BACHELORARBEIT / BACHELOR'S THESIS

The Determinants of Economic Growth after the Global Financial Crisis in OECD Countries

verfasst von / submitted by

Nik Khosravipour

angestrebter akademischer Grad / intended degree

Bachelor of Science (WU) / BSc (WU)

Matrikelnummer / student ID number:	12208981
Studium / degree programme:	Business and Economics
Betreut von / Supervisor:	Univ.Prof.Dr. Jesús Crespo Cuaresma

Wien, Juni 2026 / Vienna, June 2026

Contents

List of Abbreviations

Abbreviation	Meaning
AR	Autoregressive (as in the AR(2) serial-correlation test)
CD	Cross-sectional dependence (Pesaran CD test)
COVID-19	Coronavirus disease 2019
CR2	Bias-reduced cluster-robust variance estimator (type CR2)
GDP	Gross domestic product
GMM	Generalized method of moments
IT	Information technology
OECD	Organisation for Economic Co-operation and Development
PWT	Penn World Table
TFP	Total factor productivity

Abstract

Long-horizon projections of national income assume that the empirical relationship between growth and its structural determinants is stable over time. This thesis tests that assumption against the strongest shock in the post-war record of advanced economies. Using a balanced annual panel of 24 pre-1995 OECD members from 1996 to 2019, built from the Penn World Table 11.0, a two-way fixed-effects growth regression on the augmented Solow determinants is combined with structural-break tests at known dates, 2008 and 2010, and at unknown dates through a sup-Wald scan. Mean growth fell from 2.5% per year before the crisis to 0.6% after it, in every country of the sample, yet no break test rejects the stability of the determinant-growth relationship at any candidate year. Conditional convergence and the investment association hold across the sample exclusions and the alternative labor measure. The crisis lowered the level of output without a measurable change in the structure of the growth process, and projections built on that structure require a lower starting point rather than a new model.

Introduction

Long-horizon projections of national income rest on an assumption that is rarely tested. Models that project GDP paths decades ahead, for applications such as the social cost of carbon and pension liabilities, estimate the empirical relationship between growth and its determinants from past data and carry that relationship forward unchanged (Müller et al. 2022). It is sound exactly as long as the relationship is stable. A structural break in how growth responds to its determinants would make the past record an unreliable guide to the future, and the prediction intervals built on it would fail in a way no widening could repair.

The global financial crisis of 2008 is the strongest candidate for such a break in the recent record of advanced economies. It produced the deepest contraction in the OECD since the Second World War, and the decade that followed it looked different from the decade before: growth across the advanced economies settled visibly lower and never returned to its pre-crisis pace. The crisis literature establishes that financial crises of this kind depress the level of output permanently (Cerra and Saxena 2007) and names the hysteresis mechanisms through which a contraction passes into potential output itself (Summers 2014). Whether the crisis also changed the relationship between growth and its determinants is a separate and open question. The slopes of that relationship are the quantities long-horizon projection models hold fixed. This thesis asks the question directly: did the empirical relationships between key structural determinants of growth and GDP per capita growth change after the global financial crisis of 2008?

Design and Approach

Determinants come from the augmented Solow framework: the economy's starting income position, the investment share of GDP, human capital, and the growth of the labor force (Solow 1956; Mankiw et al. 1992). These four variables are also the ones the robustness literature defends as the durable core of growth empirics, while fiscal, monetary, and institutional correlates prove fragile to specification choices (Levine and Renelt 1992; Sala-I-Martin 1997). Twenty-four economies that belonged to the OECD before 1995 make up the sample, observed annually from 1996 through 2019 in the Penn World Table, version 11.0. It closes before the COVID-19 contraction, so that a second large shock cannot confound the reading of the first.

A two-way fixed-effects panel forms the empirical design. Country effects absorb permanent differences across economies, and year effects absorb shocks common to the sample in each year, including the level effect of the crisis itself. The question then falls entirely on the slopes. A post-crisis indicator interacted with each determinant measures whether any coefficient changed after the break, a joint Wald test on the interactions delivers the panel form of a Chow test, and a sup-Wald scan over candidate break years removes the need to commit to a date in advance. Because identification comes from within-country variation over time, the estimates are short-run conditional associations inside each economy rather than statements about long-run cross-country convergence, and the thesis reads them as such throughout.

Findings

The level and the slopes give different answers, and the contrast between them is the central result. Mean GDP per capita growth across the sample fell from 2.5% per year before 2008 to 0.6% per year after it. Every one of the 24 countries recorded a lower average after the crisis, and annual growth became more volatile. That level shift is large and visible in every descriptive cut of the data. Slopes did not move with it. The joint test on the 2008 interactions fails to reject stability. Neither does any alternative: the test at 2010 is weaker still, and the data-driven scan finds no significant break at any candidate year between 2001 and 2015. The baseline relationships, conditional convergence and a positive investment association, hold with the expected signs across every sample exclusion and specification change. One suggestive movement, a strengthening of the investment slope after 2008, does not survive the joint test.

The crisis therefore moved the level of growth and left its empirical structure intact. Conditional projections built from the estimates make the practical content of that finding concrete: paths projected under the pre-crisis and post-crisis common growth components are nearly indistinguishable, and the lasting imprint of the crisis enters through the lower output level from which any projection now launches. For the long-horizon models that motivated the question, the relationship they hold fixed came through the strongest stress test the data can offer. The chapters that follow develop this result, beginning with the literature that makes stability worth testing.

Literature Review

The Determinants of Growth

Empirical growth research rests on a regressor set that has barely changed in six decades. Solow (1956) builds the foundation: under constant returns and diminishing marginal products, an economy converges to a steady state determined by its saving rate and its rate of labor force growth, and an economy below that steady state grows faster while it approaches it. The model assigns the saving rate a positive effect on steady-state income, assigns population growth a negative one, and leaves long-run per capita growth to exogenous technology. Mankiw et al. (1992) add human capital and take the model to data in its modern form. The augmented specification accounts for roughly 80% of the cross-country variation in income per capita over 1960 to 1985, the implied factor shares come out near one third each for physical and human capital, and the model predicts conditional convergence at about 2% per year. Their estimates also carry a warning for this design: in the 22-country OECD subsample the coefficients keep their signs but lose precision, a pattern that small advanced-economy samples repeat throughout the literature.

The cross-country growth regression that operationalized this framework begins with Barro (1991). Across 98 countries, growth is negatively related to initial income once human capital is held constant, with a coefficient on 1960 income of -0.0075, and poor countries catch up only conditional on adequate schooling. The conditioning problem this approach created became a literature of its own. Levine and Renelt (1992) apply extreme-bounds analysis to the accumulated findings and conclude that almost none survive variation in the conditioning set: the investment share and initial income are robust, and the fiscal, monetary, trade, and institutional correlates are fragile. Sala-I-Martin (1997) relaxes the criterion and recovers 22 robust variables out of 59 tested, with investment again the strongest result in the exercise. The two papers disagree on where to draw the line and agree on what sits above it. A parsimonious specification built on initial income, investment, human capital, and labor force growth uses the variables this robustness literature defends. This thesis estimates that specification.

