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.

Human Capital, Residential Investment, and Productivity Dynamics: Evidence from Harmonized Vietnamese Provinces

Quang Binh Bui

binhbq@due.edu.vn

Da Nang University of Economics

Bui Phan Nha Khanh

Da Nang University of Economics

Research Article

Keywords: Labor productivity, Human capital, Residential investment, Local projections, Vietnam

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10639902/v1

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

Additional Declarations: No competing interests reported.

Abstract

This study examines how changes in human capital and household residential capital formation are associated with provincial labor-productivity dynamics in Vietnam. A balanced panel of 28 harmonized provincial units over 2010–2025 is reconstructed to the 2025 administrative boundaries. Human capital is measured by the trained-worker share, while residential investment is proxied by newly self-built floor area per capita. Panel local projections with province and year fixed effects trace cumulative productivity responses over five years. Inference combines province-clustered standard errors, Webb-weighted wild-cluster bootstrap p-values, Driscoll–Kraay covariance estimates, and a common-factor augmentation. Residential investment is followed by a persistently positive productivity path, although statistical precision varies across estimators. Changes in the trained-worker share show little short-run association but a delayed negative medium-run pattern whose significance is less stable. The results are interpreted as conditional dynamic associations, highlighting local absorptive capacity, skill matching, and spatial adjustment rather than causal treatment effects in provinces.

JEL classification: J24; R11; R21; C23; O47

1. Introduction

Productivity growth depends not only on the accumulation of labor and capital but also on whether local economies can combine skills, residential capacity, and places efficiently. Two margins are especially relevant in rapidly transforming economies. Human capital can raise productivity through individual capability and local knowledge spillovers, yet its payoff depends on skill quality, technology, and the allocation of workers across productive uses (Moretti, 2004; Hanushek and Woessmann, 2008; Gennaioli et al., 2013). Housing formation can shape where workers live, how readily they move toward expanding labor markets, and whether productive locations can accommodate growth (Diamond, 2016; Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019). These margins are usually studied in separate literatures even though both affect a region's capacity to absorb labor, skills, and investment.

The paper addresses a three-part research gap. First, the substantive literature offers limited evidence on the joint medium-run productivity dynamics associated with workforce skill formation and household residential capital formation at the subnational level. Second, much of the housing literature measures prices, rents, regulatory constraints, or supply elasticities rather than a physical flow of newly created residential space. This distinction is important in settings where self-built housing remains economically relevant and formal housing-market indicators are incomplete. Third, changes in administrative boundaries can make regional comparisons internally inconsistent unless historical observations are reconstructed on a common spatial basis. The paper therefore treats dynamic adjustment, measurement, and spatial comparability as interconnected empirical problems rather than as separate data exercises.

The human-capital literature provides a strong positive benchmark but also cautions against equating observed training with productive capability. Regional human capital is closely associated with

development, and more educated local workforces can generate productivity spillovers (Moretti, 2004; Gennaioli et al., 2013). Yet cognitive skills and education quality materially change measured human-capital stocks and their returns (Hanushek and Woessmann, 2008; Schoellman, 2012). When technologies, occupations, and skill demand adjust slowly, additional training need not translate immediately into aggregate productivity. Pritchett (2001), Acemoglu and Zilibotti (2001), Caselli and Coleman (2006), and Hsieh et al. (2019) provide complementary reasons why skill accumulation and productive use can diverge.

A parallel spatial literature emphasizes that residential capacity can affect the allocation of workers and economic activity. Housing constraints may restrict movement toward productive locations, while land-use restrictions can impede regional adjustment (Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019; Herkenhoff et al., 2018). At the same time, housing construction is inherently procyclical: productive local economies can generate higher incomes, stronger housing demand, and more building activity. Hornbeck and Moretti (2024) document how productivity gains propagate through wages and housing markets, while Greaney (2026) cautions that aggregate magnitudes attributed to housing constraints are sensitive to implementation choices. The empirical challenge is therefore to characterize persistent co-movement without assigning a causal direction that the data cannot identify.

Vietnam provides a useful setting in which these margins evolve simultaneously. Provinces differ markedly in workforce composition, population dynamics, urban development, digital readiness, and household construction activity. The 2025 administrative reorganization also creates a measurement challenge: historical observations must be placed on a consistent geographic basis before provincial comparisons over time are meaningful. This study reconstructs a balanced panel of 28 of the 34 post-reorganization provincial units for 2010–2025 by aggregating predecessor-province data to the 2025 boundaries. The harmonized geography is used as a common accounting frame, not as an exogenous policy experiment.

The study asks two related questions. How are changes in a training-based measure of human capital associated with the subsequent path of provincial labor productivity? And how is household residential capital formation, measured by newly self-built residential floor area per capita, associated with that path? A dynamic perspective is central because both channels may involve adjustment lags. Skills require time to be matched with suitable jobs, technologies, and complementary capital, while residential construction can accompany infrastructure, migration, local demand, and spatial reallocation. A contemporaneous regression alone may therefore conceal economically relevant medium-run relationships.

The paper makes four contributions. First, it brings human capital and residential capital formation into a common subnational productivity framework, allowing their adjustment paths to be evaluated on the same spatial and temporal basis. Second, it uses a physical housing-flow measure that captures household residential capital formation rather than prices or the existing housing stock. Third, it reconstructs historically consistent provincial units after the 2025 boundary changes, making the spatial

unit comparable through time while explicitly acknowledging the modifiable areal unit problem. Fourth, rather than presenting established estimators as a methodological contribution, the paper asks whether the documented response paths survive inference designed for a small number of clusters and for common provincial shocks, and whether they vary systematically with lagged local conditions. The baseline panel local projections are therefore complemented by wild-cluster bootstrap inference, Driscoll-Kraay covariance estimates, a common-factor augmentation, and pre-specified state-dependent extensions (Jordà, 2005; Cameron et al., 2008; Driscoll and Kraay, 1998; Pesaran, 2006; Chudik and Pesaran, 2015; Ramey and Zubairy, 2018).

The empirical evidence reveals two distinct patterns. Residential investment is followed by a positive productivity path across horizons, with coefficients that remain positive under the principal robustness exercises, although statistical precision varies across inference procedures. Changes in the trained-worker share display little short-run association but become negative at medium horizons; the magnitude around years three to four is comparatively stable, while statistical significance is less uniform. These results are interpreted as conditional dynamic associations rather than treatment effects. In particular, the human-capital pattern is not evidence that training is productivity-reducing; it instead motivates closer attention to skill quality, matching, reallocation, and local absorptive capacity.

The remainder of the paper proceeds as follows. Section 2 develops the literature-based motivation and empirical expectations. Section 3 describes the harmonized data, measurement choices, empirical estimand, and inference strategy. Section 4 reports the dynamic estimates and robustness analysis. Section 5 evaluates competing economic interpretations and the limits of the evidence. Section 6 concludes with policy implications and study limitations.

