econ. 70, 244–263 (2024). https://doi.org/10.17221/317/2023-agricecon

64. Zhai, C., Hu, D., Dong, J., Jin, M.: Does the digital economy amplify household income uncertainty? Evidence from China. Econ. Model. 158, 107518 (2026). https://doi.org/10.1016/j.econmod.2026.107518