The move from cross-sections to panels supplies the identification strategy. Islam (1995) shows that the cross-sectional regressions omit the country-specific technology level, that the omission biases the convergence coefficient because technology levels correlate strongly with the regressors, and that country fixed effects absorb the problem. Convergence estimates rise sharply once the fixed effects enter, reaching 9.3% per year in the OECD subsample, and the implied capital share falls to the plausible one third. Two further results from the panel literature shape the design used here. First, Islam (1995) finds that human capital loses significance once fixed effects absorb its cross-country variation, since schooling moves too slowly within countries to identify a coefficient from the remaining variation; the same mechanics operate in any within-country design, including this one. Second, Caselli et al. (1996) demonstrate how strongly the estimator drives the conclusion: ordinary least squares yields convergence near 0.6% per year, fixed effects yields 6.3%, and their preferred difference GMM yields 12.8%, while a Hausman test rejects the exogeneity of the Solow regressors outright. Smith (2024) synthesizes the econometric reasons behind this dispersion, from the bias identified by Nickell (1981) in short panels to slope heterogeneity across countries, and the practical lesson is that convergence estimates are statements about an estimator and a sample as much as about an economy. The estimates in this thesis are within-country conditional associations from a 24-year annual panel, and they are read as such throughout.

Stability as a Working Assumption

A second branch of the literature puts these relationships to work over horizons where their stability stops being a technicality. Müller et al. (2022) build a Bayesian factor model of growth in 113 countries over 118 years to generate joint predictive distributions of national incomes up to a century ahead, the horizon required for applications such as the social cost of carbon. The premise is stated in the paper's opening: the joint process followed by long-run growth over the past century is taken as a useful guide to the next one. The model treats the 2008 crisis as one observation inside a stable process, and nothing in the estimation tests whether that treatment is warranted.

The cost of being wrong about stability is measured by how little precision the stability assumption already leaves. Müller and Watson (2016) construct prediction sets for long-horizon average growth that remain valid across a wide class of persistence models, and even for the United States, with more than 150 years of data, the 90% set for 25-year-ahead average growth runs from −0.3% to 3.0%. Long-horizon inference is fragile under the best of assumptions. A break in the determinant-growth relationship would invalidate the intervals rather than merely widen them. The stability of those relationships through the largest shock in the post-war record is therefore worth testing.

A Relationship That Moves

The assumption would be harmless if the growth relationship were demonstrably fixed. The convergence literature of the past decade shows the opposite. Patel et al. (2021) document that unconditional convergence, absent from the data throughout the period that built the classic literature, emerged around 1995: for start dates before 1990 the convergence coefficient is zero or positive across every major dataset, and for start dates after 1995 poorer countries systematically grow faster, with the pattern holding in the Penn World Table, the World Development Indicators, and the Maddison data alike. A stylized fact that anchored 30 years of research reversed; the reversal was dated and robust, and no model had predicted it.

Kremer et al. (2021) decompose the reversal and locate the moving parts. The gap between conditional and unconditional convergence closed for two reasons: the correlates of growth themselves converged across countries, and the slopes of growth on the short-run policy and institutional correlates flattened. Solow regressors are the exception. Regressing the 2005 era slopes on their 1985 counterparts, the coefficients of growth on investment, population growth, and human capital line up with an R^2 of 0.95, while the slopes on the non-Solow correlates line up with an R^2 of 0.06. The Solow relationships are the stable core of an otherwise unstable empirical object. That finding sharpens the question this thesis asks rather than settling it, since the Kremer evidence compares decade-long cross-country eras, and whether the stable core also held within advanced economies, at annual frequency, through the crisis itself, is the question a within-country break test exists to answer.

The rest of the recent evidence reinforces the era-dependence and disputes its interpretation. Johnson and Papageorgiou (2020) survey the full convergence literature and conclude that developing countries as a group have made no progress on closing the income gap over 50 years, reading the recent episode as one-off level effects rather than a durable change in growth dynamics. Startz (2020) fits a Markov-switching model in which convergence is a regime rather than a fixed parameter, and finds the share of countries in the convergent state rising from roughly 9% in the early post-war decades to 76% by the 2010s, an estimate of the same transition Patel et al. (2021) date to the mid-1990s. Fiaschi and Johnson (2026) reach the era-dependence conclusion through distribution dynamics: the cross-country income distribution was bimodal and non-convergent from 1970 to 1995, unimodal and convergent from 2000 to 2010, and partially reverted thereafter, with the post-2010 shift concentrated in high-income countries, the income group this thesis samples. Mountford (2025) adds the longest perspective, estimating near-unit-root persistence in relative income positions and finding essentially no quartile mobility across 50 years. Whatever the papers conclude about the durability of the recent convergence, they share the premise that matters here: the empirical relationship between income levels and growth has shifted across eras, repeatedly and detectably. A framework whose cross-country slopes reversed once in 1995 has no automatic claim to within-country stability through 2008.

Crises and the Level of Output

Financial crises are the natural candidate for the kind of shock that could move the relationship. Reinhart and Rogoff (2008) place 2008 in an eight-century record across 66 countries in which surges of international capital mobility are systematically followed by waves of banking crises, and the credit expansion of the 2000s fits the historical pattern that preceded it. The aftermath of such crises is the part of the record that bears on this thesis. Cerra and Saxena (2007), estimating impulse responses across 190 countries, find that banking crises are followed by output losses averaging 7.5% of GDP with no subsequent return to the prior trend, rising toward 10% when a currency crisis coincides. Jordà et al. (2013) condition the severity on the boom that precedes it: across 14 advanced economies and more than 200 business cycles since 1870, recessions that follow financial crises run far below normal recession paths, and each standard deviation of excess credit accumulated in the preceding expansion lowers the five-year output path by a further 2pp. Summers (2014) supplies the mechanism that makes such losses permanent: depressed investment thins the capital stock, long unemployment spells erode labor supply, and the shortfall passes into potential output itself. This is the hysteresis channel.

This crisis literature describes movements in the level of output. Whether the crisis also moved the slopes, the strength with which investment, human capital, population growth, and the convergence term translate into annual growth, is a question the crisis literature does not address, and the productivity record complicates any reading that attributes the whole post-2008 slowdown to the crisis. Cette et al. (2015) document that total factor productivity growth decelerated before 2008 in every major advanced economy, with formal break tests placing the United States slowdown at the end of 2003 (Fernald, n.d.), and with Italy and Spain slowing from the late 1990s as euro entry cut real interest rates by 5 to 6pp and channeled cheap credit toward low-productivity firms. Part of what follows 2008 in the data was already underway before it.