2. Literature Review and Conceptual Motivation

2.1. Human capital, skill quality, and regional productivity

Human capital is a central explanation for persistent productivity differences across countries and regions, but the literature distinguishes sharply between measured educational quantity and economically productive skills. Across a large sample of subnational regions, Gennaioli et al. (2013) show a strong association between human capital and regional development, while Moretti (2004) identifies productivity spillovers from more educated local workforces. These findings support a positive benchmark in which individual skills, knowledge diffusion, and complementarities raise local productivity. However, Hanushek and Woessmann (2008) emphasize that cognitive skills and education quality matter more than simple attainment measures, and Schoellman (2012) shows that accounting for education quality changes estimates of human-capital stocks substantially.

This distinction matters when the empirical measure is the share of employed workers classified as trained. A higher trained-worker share may capture broader productive capability only imperfectly, particularly if training quality, task relevance, experience, and employer demand vary across locations.

Pritchett (2001) shows that expansions in measured education can have weak aggregate returns when institutional quality or demand for skilled labor is limited. Acemoglu and Zilibotti (2001) and Caselli and Coleman (2006) emphasize the compatibility between technologies and available skill endowments, while Hsieh et al. (2019) show that the allocation of talent across occupations can have first-order productivity consequences. The broader misallocation literature similarly demonstrates that productive resources can generate very different aggregate outcomes depending on where they are deployed (Hsieh and Klenow, 2009).

These mechanisms imply different dynamic predictions. Knowledge spillovers and skill accumulation provide a positive benchmark; measurement error can attenuate the observed relationship; skill-technology mismatch can delay the return to training; and occupational or sectoral reallocation can generate non-monotonic adjustment. The literature therefore does not justify imposing an immediate or uniformly positive response at every horizon. A dynamic empirical design is better suited to distinguishing a short-run absence of association from a medium-run adjustment pattern.

2.2. Residential capital formation, spatial allocation, and productivity

A separate literature links housing conditions to local development through worker mobility, spatial sorting, and the geographic allocation of economic activity. Glaeser and Gyourko (2018) synthesize evidence that housing-supply conditions influence households' ability to relocate toward productive labor markets. Hsieh and Moretti (2019) formalize how housing constraints can restrict movement toward high-productivity cities and generate spatial misallocation, while Diamond (2016) shows that location choices differ systematically by skill and respond jointly to wages, amenities, and housing costs. Saiz (2010) further demonstrates that physical and regulatory constraints shape local housing-supply responsiveness, and Herkenhoff et al. (2018) link land-use restrictions to aggregate economic performance.

The variable used in this study is related to, but not equivalent to, the conventional objects in that literature. Newly self-built residential floor area per capita is neither a measure of housing affordability, the existing housing stock, nor the elasticity of market supply. It is better interpreted as a revealed physical flow of household residential capital formation and local accommodation capacity. In settings where self-built housing is economically important, such a flow may co-move with construction activity, infrastructure deepening, urban expansion, and the capacity to accommodate labor inflows. Agglomeration evidence also suggests that local investment and economic density can interact with productivity through spillovers and reallocation (Greenstone et al., 2010). This interpretation provides a conceptual bridge from a physical construction measure to the broader spatial-allocation literature without treating the measure as a direct proxy for housing-market regulation.

Reverse causality remains central. Productivity growth can raise wages, local demand, land values, and construction, so a positive relationship may reflect productive places building more housing rather than housing causing productivity. Hornbeck and Moretti (2024) document how local productivity gains

propagate through wages and housing markets. Greaney (2026) also shows that the aggregate magnitude of housing-constraint effects can be sensitive to implementation choices. Accordingly, residential construction is best treated as an endogenous local adjustment margin whose timing and persistence are empirically informative even when causal direction is unresolved.

2.3. Why an emerging-economy provincial setting matters

Most influential evidence on housing constraints, skill sorting, and local productivity comes from advanced economies with deep formal housing markets and relatively stable administrative geographies. Directly transferring those magnitudes to an emerging economy would be inappropriate. Vietnam differs in the importance of self-built housing, the pace of structural transformation, the role of provincial institutions, and the extent to which labor, infrastructure, and urban systems adjust simultaneously. These features strengthen the case for examining physical residential formation and training-based human capital as local adjustment margins while avoiding structural interpretations imported from other settings.

The provincial setting also raises a measurement issue that is rarely central in the existing literature. Administrative consolidation can mechanically alter regional totals, ratios, and per-capita measures if historical observations are combined across incompatible boundaries. A harmonized spatial frame is therefore a prerequisite for interpreting longitudinal provincial dynamics. This study treats boundary harmonization as a measurement contribution rather than as a source of identification: the 2025 map defines the comparable unit, while historical economic variation remains observational.

2.4. Empirical expectations

The literature implies directional benchmarks but not deterministic hypotheses. Productive skills should, in principle, support productivity, yet the observed response to a higher trained-worker share may be weak, delayed, or non-monotonic when training quality, technology, occupational demand, and talent allocation adjust at different speeds. Residential capital formation may facilitate local absorption and spatial adjustment, suggesting a positive conditional association, but simultaneity and procyclicality make a causal sign inappropriate to impose ex ante. The empirical objective is therefore to estimate the timing, persistence, and robustness of both response paths rather than to search for a predetermined sign at a single horizon.

The same mechanisms also leave room for state dependence. Prior residential capital formation may proxy local accommodation conditions under which additional skills are absorbed, while prior human capital may proxy capabilities that accompany residential expansion. Because neither channel implies an unambiguous interaction sign, the heterogeneity analysis asks whether response paths differ across pre-specified local states rather than searching for a regime that maximizes statistical significance.

3. Data and Empirical Strategy

3.1. Data, boundary harmonization, and sample construction

The empirical dataset is a strongly balanced panel of 28 of Vietnam's 34 provincial-level administrative units observed annually from 2010 to 2025, yielding 448 province-year observations (N = 28, T = 16). Geographic units follow the post-reorganization structure established by Resolution No. 202/2025/QH15 of the National Assembly, adopted on 12 June 2025. Historical observations originally published for predecessor provinces are reconstructed to these 2025 boundaries for every year in the sample. The purpose is to maintain a time-consistent spatial unit rather than combine observations defined on different administrative maps.

Most socioeconomic series are reconstructed from annual publications of the General Statistics Office of Vietnam. For merged units, original totals are aggregated whenever possible, and ratios or per-capita measures are computed only after aggregation. This avoids unweighted averaging across predecessor provinces that differ substantially in economic scale. Six units are excluded because a complete balanced historical series could not be reconstructed. Based on the study's sample accounting, they jointly represent approximately 2.2–2.5% of national GDP, 7-7.5% of population, and 2-2.5% of enterprises. The relatively small GDP and enterprise shares reduce, but do not eliminate, concern about coverage; the non-trivial population share means that the balanced sample should not be interpreted as mechanically representative of every post-2025 provincial unit.

Boundary harmonization improves comparability but does not remove the modifiable areal unit problem. All estimates are conditional on the 2025 spatial partition, and the new boundaries are treated as an accounting framework rather than an exogenous shock. Historical series are therefore interpreted within this harmonized spatial frame, and the resulting estimates remain conditional on the aggregation choices used to construct the post-2025 provincial units.