The two bodies of evidence assemble the test this thesis runs. Convergence research shows that growth relationships shift across eras, so stability through 2008 is an open empirical question rather than a safe default. The crisis literature shows that 2008 delivered a shock large enough to move the level of output permanently, the natural candidate date for a break in the slopes. Forecasting models that hold these relationships fixed for decades ahead are where the stakes originate. None of the three provide a direct test of whether the within-country relationship between the Solow determinants and annual growth changed at the crisis in the economies where the crisis hit hardest. That test, on an annual panel of 24 OECD economies with formal break statistics at known and unknown dates, is the contribution of the chapters that follow.

Data and Variables

Data and Sample

Penn World Table version 11.0 (Feenstra et al. 2015) is the primary data source. As the standard reference for cross-country macroeconomic comparison, it expresses output, capital, and prices in common units across countries and over time. A single supplementary series, population growth, is cross-checked against the World Bank's World Development Indicators so that the demographic measure does not rest on a single source.

Twenty-four economies that belonged to the OECD before 1995, observed annually from 1996 through 2019, form the sample. The membership cut keeps the panel inside a group of advanced economies with broadly comparable institutions and data quality, and it excludes the transition economies that joined later, whose growth dynamics over this period were driven by convergence from a far lower base. It opens in 1996 because the dependent variable is a growth rate and the first year is lost to lagging. It closes in 2019 to hold the COVID-19 contraction outside the sample, since a second shock at the very end of the series would confound the question this thesis asks about the first one.

The resulting panel is balanced, with every country observed in all 24 years, for 576 country-year observations and no missing values in any of the constructed variables. Turkey and Mexico sit at the edge of the advanced-economy group and carry higher growth volatility than the Western European core; both are retained in the main sample and removed in a robustness check. Iceland is retained on the same terms despite its small size; Luxembourg, a micro-state, is excluded, the standard treatment in OECD samples.

Three countries contain extreme observations that reflect real distortions in the recorded data rather than coding errors. Ireland records GDP per capita growth of 23.4% in 2015, the maximum in the sample, together with implausibly high investment shares from 2015 onward; both trace to the relocation of multinational intellectual property into Irish national accounts, the episode known as *leprechaun economics*. Turkey records a contraction of 6.9% in 2001, driven by its banking and currency crisis. In Iceland, the collapse of the banking sector produced a contraction of 8.0% in 2009. All three observations are retained in the main sample, and separate robustness checks remove Ireland, Turkey together with Mexico, and Iceland in turn.

Variable Construction

The dependent variable is the annual growth rate of real GDP per capita. Real GDP per capita is constructed as real GDP at constant national prices divided by population, and its growth rate is the proportional change from the previous year. Real GDP at constant national prices is the appropriate PWT series for measuring growth within a country over time, since it holds prices fixed and isolates the change in real output.

The explanatory variables follow the augmented Solow framework. The Solow model, extended to include human capital by Mankiw et al. (1992), identifies the accumulation of physical and human capital and the growth of the labor force as the proximate determinants of output, conditional on the economy's distance from its own steady state (Solow 1956). Four regressors operationalize this framework. The convergence term is the log of lagged real GDP per capita, which captures the economy's starting position in each period. The investment share is the share of gross capital formation in GDP, taken directly from the PWT. Human capital enters as the log of the PWT human capital index, a measure built from average years of schooling and an assumed return to education. The labor input enters as population growth, the Solow rate n.

Table 2 reports the definitions and Table 3 the descriptive statistics over the estimation sample. Mean growth across the panel is 1.6% per year, with a standard deviation of 2.6%, and the range runs from the Greek contraction of 2011 at the low end to the Irish observation of 2015 at the high end. The investment share averages 25.6% of GDP, and its maximum of 63.7% is again the Irish distortion. Population growth is low and stable, averaging 0.7% per year.

Table 2: Variable definitions.

Variable	Definition
GDP per capita growth	Proportional change in real GDP per capita (real GDP at constant national prices divided by population). Dependent variable.
Log GDP per capita, lagged	Natural log of the previous year's real GDP per capita. Convergence term.
Investment share	Share of gross capital formation in GDP, from the PWT.
Log human capital	Natural log of the PWT human capital index.
Population growth	Proportional change in population. The Solow rate n.
Employment growth	Proportional change in persons engaged. Alternative labor input, used in robustness only.
Post-crisis indicator	Equals one from the break year onward. Set at 2008 in the baseline and at 2010 as an alternative.

Table 3: Descriptive statistics over the estimation sample (1996–2019, 24 countries).

Variable	N	Mean	SD	Min	Max
GDP p.c. growth	576	0.016	0.026	-0.097	0.234
Log GDP p.c. (lag)	576	10.725	0.380	9.456	11.587
Investment share	576	0.256	0.050	0.136	0.637
Log human capital	576	1.143	0.143	0.632	1.328
Population growth	576	0.007	0.006	-0.007	0.029
Employment growth	576	0.012	0.017	-0.082	0.066
Post-2008 dummy	576	0.500	0.500	0.000	1.000

Variable	N	Mean	SD	Min	Max
Post-2010 dummy	576	0.417	0.493	0.000	1.000

Table 4 reports the same variables split at 2008. Mean growth falls from 2.5% to 0.6% per year and its dispersion widens, while the investment share, human capital, and population growth stay close to flat across the two windows.

Table 4: Sample means before and after 2008, standard deviations in parentheses. Mean GDP per capita growth falls from 2.5% to 0.6% per year while its dispersion widens.

Variable	Pre-2008 (1996–2007)	Post-2008 (2008–2019)
GDP p.c. growth	0.025 (0.019)	0.006 (0.029)
Investment share	0.253 (0.048)	0.260 (0.051)
Log human capital	1.113 (0.152)	1.172 (0.128)
Population growth	0.007 (0.005)	0.007 (0.006)

The PWT also reports employment, and employment growth is the more common labor measure in applied work. Employment growth is procyclical and moves mechanically with output over the business cycle, which inflates its apparent association with annual GDP growth and conflates the labor channel with the cycle the analysis is trying to read. Population growth is the quantity the Solow model actually specifies, and it is close to exogenous at annual frequency. It is the baseline, and employment growth is reported only as a robustness check.

Empirical Methodology

The augmented Solow model was developed to explain long-run, cross-country differences in income, and the empirical literature that grew from it estimates convergence across countries from the between-country variation in the data (Barro 1991; Sala-I-Martin 1997; Levine and Renelt 1992). This design differs in kind: it estimates an annual panel with country fixed effects, so the identification comes from within-country variation over time. Accordingly, the object of interest is the short-run conditional comovement between the Solow determinants and growth inside each economy, and whether that comovement shifted after 2008. This panel approach to growth, with country effects absorbing fixed heterogeneity, follows Islam (1995). Regressors are borrowed from the growth literature for their economic content.

Empirical Specification

The baseline is a two-way fixed-effects regression of growth on the four determinants:

$$g_{it} = \beta_1 \ln y_{i,t-1} + \beta_2 s_{it} + \beta_3 \ln h_{it} + \beta_4 n_{it} + \alpha_i + \lambda_t + \varepsilon_{it}$$

Here g_{it} is GDP per capita growth in country i and year t, $\ln y_{i,t-1}$ is lagged log income, s_{it} is the investment share, $\ln h_{it}$ is log human capital, and n_{it} is population growth. Country effects α_i absorb all time-invariant heterogeneity across economies, including institutions, geography, and any permanent difference in average growth. Year effects λ_t absorb shocks common to the whole sample in a given year. The global financial crisis produced a sharp common fall in growth around 2009, and the year effects absorb that level shift in full. Year effects cannot absorb a change in the slopes, in how strongly growth responds to investment, human capital, convergence, and population. That change is the quantity of interest, and it is recovered through interactions rather than through a level dummy.

The structural-break specification adds interactions between each regressor and a post-crisis indicator:

$$g_{it} = \beta' x_{it} + \gamma'(D_t \cdot x_{it}) + \alpha_i + \lambda_t + \varepsilon_{it}$$

where x_{it} collects the four regressors and D_t equals one from the break year onward. Each element of γ measures the change in the corresponding coefficient after the break. A shift in that relationship after the crisis implies that γ differs from zero. Because the year effects already span the post-crisis indicator, D_t enters only through its interactions, and the level effect of the crisis remains in λ_t.

Estimation and Inference

Estimation proceeds by the within (two-way fixed-effects) estimator. Fixed effects are chosen over random effects, confirmed rather than assumed. A Hausman test rejects the random-effects specification decisively (chi-squared = 125.2, $p < 0.001$). Income and investment plainly correlate with the country effects, so the rejection is the expected result.

Inference requires more care than estimation. With only 24 clusters, cluster-robust standard errors are consistent as the number of clusters grows but biased downward and over-rejecting when clusters are few, so the conventional clustered standard error is unreliable at this sample size (Cameron and Miller 2015). The baseline inference therefore uses the CR2 cluster-robust estimator with Satterthwaite degrees of freedom, a small-sample correction designed for this setting (Bell and McCaffrey 2002; Pustejovsky and Tipton 2018). CR2 adjusts the residual-based variance for the leverage of each cluster and approximates the sampling distribution with a t reference whose degrees of freedom are estimated rather than fixed at the number of clusters minus one.

Advanced economies share business cycles, and a common global shock will correlate the errors across countries within a year. On a one-way specification with country effects alone, Pesaran's test rejects cross-sectional independence overwhelmingly (CD statistic = 34.9, $p < 0.001$) (Pesaran 2021). Year effects exist to capture this common component, and they do: on the two-way specification the same test finds no remaining dependence (CD = -0.42, p = 0.68). Driscoll-Kraay standard errors, which are heteroskedasticity- and autocorrelation-consistent and remain valid under cross-sectional dependence (Driscoll and Kraay 1998), are reported alongside the CR2 estimates as a guard against any dependence the year effects miss. A Wooldridge test further confirms serial correlation in the idiosyncratic errors ($F = 28.5$, $p < 0.001$) (Wooldridge 2010), which both the CR2 and the Driscoll-Kraay estimators accommodate. Both sets of standard errors agree on the substantive conclusions, so the inference does not hinge on a particular variance assumption.

Testing for a Structural Break

Break detection proceeds first at a known date and then with the date chosen by the data. The known-date test asks whether the relationship changed at a specific, economically motivated year. Two natural candidate dates present themselves: 2008, the onset of the financial crisis, and 2010, the onset of the euro area sovereign debt crisis that followed it. Baseline estimates set the break at 2008 and the alternative at 2010, and both are reported. At a known break date the interaction model above is a Chow test in panel form (Chow 1960), and the hypothesis that the relationship is stable is the joint hypothesis that all interaction coefficients in γ equal zero. This is tested with a CR2-based Wald test on the interaction block.

A simpler version of the known-date comparison estimates the model separately on the pre-crisis window, 1996 through 2007, and the post-crisis window, 2010 through 2019, with the crisis trough excluded, and compares the two coefficient vectors directly. That split-sample comparison is reported alongside the interaction tests. Splitting the sample re-estimates the country effects on windows of 10 to 12 years, which leaves little within-country variation behind them, and shortening the panel mechanically inflates the dynamic-panel bias on the convergence term. Under the discipline of common country effects and a single error process, the interaction model performs the same comparison, and the formal inference rests on it.

Unknown-date testing avoids pre-committing to a year. Choosing the break by hand risks reading the crisis into whichever year is most convenient, so the break date is instead determined by the data. Following Andrews (1993) and Bai and Perron (1998), the interaction model is re-estimated with the break placed at each candidate year across the central portion of the sample, trimming the first and last few years where a split would leave too few observations on one side. At each candidate date, the Wald statistic on the interaction block is recorded, and the maximum over all candidates, the sup-Wald statistic, identifies the most likely break. Because the distribution of the maximized statistic is nonstandard, the sup-Wald value is assessed against the critical values in Andrews (1993) rather than against the conventional Wald reference, and the per-year p-values from the scan are read as descriptive only.

Dynamic Panel Bias and Robustness

Lagged income enters as a near-dynamic regressor, and the within estimator of its coefficient is biased toward zero in a panel with a short time dimension, the bias identified by Nickell (1981). With 24 years the bias is moderate, but it falls on precisely the coefficient that captures conditional convergence, so it cannot be ignored. Difference GMM is the standard remedy: it removes the country effects by first differencing and instruments the transformed convergence term with deeper lags of income (Arellano and Bond 1991; Caselli et al. 1996). This thesis reports difference GMM as a robustness check, with the instrument set collapsed to keep the count of instruments small relative to the number of countries. On a panel this short the GMM estimator is weakly identified, a limitation taken up with the results, so the within estimator remains the workhorse and the GMM serves to bound the direction of the convergence bias.

Remaining robustness checks address the sample and the variable choices flagged above. Ireland, Turkey together with Mexico, and Iceland are each excluded in turn to confirm that no single distorted or volatile country drives the result. Employment growth replaces population growth as the labor input, since the baseline should not depend on the choice of labor measure. Moving the break date from 2008 to 2010 tests whether the structural-break conclusion hinges on the timing of the split.