3.2. Variable construction and measurement scope

The dependent variable, lnLP, is the natural logarithm of provincial labor productivity, calculated as harmonized GRDP at constant 2010 prices divided by harmonized employment. GRDP and employment are first summed across predecessor provinces before the ratio is constructed. Human capital (HC) is measured by the percentage of employed workers who have received training. For merged units, trained workers are aggregated in the numerator and total employment in the denominator. The coefficient is therefore interpreted per one-percentage-point change in the trained-worker share. This indicator captures an observable training-based component of human capital, but not education quality, cognitive skills, experience, task relevance, or the broader stock of productive capabilities (Hanushek and Woessmann, 2008; Gennaioli et al., 2013).

The second focal variable is newly self-built residential floor area per capita. GSO-reported floor area is summed across predecessor provinces and expressed as lnHouse_pc = ln(floor area) - ln(population), which is equivalent to the logarithm of residential floor area per capita. The measure is a flow indicator

of household residential capital formation. It is not a measure of the housing stock, affordability, quality, market supply elasticity, or welfare. Because residential development can both influence and respond to local economic conditions, mobility, and spatial allocation, the variable is treated as endogenous and its coefficient is not assigned a causal interpretation.

Controls include log population, ICT readiness, log urban population, and CPI. The ICT index is reconstructed from annual infrastructure, human-resource, and ICT-application components published by the Ministry of Information and Communications and aggregated using predecessor-province GRDP shares. This weighting scheme is transparent but not innocuous because GRDP also enters the numerator of labor productivity. ICT is therefore treated strictly as a conditioning variable rather than as a focal causal regressor, and its coefficient is not interpreted substantively. Urban population is retained as a scale measure, not an urbanization rate. CPI is reconstructed using predecessor-province shares in retail sales of goods and consumer-service revenues as weights.

3.3. Empirical estimand and identification scope

The paper does not treat a static level coefficient as the primary estimand. The central object is the cumulative conditional response of labor productivity following a change in each focal variable. This choice aligns the empirical design with the possibility that skill matching, migration, residential construction, and complementary investment unfold over several years. Province fixed effects remove time-invariant local characteristics, while year fixed effects absorb nationwide shocks common to all provinces. These controls improve comparability but do not resolve reverse causality, anticipatory behavior, measurement error, or omitted time-varying local shocks. The estimated coefficients are therefore interpreted as conditional dynamic associations, not identified treatment effects.

3.4. Dynamic local projections and inference

The dynamic specification follows Jordà's (2005) local-projection approach. For horizon k = 0,...,5, define the cumulative outcome as $Y_{i,t+k}^{(k)} \equiv \ln LP_{i,t+k} - \ln LP_{i,t-1}$. Each horizon-specific regression is estimated as:

$$Y_{i,t+k}^{(k)} = \beta_k \Delta HC_{i,t} + \theta_k \Delta \ln House_pc_{i,t} + \Gamma_k' \Delta X_{i,t-1} + \mu_i^{(k)} + \lambda_t^{(k)} + \varepsilon_{i,t}^{(k)}$$

where DeltaHC is the annual percentage-point change in the trained-worker share, Delta lnHouse_pc is the annual log change in residential construction per capita, DeltaX(i,t-1) contains one lag of the differenced state vector, and mu and lambda denote province and year fixed effects.

The horizon is capped at five years for two reasons. Economically, five years is long enough to capture medium-run adjustment in skill matching, residential construction, and local reallocation. Statistically, longer horizons would intensify endpoint attrition in a panel with only 16 annual observations per province. The baseline estimation sample declines from 392 observations at k = 0 to 252 at k = 5; a common-sample specification is therefore reported as a robustness check. At multi-year horizons,

cumulative outcomes overlap mechanically across adjacent observations, so province-clustered inference is used to allow arbitrary within-province dependence induced by this construction.

Because the panel contains only 28 provincial clusters, Webb-weighted wild-cluster bootstrap p-values with 4,999 replications are treated as the principal small-cluster inference check (Cameron et al., 2008). Pointwise confidence intervals and p-values are reported by horizon, but the manuscript avoids treating an isolated threshold crossing as decisive evidence. Interpretation instead emphasizes the sign, economic magnitude, persistence across neighboring horizons, and stability across pre-specified sensitivity checks. This is important because testing multiple horizons creates a natural multiple-inference concern even when each horizon-specific regression is correctly estimated.

3.5. Robustness strategy and common-factor sensitivity

Robustness exercises are chosen to address specific threats rather than to search across arbitrary specifications. The lag structure is expanded from one to two years; all differenced right-hand-side variables are winsorized at the 1st and 99th percentiles; 2020–2021 is excluded to assess pandemic sensitivity; and a common estimation sample is imposed across horizons. An additional extreme-outlier exclusion is treated only as an appendix sensitivity because it removes a non-trivial share of observations. The preferred baseline is never selected by deleting observations according to the future outcome.

Cross-sectional dependence is addressed in two complementary sensitivity analyses. Driscoll-Kraay standard errors provide covariance-robust inference in the presence of general cross-sectional correlation (Driscoll and Kraay, 1998). Because T = 16 is relatively short, these results are treated as complementary inference rather than as the preferred estimator. A CCE-inspired common-factor augmentation adds cross-sectional averages to absorb latent common components motivated by Pesaran (2006) and Chudik and Pesaran (2015). This exercise is used as a specification sensitivity, not as a claim that the full dynamic CCE asymptotic framework is satisfied. Neither procedure identifies geographic spillovers; a spatial-lag interpretation would require a separately justified spatial-weights matrix.

3.6. State-dependent and standardized-shock extensions

To probe heterogeneity without using outcome-selected thresholds, two state-dependent local-projection extensions classify each observation using lagged level variables and cutoffs fixed from the horizon-zero baseline estimation sample. For the human-capital shock, the high-residential state equals one when lagged lnHouse_pc exceeds 0.0178, the sample median. For the residential-investment shock, the high-human-capital state equals one when the lagged trained-worker share exceeds 17.825%. These cutoffs are held fixed across all horizons. Each extension includes the state indicator, the non-focal contemporaneous shock, the baseline lagged first-difference state vector, province fixed effects, and year fixed effects. For the human-capital extension, the horizon-specific equation can be written as:

Relative to the baseline specification, the human-capital extension replaces the single ΔHC slope with two interactions, ΔHC × (1 - S_House) and ΔHC × S_House, where S_House is the lagged high-residential-investment indicator. All other baseline controls and fixed effects are retained.

The residential-investment extension is defined analogously by allowing the slope on DeltalnHouse_pc to vary with the lagged human-capital state. At each horizon, the relevant null is equality of the low- and high-state response coefficients. Equality p-values are obtained from 4,999 Webb-weighted restricted-residual wild-cluster bootstrap replications at the province level. A whole-path omnibus test applies common cluster weights across horizons and compares the sum of squared horizon-specific bootstrap t statistics with its restricted-bootstrap distribution. These exercises are heterogeneity diagnostics rather than identified regime treatment effects (Ramey and Zubairy, 2018).