Conditional Projections

As a final step, the estimates are turned into conditional projections of GDP per capita paths, the use to which long-horizon forecasting puts relationships of this kind (Müller et al. 2022). These projections take the full-sample coefficients, hold the investment share, human capital, and population growth at their 2015 to 2019 country means, and iterate the growth equation forward from the observed 2019 level, with the convergence term updating along the path. Within each era, the common component of growth, the country effect together with the era-average year effect, is recovered from the pre-crisis window, the post-crisis window, and the full sample in turn, so the projected paths differ only in the era from which the common component is drawn. The exercise is illustrative. Because the within model contains no trend-productivity term, the levels of the projected paths understate long-run growth, and no single path should be read as a forecast. The results are read across countries rather than along time: for each country the 2019 launch level of GDP per capita is set against the annualized growth rate its projection implies over the horizon, which renders the estimates as a convergence-style cross-section, with growth annualized over the whole horizon as that construction requires, and the era comparison enters as a shift in the common growth component.

Empirical Results

The Fall in Growth

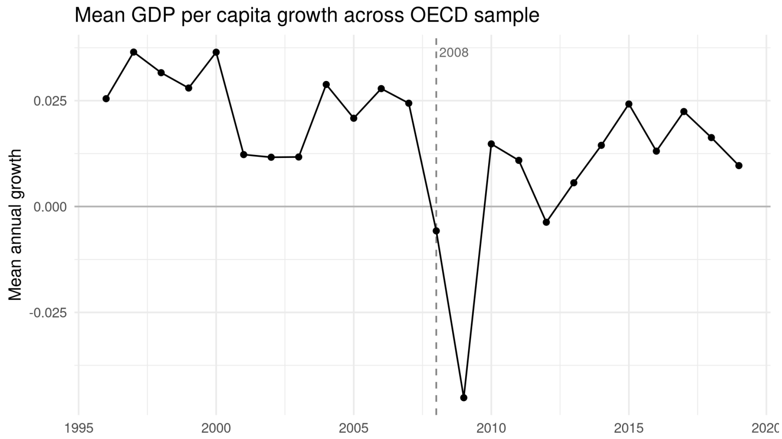

Figure 1: Mean annual GDP per capita growth across the sample, 1996–2019.

Average growth in the sample dropped sharply at the crisis. Mean GDP per capita growth ran at 2.5% per year across the pre-2008 window and fell to 0.6% per year from 2008 onward, a decline of almost 2pp in the typical annual growth rate. Volatility rose in tandem, as the standard deviation of annual growth widened from 1.9pp before the crisis to 2.9pp after it. Figure 1 plots the cross-country mean by year, and the drop around 2009 is the largest single movement in the series.

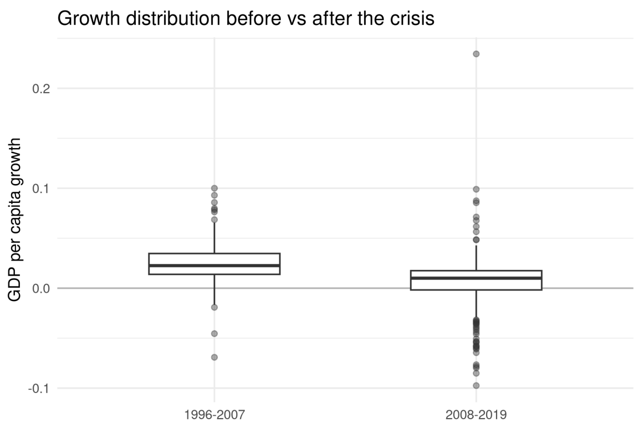

Figure 2: Growth distribution before and after the crisis.

This decline is broad-based. Every one of the 24 countries recorded a lower average growth rate after 2008 than in the years before it, so the aggregate fall reflects a common movement across the sample rather than a few large contractions pulling down the mean. Figure 2 shows the growth distribution before and after the crisis, with the post-2008 distribution shifted down and spread wider, and the per-country paths reported in the appendix confirm the same pattern economy by economy.

This level shift is the part of the crisis the two-way fixed-effects model removes by construction. Year effects absorb the common fall in 2009 and the slow recovery that followed, so the regressions below speak to a separate question: whether the strength of the association between the structural determinants and growth changed once the level shift is taken out.

Baseline Estimates

Table 5: Conditional growth regressions and the post-2008 structural break. Dependent variable: GDP per capita growth. Column (5) replaces the contemporaneous investment share with its one-year lag. CR2 cluster-robust standard errors (by country, Satterthwaite dof) in parentheses; country and year fixed effects included. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Baseline	(2) Break 2008	(3) Break 2010	(4) Emp.	(5) Inv. lag
Log GDP p.c. (lag)	-0.074***	-0.089***	-0.102***	-0.074***	-0.054**
	(0.021)	(0.026)	(0.029)	(0.015)	(0.020)
Investment share	0.163**	0.116*	0.129*	0.101**	
	(0.050)	(0.053)	(0.056)	(0.036)	
Investment share (lag)					0.057
					(0.055)
Log human capital	-0.033	-0.004	-0.066	0.018	0.003
	(0.111)	(0.095)	(0.072)	(0.094)	(0.115)
Population growth	-0.349	-0.854	-0.727		-0.024
	(0.316)	(0.489)	(0.508)		(0.574)
Employment growth				0.635***	
				(0.077)	
Log GDP p.c. (lag) × post-2008		-0.017			
		(0.012)			
Investment share × post-2008		0.108*			
		(0.055)			

	(1) Baseline	(2) Break 2008	(3) Break 2010	(4) Emp.	(5) Inv. lag
Log human capital × post-2008		0.042			
		(0.038)			
Population growth × post-2008		0.914			
		(0.668)			
Log GDP p.c. (lag) × post-2010			-0.012		
			(0.017)		
Investment share × post-2010			0.085		
			(0.070)		
Log human capital × post-2010			-0.004		
			(0.041)		
Population growth × post-2010			1.133		
			(0.774)		
Observations	576	576	576	576	552
Within R^2	0.091	0.120	0.130	0.275	0.022

In Table 5, the first column reports the baseline two-way fixed-effects regression. Two of the four determinants carry a clear within-country association with growth. Convergence enters negative and significant, with a coefficient of -0.074 (CR2 p = 0.009), so economies that enter a year further above their own recent income path grow more slowly over the year. That is the conditional convergence the augmented Solow model predicts. At annual frequency, the magnitude is small. In plain terms, an economy entering a year 10% above its recent income path grows about 0.7pp more slowly over that year. Investment share enters positive and significant, with a coefficient of 0.163 (CR2 p = 0.014); a 10pp increase in the investment share is associated with roughly 1.6pp of additional annual growth. Human capital and population growth carry no measurable association at annual frequency, with p-values of 0.77 and 0.30. Both variables move slowly and contribute little within-country variation over a 24-year window, so the nulls are expected.

Both inference methods agree. Driscoll-Kraay standard errors, robust to the cross-sectional dependence documented in the previous chapter, return the same two significant coefficients and sharpen them slightly, with the convergence term at p = 0.004 and the investment share at $p < 0.001$. At 0.091, the within R^2 is modest by construction, since the country and year effects already remove most of the variation in annual growth and the regressors explain only what remains.