To make the economic magnitudes of the two focal shocks directly comparable, the baseline coefficients are also rescaled to one sample standard deviation of each annual change. The standard deviations are 1.640 percentage points for DeltaHC and 0.1396 log points for DeltalnHouse_pc. This is a pure rescaling of the baseline local projections and therefore does not change their t statistics or bootstrap p-values.

As an additional Online Appendix sensitivity, the binary state indicator is replaced by a fixed logistic transition function, $G(q) = 1/[1 + \exp\{-\gamma(q - c)\}]$. The transition center c is set to the midpoint of the 25th and 75th percentiles of the lagged state, and $\gamma = 2 \ln(3)/(Q75 - Q25)$, so that the transition weight equals 0.25 at Q25 and 0.75 at Q75. The transition parameters are calibrated from the state-variable distribution rather than selected from the productivity response, following the smooth-state logic used in nonlinear dynamic analysis (Auerbach and Gorodnichenko, 2012).

4. Empirical Results

The results reported in this section are based on the harmonized balanced panel of 28 provincial units over 2010–2025 (448 unit-year observations). The main text reports only the estimates needed to establish the empirical pattern; full diagnostics, horizon-by-horizon estimates, and sensitivity analyses are provided in the Online Appendix.

4.1. Descriptive evidence and preliminary diagnostics

Table 1 reports compact descriptive statistics for the balanced panel. Average log labor productivity is 4.219, while the mean share of employed workers with formal training is 19.53%. Residential construction per capita exhibits substantial cross-sectional and temporal dispersion, consistent with pronounced heterogeneity in household housing investment across provinces.

Table 1
Descriptive statistics

Variable	Definition	Mean	SD	Min	Max
lnLP	Log labor productivity	4.219	0.544	2.633	5.774
Human capital, h	Trained-worker share (%)	19.534	8.923	3.300	55.000
lnHouse_pc	Log residential floor area per capita	-0.022	0.483	-2.279	1.590
lnPopulation	Log total population	7.890	0.530	6.997	9.508
ICT	ICT readiness index	0.497	0.096	0.267	0.711
lnUrbanPop	Log urban population	6.642	0.771	5.243	9.289
CPI	Consumer price index	91.860	3.882	81.900	101.500

Notes: N = 448 for all level variables. Variables are harmonized to the post-2025 provincial boundaries described in Section 3.1. Human capital is measured in percentage points. lnHouse_pc is the logarithm of newly self-built residential floor area per capita. Full distributional statistics and statistics for first differences appear in Online Appendix Table OA.1; correlation matrices appear in Table OA.2.

The preliminary diagnostics support the empirical strategy. Pairwise correlations among the variables in first differences are generally modest; the largest absolute correlation is 0.368, between changes in residential investment per capita and population. The mean variance-inflation factor is 1.06, providing little indication of problematic multicollinearity. The Modified Wald test strongly rejects homoskedasticity (chi-squared(28) = 7,524.96, $p < 0.001$), whereas the Wooldridge test does not reject the null of no first-order serial correlation ($p = 0.809$). The Pesaran CD statistic provides marginal evidence of cross-sectional dependence (CD = 1.845, $p = 0.065$; average absolute residual correlation = 0.281). These diagnostics motivate province-clustered inference, wild-cluster bootstrap p-values, and the Driscoll-Kraay and common-correlated-effects checks reported below. CIPS tests strongly reject a unit root for every first-differenced variable at the 1% level (Online Appendix Tables OA.3-OA.4).

4.2. Baseline dynamic responses

The principal estimates are obtained from panel local projections. At horizon k, the dependent variable is the cumulative change in log labor productivity from t-1 to t + k. Each regression includes province and year fixed effects and one lag of the differenced state vector. Because the sample contains only 28 provincial clusters, inference in the main text gives priority to Webb-weighted wild-cluster bootstrap p-values.

Table 2
Baseline panel local-projection estimates

Horizon	Human-capital response (%)	Clustered SE	Wild p	Residential-investment response (%)	Clustered SE	Wild p	N
0	0.236	0.311	0.473	0.036	0.043	0.430	392
1	0.065	0.563	0.928	0.092	0.068	0.203	364
2	-0.162	0.677	0.836	0.106	0.076	0.217	336
3	-1.198	0.713	0.140	0.119	0.088	0.238	308
4	-1.385	0.755	0.081	0.184	0.102	0.129	280
5	-0.971	0.845	0.287	0.167	0.123	0.215	252

Notes: For human capital, the response is the cumulative percentage change in labor productivity associated with a one-percentage-point increase in the trained-worker share. For residential investment, the response corresponds to a 1% increase in residential floor area per capita. All specifications include province and year fixed effects plus one lag of the differenced state vector. Wild p-values use 4,999 Webb-weighted wild-cluster bootstrap replications. The full coefficient, confidence-interval, and standardized-shock results are reported in Online Appendix Table OA.5.

The response to an increase in the trained-worker share is close to zero over the first two years but turns negative at medium horizons. The cumulative response reaches its largest magnitude after four years: a one-percentage-point increase in the trained-worker share is associated with approximately 1.385% lower cumulative labor-productivity growth (wild-bootstrap p = 0.081). The response subsequently attenuates to -0.971% at year five. The timing therefore points to a delayed, rather than contemporaneous, relationship.

This pattern should not be interpreted as evidence that human capital itself reduces productivity. The training variable captures one observable component of workforce human capital and does not measure skill quality or the entire stock of productive capabilities. The delayed negative response is therefore better viewed as an empirical puzzle potentially consistent with adjustment frictions, skill-job mismatch, labor reallocation, migration, or limited local absorptive capacity. These mechanisms remain interpretations rather than identified causal channels.

Residential investment displays a different trajectory. Its estimated coefficient is positive at every horizon, increasing from 0.036% contemporaneously to 0.184% at the four-year horizon before moderating to 0.167% at year five. The baseline wild-bootstrap p-values do not reject zero at conventional thresholds, so the baseline specification is more informative about the persistence and direction of the response path than about its statistical precision.

Because the two shocks are measured in different units, the standardized scaling helps compare their economic magnitudes. At the four-year horizon, a one-standard-deviation increase in the trained-worker

share change is associated with a -2.271% cumulative productivity response, while a one-standard-deviation increase in residential investment per capita is associated with a + 2.568% response. The statistical inference is unchanged by this rescaling. The complete standardized paths are reported in Online Appendix Table OA.9.

4.3. Cross-sectional dependence and alternative inference

Because provincial economies are exposed to nationwide shocks and may co-move over time, the baseline estimates are re-evaluated using Driscoll-Kraay inference and a CCE-inspired common-factor augmentation. Driscoll-Kraay estimation leaves the point estimates unchanged but modifies their estimated precision. At the four-year horizon, the human-capital coefficient remains − 1.385% but is not significant at the 10% level ($p_{DK} = 0.109$). The negative medium-run pattern is therefore more robust in sign and magnitude than in statistical precision.