Testing the Structural Break

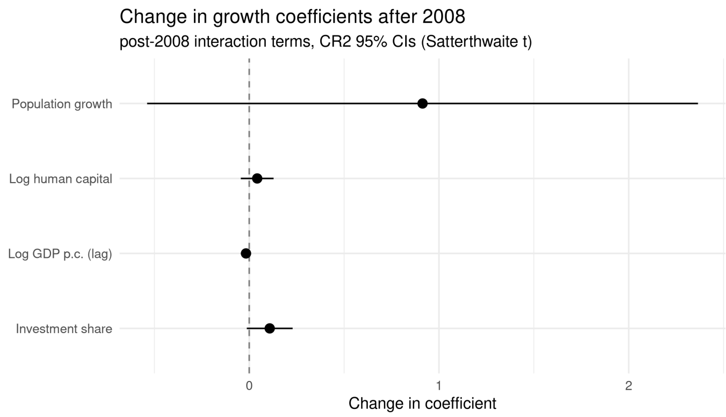

Figure 3: Change in the growth coefficients after 2008 (post-2008 interaction terms, CR2 95% intervals).

In the interaction model, the question falls directly on the post-2008 change in each coefficient. None of the four interaction terms is individually significant at conventional levels, as the second column of Table 5 shows. Investment interaction is the largest and the only one that approaches significance, at 0.108 with a p-value of 0.075, lifting the investment slope from 0.116 before the crisis to 0.224 after it. Convergence interaction is small and insignificant, at −0.017 with a p-value of 0.18. Figure 3 plots the four interaction estimates with their CR2 confidence intervals, and every interval crosses zero.

That investment result is the one hint of a post-crisis shift, and it does not hold up. A single interaction at a p-value of 0.075 is weak on its own, and it has to be read alongside the three interactions that show nothing.

Joint testing settles the matter. A hypothesis that all four interactions equal zero, the panel form of a Chow test on the crisis date, is not rejected ($F(4, 7.6) = 1.30$, $p = 0.351$). Across the 2008 break, the relationship between the structural determinants and growth is statistically stable. Moving the break to 2010 weakens the picture further: the third column of Table 5 shows every interaction insignificant, with the investment interaction falling to a p-value of 0.25, and the joint test again fails to reject stability ($F(4, 7.8) = 1.03$, $p = 0.45$).

The Split-Sample Comparison

Table 6: Split-sample estimation: baseline re-estimated separately on the pre-crisis (1996–2007) and post-crisis (2010–2019) windows, excluding 2008–09. CR2 cluster-robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

	Full sample	Pre: 1996–2007	Post: 2010–2019
Log GDP p.c. (lag)	-0.074***	-0.121***	-0.304***
	(0.021)	(0.030)	(0.066)
Investment share	0.163**	0.180**	0.312***
	(0.050)	(0.062)	(0.049)
Log human capital	-0.033	-0.021	0.471**
	(0.111)	(0.139)	(0.178)
Population growth	-0.349	-1.069*	0.359
	(0.316)	(0.571)	(0.614)
Observations	576	288	240
Within R^2	0.091	0.164	0.301

Table 6 reports the simpler comparison described in the previous chapter: the baseline re-estimated separately on the pre-crisis window from 1996 through 2007 and the post-crisis window from 2010 through 2019, with the crisis trough excluded. Point estimates move. Convergence deepens from -0.121 to -0.304, the investment share rises from 0.180 to 0.312, and log human capital turns from a null -0.021 to a significant 0.471.

Read at face value, the movement would suggest the instability the joint tests fail to find. Split-sample estimation is the weaker design, and the movement traces to what it gives up. Each window re-estimates the country effects on 10 to 12 observations per country, so the slopes rest on a fraction of the within variation that identifies the interaction model. Dynamic-panel bias of the within estimator grows as the panel shortens (Nickell 1981), and the deepening of the convergence coefficient from the 24-year panel to the 10-year window is of the order that bias implies under a stable relationship. Human capital moves almost linearly within each country over a single decade, so its within-window coefficient rests on thin differences in national schooling trends, and a sign flip across windows signals fragile identification rather than a new growth channel. In the interaction model, the same pre-post comparison holds the country effects common and uses the full panel, and its joint test carries the inference: no significant break.

Locating the Break

A known-date test could miss a shift that occurred away from 2008 or 2010, so the break is also left to the data. The sup-Wald scan re-estimates the interaction model at every candidate year from 2001 through 2015 and records the joint Wald statistic on the interactions at each. No candidate year produces a significant change. At 2001, the maximum statistic falls, with $F = 2.77$ and a naive p-value of 0.108, short of the 10% level before any correction. In Wald form the maximized statistic is 11.08 on four restrictions, below the Andrews (1993) asymptotic critical value for four parameters under the 15% trimming the scan applies. That critical value is 14.3 at the 10% level and 16.2 at the 5% level. A maximum-selected statistic lies to the right of the conventional Wald reference, so accounting for the search only widens that margin, and the maximum is not significant. Crisis years carry even smaller statistics, with 2008 at F = 1.30 and 2009 at F = 1.34. Table 9 reports the full per-year scan and Figure 6 plots it, and the data identifies no break in the relationship anywhere in the sample.

Robustness

Table 7: Robustness: baseline re-estimated dropping flagged countries. CR2 cluster-robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

	Baseline	No IRL	No TUR+MEX	No ISL
Log GDP p.c. (lag)	-0.074***	-0.069*	-0.099***	-0.078**
	(0.021)	(0.033)	(0.023)	(0.029)
Investment share	0.163**	0.141*	0.144**	0.186**
	(0.050)	(0.064)	(0.056)	(0.062)
Log human capital	-0.033	-0.026	-0.160*	-0.048
	(0.111)	(0.110)	(0.079)	(0.128)
Population growth	-0.349	-0.080	-0.395	-0.449
	(0.316)	(0.375)	(0.474)	(0.321)
Observations	576	552	528	552
Within R^2	0.091	0.089	0.106	0.087

Both significant baseline coefficients survive the sample exclusions. Table 7 re-estimates the baseline dropping the countries flagged for data-quality concerns in turn. Removing Ireland leaves the convergence term at -0.069 ($p = 0.085$) and the investment share at 0.141 ($p = 0.060$), both slightly softened but unchanged in sign and magnitude. Removing Turkey and Mexico together returns coefficients of -0.099 ($p = 0.008$) and 0.144 ($p = 0.043$), and removing Iceland returns -0.078 ($p = 0.030$) and 0.186 ($p = 0.014$). No single distorted or volatile country drives the baseline result.

Labor behaves as the earlier reasoning predicted. Replacing population growth with employment growth, in the fourth column of Table 5, lifts the within R^2 from 0.091 to 0.275 and gives employment growth a large, highly significant coefficient of 0.635 ($p < 0.001$). This rise reflects the mechanical procyclicality that kept employment growth out of the baseline, since employment and output move together over the cycle. Both the convergence term (-0.074) and the investment share (0.100) hold under this substitution, so the baseline associations do not depend on the choice of labor measure.

Lagging the investment share by one year separates it from contemporaneous output. In the fifth column of Table 5, the investment share is replaced with its one-year lag: the coefficient falls to 0.057 and loses significance ($p = 0.34$), while the convergence term holds at -0.054 ($p = 0.035$). Investment association is the only baseline coefficient sensitive to this timing, and the discussion takes up what that implies.

Table 8: Robustness: difference GMM (Arellano-Bond), instruments collapsed at lags 2–3. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

	Difference GMM
Log GDP p.c. (lag)	0.676
	(1.824)
Investment share	-0.013
	(0.374)
Log human capital	-0.494
	(0.979)
Population growth	-1.565
	(3.780)
Observations	528
Countries	24
Sargan p	0.93
AR(2) p	0.53

Reported in Table 8, the difference-GMM estimate confirms only that the convergence coefficient cannot be pinned down on a panel this short. Its specification passes on its own terms: the Sargan test does not reject instrument validity (p = 0.93) and the test for second-order serial correlation is clear (AR(2) p = 0.53). The coefficients carry very large standard errors all the same, and the convergence term comes back at 0.68 with a p-value of 0.71, wrong in sign and statistically indistinguishable from zero. GMM neither confirms nor overturns the within estimate of convergence, and the size of the Nickell-bias correction cannot be read from it. Within estimation remains the basis for the conclusions.

Across every specification and robustness check, one picture holds. Growth fell substantially after 2008 as a shift in its level, and the relationship between the structural determinants and growth did not measurably change.

Discussion

A Level Shift without a Slope Break

The 2008 crisis lowered the level of output across the 24-country sample without measurably altering the relationship between the Solow determinants and annual growth. Mean GDP per capita growth fell by almost 2pp between the windows on either side of the crisis, and every country recorded a lower average after 2008 than before it. Growth distribution shifted down and widened. The joint Chow test on the 2008 break fails to reject slope stability ($F = 1.30$, $p = 0.35$); at 2010 the case for a break is weaker still ($F = 1.03$, $p = 0.45$), and the sup-Wald scan finds none at any candidate year. Both the descriptive record and the regression evidence agree. What moved was the level of growth; the slopes stayed where they were.

The pattern matches what the crisis literature predicts when a banking shock causes a permanent output loss. Cerra and Saxena (2007), covering 190 countries over four decades, find that banking crises are followed by output losses averaging around 7.5% of GDP with no subsequent return to the prior trend, and that combined banking and currency crises produce losses closer to 10%. Across the OECD, 2008 was a banking crisis. Their estimates concern the level of output, and the record here implies a loss of the same character. Crisis-attributable loss and the gap against the pre-crisis trend are different magnitudes. The part the crisis itself imposed is the common level loss the year effects absorb across 2008 and 2009, about 7% of output, close to the Cerra banking-crisis average. A wider gap measures something different: a growth shortfall of almost 2pp per year, cumulated over the 12 post-crisis years, reaches one fifth of the output level by 2019 relative to a continuation of the pre-crisis pace, and it folds in a productivity slowdown that predates the crisis together with a recovery that never regained the earlier trend. Year effects in the regression absorb the common downward shift, and the slopes estimated net of it describe a determinant-growth relationship the crisis left intact.

Credit-expansion history shaped the depth of the loss. Jordà et al. (2013), working across 14 advanced economies and 223 business cycles since 1870, find that financial-crisis recessions run substantially deeper than normal recessions, and that excess credit built up during the preceding expansion compounds the damage. A one-standard-deviation increase in prior credit growth lowers the five-year output path by a further 2pp. Against that historical distribution, the OECD credit build-up of the 2000s placed the 2008 contraction toward the severe end. Part of the persistence in the post-2008 growth shortfall traces to the leverage built up before the shock.

Hysteresis is the mechanism through which the level loss becomes permanent (Summers 2014). A severe contraction reduces investment and lowers the capital stock, and prolonged unemployment erodes skills and withdraws workers from the labor force. Weaker innovation spending depresses the path of total factor productivity. Each channel operates on the level of potential output, a different object from the parameters that map the determinants into annual growth. This thesis estimates those parameters, and they are stable. Kremer et al. (2021) find the same separation across the full cross-country record: the slopes of growth on the Solow fundamentals of investment, population growth, and human capital carry an R^2 of 0.95 across the eras on either side of the crisis, while the slopes on non-Solow policy and institutional correlates carry an R^2 of 0.06. Its slopes form the stable core, and these results extend that finding to the within-country, annual dimension of the OECD panel.

Among the four interactions, only the investment coefficient moves enough to invite interpretation. It rises by 0.108 ($p = 0.075$), lifting the implied slope from 0.116 before the crisis to 0.224 after it. In a low-growth environment, years of elevated investment may translate more reliably into above-average output, tightening the within-country comovement of investment and growth. A third reading is harder to dismiss, since the contemporaneous link between investment and output examined below could itself have tightened after the crisis, raising the measured comovement with no change in the underlying slope. Sampling noise around a stable slope is an equally consistent reading. It does not survive the joint test, and it should not be carried past what the evidence supports.

A Pre-existing Component

Not all of the post-2008 growth shortfall originated in the crisis. Cette et al. (2015), using industry-level total factor productivity data for the United States and the four largest euro area economies, document that TFP growth slowed in each before 2008. In the United States, formal break tests place the slowdown at 2003:Q4, concentrated in IT-producing and IT-intensive industries where the returns to a decade of technology adoption had begun to fall (Fernald, n.d.). In Italy and Spain the deceleration runs earlier and through a different channel: the convergence of nominal interest rates after euro area entry cut real rates by 5 to 6pp between the mid-1990s and 2005, easing credit for low-productivity firms and raising capital misallocation. Aggregate TFP growth in both economies was near zero before the financial crisis arrived.

Italy and Spain belong to the 24-country sample, and so does Greece, which records the deepest contraction in the panel, 9.7% in 2011, during the sovereign debt crisis that followed the banking shock; the alternative break date at 2010 is set to capture that second phase. The lower average growth recorded across the southern euro area after 2008, and across the sample as a whole, combines two forces: the crisis-induced level fall documented by Cerra and Saxena (2007) and a productivity deceleration already underway. Year effects in the two-way design absorb the common component of that earlier trend, so the regression slopes are not contaminated by it. Both forces together are reflected in the intercept shift in the descriptive statistics, and the tests here are not built to separate them. Some of the lower mean growth after 2008 would therefore have appeared without the financial crisis, and the crisis accounts for only part of the gap.