The evidence for residential investment is stronger under Driscoll-Kraay inference. Its coefficient is statistically significant from horizons one through five, with p-values of 0.035, 0.015, 0.045, 0.041, and 0.003, respectively. At the four-year horizon the coefficient remains 0.184%, but the Driscoll-Kraay p-value falls to 0.041. Thus, allowing the covariance matrix to accommodate general cross-sectional dependence reinforces the evidence of a persistent positive medium-run association between residential investment and labor productivity.

The CCE-augmented specification absorbs latent common components through cross-sectional averages. The estimated magnitudes remain extremely close to the baseline. At the four-year horizon, the human-capital response is -1.380% compared with − 1.385% in the baseline, while the residential-investment response is 0.195% compared with 0.184%. The corresponding wild-bootstrap p-values are 0.083 and 0.110. Moreover, after CCE augmentation the Pesaran CD test does not reject cross-sectional independence at horizons one through five. The stability of the coefficients reduces, but does not eliminate, concern that the principal medium-run response patterns are driven mainly by omitted nationwide common factors.

4.4. Robustness of the medium-run estimates

Table 3 summarizes the main robustness exercises at the four-year horizon, where the negative human-capital response reaches its largest baseline magnitude and the residential-investment response is close to its peak. The table deliberately reports a compact set of theoretically motivated checks; all horizons and additional sensitivities are deferred to Online Appendix Tables OA.6-OA.8.

Table 3
Robustness at the four-year horizon

Specification	Human-capital response (%)	p-value	Residential-investment response (%)	p-value
Baseline FE + Wild bootstrap	-1.385	0.081	0.184	0.129
Two lags + Wild bootstrap	-1.399	0.207	0.254	0.082
Winsorized RHS + Wild bootstrap	-1.544	0.127	0.173	0.113
Excluding 2020–2021 + Wild bootstrap	-1.427	0.098	0.102	0.280
Common horizon sample + Wild bootstrap	-1.245	0.165	0.114	0.261
Driscoll-Kraay	-1.385	0.109	0.184	0.041
CCE-FE + Wild bootstrap	-1.380	0.083	0.195	0.110

Notes: Responses are expressed in the same units as Table 2. p-values are Webb wild-cluster bootstrap p-values except for the Driscoll-Kraay row, which reports t-based DK p-values. The winsorized specification caps all first-difference right-hand-side variables at the 1st and 99th percentiles. The common-sample specification fixes the estimation sample across horizons. Full horizon-by-horizon estimates, extreme-RHS exclusion, and the future-level LP are reported in Online Appendix Table OA.6.

The human-capital estimates are economically stable. Across the principal specifications in Table 3, the four-year response lies roughly between − 1.25% and − 1.54%. Statistical precision is less stable: the relationship is marginally significant in the baseline, when the pandemic years are excluded, and under CCE augmentation, but not under the two-lag, winsorized, common-sample, or Driscoll-Kraay specifications. The appropriate conclusion is therefore a robust negative medium-run pattern with moderate statistical support, rather than a universally significant negative effect.

Residential-investment estimates are likewise stable in sign and remain positive across the main specifications. At the four-year horizon, the magnitude ranges from about 0.10% to 0.25%. Statistical precision varies with the inference procedure: the coefficient is significant at 5% under Driscoll-Kraay inference and at 10% in the two-lag wild-bootstrap specification, but does not reach conventional significance under the baseline wild-bootstrap or CCE inference. The evidence therefore supports a persistent positive productivity association while indicating that exact statistical precision is specification-dependent.

The main conclusions are not driven by the pandemic period. Removing 2020–2021 leaves the four-year human-capital response at -1.427% (wild-bootstrap p = 0.098). Winsorizing the full set of differenced right-hand-side variables changes the magnitudes only modestly. More aggressive exclusion of extreme

observations is deliberately relegated to the Online Appendix because it removes a non-trivial fraction of the regression sample and is better interpreted as a sensitivity analysis than as an alternative baseline.

4.5. State-dependent local projections

Table 4 examines whether the baseline dynamic responses differ across pre-specified lagged local states. The state cutoffs are not estimated from productivity outcomes and remain fixed across horizons, so the exercise is designed to characterize heterogeneity rather than to identify an optimal threshold.

Table 4
State-dependent panel local-projection estimates

Horizon	HC response: low housing (%)	HC response: high housing (%)	Wild p: equality	House response: low HC (%)	House response: high HC (%)	Wild p: equality	N
0	0.240	0.107	0.877	0.016	0.047	0.621	392
1	0.020	-0.123	0.902	-0.000	0.149	0.228	364
2	0.198	-0.709	0.292	0.024	0.158	0.348	336
3	-0.318	-2.059	0.060	0.022	0.179	0.221	308
4	-0.536	-2.257	0.090	0.114	0.221	0.517	280
5	-0.545	-1.868	0.349	0.198	0.132	0.737	252

Notes: HC responses are cumulative percentage productivity responses to a one-percentage-point increase in the trained-worker share. House responses are approximate percentage productivity responses to a 1% increase in residential investment per capita. Low/high housing is defined using lagged lnHouse_pc with a fixed median cutoff of 0.0178; low/high HC is defined using lagged HC with a fixed median cutoff of 17.825%. All models include the state indicator, the non-focal contemporaneous shock, one lag of the differenced state vector, province fixed effects, and year fixed effects. Equality p-values use 4,999 Webb-weighted restricted-residual wild-cluster bootstrap replications. Whole-path equality p-values are 0.214 for the HC responses across housing states and 0.460 for the residential-investment responses across HC states. Clustered standard errors and full details appear in Online Appendix Table OA.10.

For human capital, the low- and high-housing response paths are similar through the first two years, but a larger negative response appears in the high-housing state at medium horizons. At year three, the cumulative response is -0.318% in the low-housing state and − 2.059% in the high-housing state (wild-bootstrap equality p = 0.060). At year four, the corresponding estimates are − 0.536% and − 2.257% (p = 0.090). The whole-path equality test, however, does not reject a common response profile (p = 0.214). The evidence therefore suggests that the delayed negative average response is more concentrated in higher lagged residential-investment states around years three and four, but it does not establish a sharply distinct housing regime.

For residential investment, the point estimates are generally larger under high lagged human capital from years one through four. For example, the year-three response is 0.022% in the low-HC state and 0.179% in the high-HC state. None of the horizon-specific equality tests provides strong evidence of a difference, and the whole-path equality p-value is 0.460. The median-split results therefore do not show that prior human capital systematically conditions the residential-investment productivity path.

The smooth-transition sensitivity in Online Appendix Table OA.11 produces stronger local contrasts but is kept secondary. Human-capital response differences across the smooth housing state have wild-bootstrap p-values of 0.050, 0.045, and 0.046 at years two, three, and four, respectively, while the whole-path p-value is 0.065. For residential investment, the largest local contrast occurs at year three ($p = 0.034$), whereas the whole-path p-value is 0.156. Taken together, these results provide at most suggestive evidence of smooth state dependence and do not replace the average baseline responses as the paper's primary estimand.

4.6. Overall interpretation

Taken together, the estimates reveal two distinct medium-run patterns. First, higher residential construction per capita is consistently followed by a positive cumulative labor-productivity response. The sign and magnitude survive alternative lag structures, common-factor controls, outlier treatments, and covariance estimators, although statistical precision differs across methods. This pattern is consistent with residential capital formation co-moving with broader local investment, construction activity, infrastructure deepening, and demand conditions. Because local economic expansion may itself stimulate residential construction, however, the estimates should be interpreted as dynamic conditional associations rather than causal effects.

Second, increases in the trained-worker share are followed by a delayed negative productivity response concentrated around years three and four. The magnitude is stable across specifications, but statistical significance is only moderate. The finding therefore constitutes an empirical puzzle rather than evidence that human-capital accumulation is productivity-reducing. In this setting, it is more cautiously interpreted as evidence that an expansion of measured training does not automatically translate into higher provincial productivity when labor-market matching, sectoral demand for skills, migration, or local absorptive capacity are imperfect.

The heterogeneity exercises sharpen rather than overturn these average patterns. The median-split local projections place the more negative medium-run human-capital responses in high, not low, lagged housing states, while residential-investment responses do not differ reliably by prior human capital. The smooth-transition sensitivity strengthens some horizon-specific contrasts but yields only marginal whole-path evidence for the human-capital response. A simple interpretation in which residential capacity mechanically buffers the human-capital adjustment, or human capital uniformly amplifies the housing response, is therefore not supported by these data.

More broadly, the stability of the medium-run estimates after common-factor augmentation reduces concern that the central response patterns are simply artifacts of nationwide common shocks, but it does not eliminate that possibility. Nevertheless, the observational nature of the data, potential endogeneity of both focal variables, and harmonization of historical observations to 2025 administrative boundaries limit causal interpretation. The results should accordingly be read as evidence on conditional dynamic relationships across harmonized Vietnamese provincial economies, rather than as identified treatment effects.

5. Discussion

5.1. Residential investment: enabling capacity or a procyclical marker?

The most consistent feature of the residential-investment estimates is their sign rather than any single p-value. The coefficient is positive at every baseline horizon and remains positive across the principal robustness exercises. At the four-year horizon, the estimate is 0.184% in the baseline and 0.195% under the common-factor augmentation, while Driscoll-Kraay inference yields a p-value of 0.041. This stability is compatible with an enabling-capacity interpretation: expanding residential space may accompany the ability of productive locations to absorb workers, construction activity, infrastructure, and complementary local investment (Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019; Diamond, 2016; Herkenhoff et al., 2018).

An equally plausible interpretation is procyclicality. Local productivity growth can raise incomes, employment, land demand, and the incentive to build, making self-built floor area a marker of broader local expansion rather than a causal input into productivity. Hornbeck and Moretti (2024) show that productivity shocks propagate through wages and housing markets, and Greaney (2026) cautions against treating particular aggregate housing-constraint magnitudes as structurally settled. The present design cannot discriminate between 'housing enables adjustment' and 'productive places build more housing.' The contribution is therefore to document the timing and persistence of the conditional association, not to identify a housing-supply multiplier.

The state-dependent extensions provide no decisive evidence that the residential-investment path requires a high-human-capital environment. Although point estimates are generally larger in the high-HC state from years one through four, the median-split whole-path equality test has $p = 0.460$ and the smooth-transition counterpart has $p = 0.156$. The evidence therefore does not establish a stable complementarity between the two measured forms of local capacity.

5.2. Human capital: a delayed pattern with competing explanations

The trained-worker share displays a different dynamic profile. The short-run coefficients are close to zero, whereas the response turns negative around years three and four. At the four-year horizon, the baseline estimate is -1.385% and the common-factor-augmented estimate is -1.380%. The sign and economic magnitude remain similar under several alternative specifications, but statistical precision is not uniform: the wild-bootstrap baseline p-value is 0.081, the common-factor sensitivity p-value is 0.083, and the Driscoll-Kraay p-value is 0.109. The evidence is therefore better described as a stable negative medium-run pattern with moderate statistical support than as a negative human-capital effect.

Several mechanisms are consistent with this pattern, and the data do not identify among them. The trained-worker share measures neither skill quality nor task relevance, so measurement error can weaken the link with productive capability (Hanushek and Woessmann, 2008; Schoellman, 2012). Training supply may expand faster than employer demand or technology adoption, generating temporary skill-job mismatch (Pritchett, 2001; Acemoglu and Zilibotti, 2001; Caselli and Coleman, 2006). Reallocation across firms, occupations, and sectors can also delay the productive use of talent (Hsieh et al., 2019). Other, less structural explanations remain possible: provinces may expand training in response to weak productivity, skilled workers may migrate, employment composition may change the denominator of the training share, or harmonization and extreme observations may introduce noise. The results therefore challenge a mechanical quantity-based interpretation of human capital, not the proposition that productive skills are valuable.

The state-dependent results also narrow one potential explanation for the human-capital pattern. Under the median split, the delayed negative response is more pronounced in provinces with higher lagged residential investment at years three and four, although equality of the full response paths is not rejected ($p = 0.214$). The smooth-transition sensitivity yields stronger horizon-specific differences but only marginal whole-path evidence ($p = 0.065$). Thus, insufficient residential formation is not supported as a simple explanation for the delayed human-capital response. Higher residential formation may itself accompany expansion, migration, sectoral reallocation, or other local adjustments, but these channels are not identified here.

5.3. What the evidence does and does not establish

The common-factor sensitivity is useful because the medium-run point estimates remain close to the baseline after cross-sectional averages are added, and residual cross-sectional independence is not rejected at horizons one through five. This reduces concern that the documented patterns are driven solely by nationwide common shocks, but it does not eliminate omitted common factors or establish a mechanism. Likewise, changes in statistical significance across wild-bootstrap, Driscoll-Kraay, and sample-based checks caution against inference from a single threshold.

The heterogeneity specifications are intentionally disciplined: the binary cutoffs are sample medians fixed from the horizon-zero estimation sample, and the smooth-transition parameters are calibrated from the lagged-state distributions rather than optimized against productivity outcomes. Even with these safeguards, multiple horizon-specific comparisons and only 28 clusters mean that isolated equality p-

values should be interpreted as suggestive local evidence rather than as identification of economic regimes.

Three claims are therefore deliberately excluded. The estimates do not establish that additional self-built housing causes productivity growth; they do not establish that training reduces productivity; and they do not show that migration, mismatch, or any other proposed channel mediates the observed paths. What the evidence does establish more credibly is the shape, timing, and specification stability of two conditional dynamic relationships across harmonized provincial economies. Mechanism discrimination requires external variation or richer micro-level data.

6. Conclusion and Policy Implications

This study examines the medium-run relationships linking labor productivity, a training-based measure of human capital, and household residential capital formation across 28 harmonized Vietnamese provincial units over 2010–2025. Two empirical patterns emerge. Residential construction per capita is followed by a persistently positive productivity path whose sign and magnitude are stable across the main specifications, although statistical precision varies across estimators. Changes in the trained-worker share show little short-run association but a delayed negative pattern around years three and four; the magnitude is comparatively stable, but significance is less uniform. These findings describe conditional dynamic relationships rather than causal treatment effects.

The state-dependent extensions do not establish sharply separated regimes. The human-capital response is more negative in the high lagged housing state around years three and four, but equality of the complete median-split paths is not rejected; residential-investment responses likewise do not differ reliably across lagged human-capital states. The smooth-transition results are somewhat more suggestive for human capital, but the whole-path evidence remains marginal. These findings reinforce the interpretation of the main estimates as average conditional adjustment patterns rather than regime-specific causal effects.

The evidence-based policy message is therefore narrower than a construction or training target. For housing, the results do not justify maximizing floor area. They are consistent with attention to whether residential capacity, serviced land, infrastructure, and local labor-market expansion adjust together as productive locations grow. In policy terms, the findings caution against treating housing as separate from employment and spatial planning, while the observational design prevents ranking specific supply interventions.

For human capital, the policy message is not to train fewer workers, but to avoid treating the trained-worker share as a sufficient performance target. Expanding measured training is more likely to translate into productivity when skill quality, employer demand, technology, occupational structure, and worker mobility are aligned. The findings are therefore consistent with a policy emphasis on employer-training linkages, information on local skill demand, task-relevant curricula, and the productive deployment of

skills rather than quantity targets alone (Hanushek and Woessmann, 2008; Acemoglu and Zilibotti, 2001; Hsieh et al., 2019).

Several limitations define the boundary of these conclusions. First, neither focal variable is externally identified, so reverse causality and omitted time-varying provincial shocks remain possible. Second, the trained-worker share is an incomplete proxy for human capital. Third, self-built floor area is not total housing supply, affordability, or housing quality. Fourth, all historical series are conditional on the 2025 harmonized geography and thus remain subject to the modifiable areal unit problem. Fifth, N = 28 and T = 16 limit precision, especially at longer horizons and under alternative covariance estimators. Sixth, six post-2025 provincial units are outside the balanced sample. Seventh, the state-dependent estimates necessarily depend on pre-specified state definitions and have limited power for detecting whole-path differences. Future research should exploit policy discontinuities or other external sources of variation and combine provincial aggregates with migration, firm, occupational, or housing-market microdata to identify the mechanisms behind the dynamic patterns.

Declarations

Author Contribution

B.Q.B. contributed to conceptualization, methodology, formal analysis, and writing the original draft. B.P.N.K. contributed to data curation, formal analysis, and writing—review and editing. Both authors reviewed and approved the final manuscript.

Data Availability

The study uses publicly available official data from the General Statistics Office of Vietnam and the Ministry of Information and Communications. The authors harmonized the original provincial series to Vietnam's 2025 post-reorganization geography and constructed the variables following the procedures documented in the manuscript and Online Appendix. The processed dataset and replication information are available from the corresponding author upon reasonable request, subject to the terms and conditions of the original data providers.

References

1. Acemoglu D, Zilibotti F (2001) Productivity differences. Quart J Econ 116(2):563–606. https://doi.org/10.1162/00335530151144104
2. Auerbach AJ, Gorodnichenko Y (2012) Measuring the output responses to fiscal policy. Am Economic Journal: Economic Policy 4(2):1–27. https://doi.org/10.1257/pol.4.2.1
3. Cameron AC, Gelbach JB, Miller DL (2008) Bootstrap-based improvements for inference with clustered errors. Rev Econ Stat 90(3):414–427. https://doi.org/10.1162/rest.90.3.414

4. Caselli F, Coleman WJII (2006) The world technology frontier. Am Econ Rev 96(3):499–522. https://doi.org/10.1257/aer.96.3.499
5. Chudik A, Pesaran MH (2015) Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. J Econ 188(2):393–420. https://doi.org/10.1016/j.jeconom.2015.03.007
6. Diamond R (2016) The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. Am Econ Rev 106(3):479–524. https://doi.org/10.1257/aer.20131706
7. Driscoll JC, Kraay AC (1998) Consistent covariance matrix estimation with spatially dependent panel data. Rev Econ Stat 80(4):549–560. https://doi.org/10.1162/003465398557825
8. General Statistics Office of Vietnam (Various years). Statistical Yearbooks of Vietnam and annual provincial socio-economic statistical publications. Hanoi: Statistical Publishing House
9. Gennaioli N, La Porta R, Lopez-de-Silanes F, Shleifer A (2013) Human capital and regional development. Quart J Econ 128(1):105–164. https://doi.org/10.1093/qje/qjs050
10. Glaeser EL, Gyourko J (2018) The economic implications of housing supply. J Economic Perspect 32(1):3–30. https://doi.org/10.1257/jep.32.1.3
11. Greenstone M, Hornbeck R, Moretti E (2010) Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. J Polit Econ 118(3):536–598. https://doi.org/10.1086/653714
12. Greaney B (2026) Housing constraints and spatial misallocation: Comment. Am Economic Journal: Macroeconomics 18(2):409–428. https://doi.org/10.1257/mac.20230141
13. Hanushek EA, Woessmann L (2008) The role of cognitive skills in economic development. J Econ Lit 46(3):607–668. https://doi.org/10.1257/jel.46.3.607
14. Herkenhoff KF, Ohanian LE, Prescott EC (2018) Tarnishing the Golden and Empire States: Land-use restrictions and the U.S. economic slowdown. J Monet Econ 93:89–109. https://doi.org/10.1016/j.jmoneco.2017.11.001
15. Hornbeck R, Moretti E (2024) Estimating who benefits from productivity growth: Local and distant effects of city productivity growth on wages, rents, and inequality. Rev Econ Stat 106(3):587–607. https://doi.org/10.1162/rest_a_01208
16. Hsieh C-T, Hurst E, Jones CI, Klenow PJ (2019) The allocation of talent and U.S. economic growth. Econometrica 87(5):1439–1474. https://doi.org/10.3982/ECTA11427
17. Hsieh C-T, Klenow PJ (2009) Misallocation and manufacturing TFP in China and India. Quart J Econ 124(4):1403–1448. https://doi.org/10.1162/qjec.2009.124.4.1403
18. Hsieh C-T, Moretti E (2019) Housing constraints and spatial misallocation. Am Economic Journal: Macroeconomics 11(2):1–39. https://doi.org/10.1257/mac.20170388
19. Jordà Ò (2005) Estimation and inference of impulse responses by local projections. Am Econ Rev 95(1):161–182. https://doi.org/10.1257/0002828053828518

20. Ministry of Information and Communications of Vietnam (Various years). Provincial ICT index reports and component data. Hanoi
21. Moretti E (2004) Workers' education, spillovers, and productivity: Evidence from plant-level production functions. Am Econ Rev 94(3):656–690. https://doi.org/10.1257/0002828041464623
22. National Assembly of Vietnam (2025) Resolution No. 202/2025/QH15 on the reorganization of provincial-level administrative units, 12 June 2025
23. Pesaran MH (2006) Estimation and inference in large heterogeneous panels with a multifactor error structure. Econometrica 74(4):967–1012. https://doi.org/10.1111/j.1468-0262.2006.00692.x
24. Pritchett L (2001) Where has all the education gone? World Bank Economic Rev 15(3):367–391. https://doi.org/10.1093/wber/15.3.367
25. Ramey VA, Zubairy S (2018) Government spending multipliers in good times and in bad: Evidence from US historical data. J Polit Econ 126(2):850–901. https://doi.org/10.1086/696277
26. Saiz A (2010) The geographic determinants of housing supply. Quart J Econ 125(3):1253–1296. https://doi.org/10.1162/qjec.2010.125.3.1253
27. Schoellman T (2012) Education quality and development accounting. Rev Econ Stud 79(1):388–417. https://doi.org/10.1093/restud/rdr025

Figures

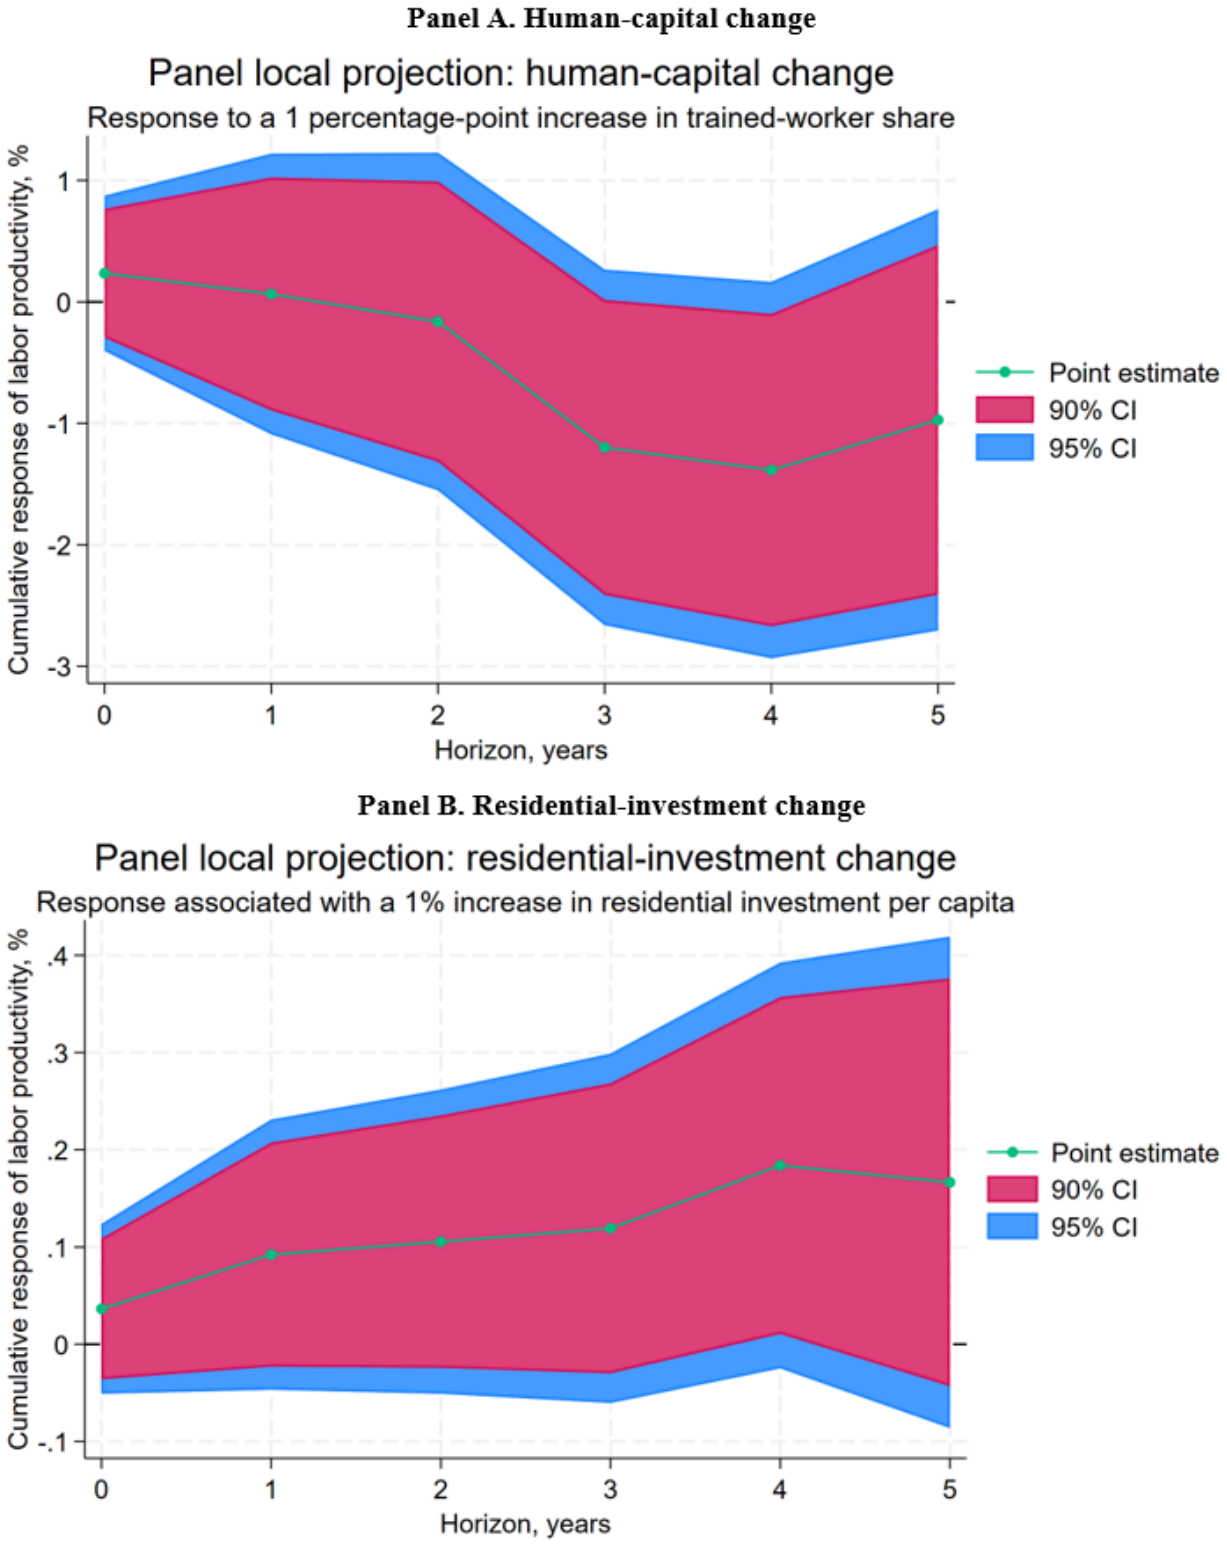


Figure 1

Dynamic responses of provincial labor productivity

Notes: Solid lines show point estimates; shaded regions show 90% and 95% confidence intervals based on province-clustered standard errors. Statistical significance discussed in the text is based primarily on

wild-cluster bootstrap p-values in Table 2. Numerical estimates are reported in Online Appendix Table OA.5.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- OnlineAppendix.docx