Implications for Long-Horizon Projections

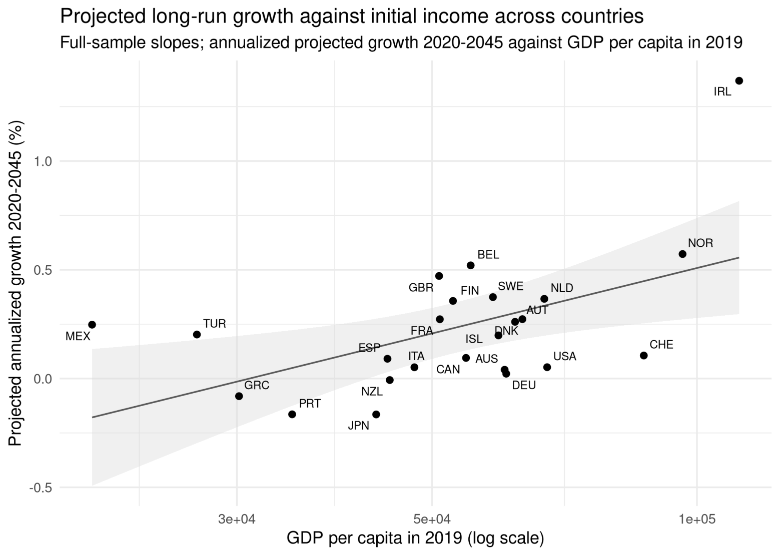

Figure 4: Projected long-run growth against initial income. For each country, GDP per capita in 2019 on the horizontal axis is set against the annualized growth rate the full-sample projection implies over 2020 to 2045 on the vertical axis. The line is the cross-country ordinary least squares fit.

Slope stability speaks to the use that motivated the question. Figure 4 presents the projection as a cross-section: each country appears once, its observed 2019 level of GDP per capita set against the annualized growth rate the full-sample estimates imply for it over the horizon. These points are a deterministic transform of the within estimates, not an estimated cross-country relation, so the figure reports the ordering the model implies. That ordering shows no unconditional convergence over the coming decades, since initial income does not pull projected growth down. A descriptive line through the points tilts faintly upward, by about 0.4pp of annualized growth per log point of income, yet the tilt halves and turns insignificant once Ireland is set aside, so the null carries the weight. Whatever upward tilt remains is in part mechanical, because the country growth component enters both the 2019 income level on the horizontal axis and the flow projected from it. The null stands against the within convergence slope the regression identifies, -0.074, an annual coefficient conditional on the country effect; the two are different objects and only the sign contrast between them is meaningful. The readings are the familiar conditional and unconditional forms of convergence (Islam 1995): the panel identifies movement toward a steady state specific to each country, and those steady states differ widely enough that the cross-section orders countries by income with no catch-up. Ireland sits at the upper right, its highest income paired with the fastest projected growth, the same national-accounts distortion the data chapter records as *leprechaun economics*.

Because the projection holds the slopes common to every country, the cross-country spread in Figure 4 comes from two sources, the country growth components the panel absorbs into its fixed effects and cross-country differences in income and the held regressors working through those common slopes. Together these terms carry the relative growth dynamics across the sample more than the determinant slopes do, and the figure shows how large they are beside the shallow cross-sectional gradient. Part of the spread is estimation noise, since the country effects come from a short panel, so the per-country positions are read as indicative rather than precise. Were the convergence parameter itself to differ across countries, the points would not array along a single line and the relative ranking could shift; the pooled specification cannot identify such slope heterogeneity, and a split between the crisis-hit periphery and the core is the natural form it would take, which the next chapter records as the leading extension. Varying the common component between the pre-crisis and the post-crisis window changes the sample-mean projected growth rate by almost nothing, yet individual country components move in offsetting directions, Greece falling the most and Turkey rising the most on windows too short to date these shifts with precision, so the stability of the average is not stability of every country's position.

The crisis enters the projection through the level at which each country begins. Year effects for 2008 and 2009 sum to a common loss of about 7% of output, close to the average banking-crisis loss in Cerra and Saxena (2007), and no later year effect runs above average to recover it, so every path now launches from a level permanently below the pre-crisis trend. For long-horizon projection of the kind Müller et al. (2022) build, the determinant slopes that such models hold fixed held through the largest macroeconomic shock in the sample period, and the imprint of the crisis is confined to the starting point from which those slopes carry growth forward.

Limitations

Statistical power is the binding limitation. The CR2 estimator with Satterthwaite degrees of freedom handles the few-clusters problem, but it does so by cutting the effective degrees of freedom for inference: the joint break tests run on roughly seven to eight degrees of freedom rather than the 23 the cluster count might suggest. That cost in power is concrete. The most precisely estimated of the four interactions, on the investment share, carries a standard error of 0.055, and at the degrees of freedom available a post-crisis change in that slope would have to reach 0.169 to be detected with 80% probability under a two-sided 5% test, larger than the baseline investment slope of 0.163 itself. No test of this size can resolve a change much smaller than the slope, and the three interactions estimated less precisely carry still less power; the observed investment shift of 0.108 sits below that threshold. The pooled test compounds the problem, since it estimates one average shift across the 24 economies, and offsetting breaks in different country groups, the euro area periphery against the core, could cancel in the mean. Failure to reject stability is consistent with a stable relationship, and it does not prove one. A larger cross-section of advanced economies, or a longer series, would sharpen the tests; this panel cannot rule out small slope shifts with confidence.

The null results for human capital and population growth reflect the design rather than evidence against the Solow model. Both variables move slowly within each country at annual frequency, and the country fixed effects absorb their levels, leaving too little within-country variation over 24 years to identify their coefficients with precision. P-values of 0.77 and 0.30 follow from that constraint. The Solow model assigns a central role to human capital and labor force growth, and the cross-country literature confirms both as robust growth correlates (Mankiw et al. 1992; Levine and Renelt 1992). At annual within-country frequency, the design cannot test those effects, and the nulls say nothing about their long-run importance.

The investment coefficient carries an endogeneity concern. Investment enters contemporaneously as a share, and it responds to current output as well as shaping it. Country fixed effects remove the time-invariant part of that feedback, but the year-to-year within-country correlation between investment and growth is not purged of reverse causation. The coefficient of 0.163 is a conditional within-country association, and reading it as causal requires the year-to-year variation in investment to be sufficiently exogenous to current output. Its sign is not obvious, since higher expected growth raises investment and inflates the estimate, while forward-looking investment that responds to lagged rather than current growth attenuates it. Across sample exclusions and the labor-variable substitution, the coefficient holds, but it does not survive lagging. With the investment share entered one year back, in the fifth column of Table 5, the coefficient falls to 0.057 and loses its significance. Part of the contemporaneous association therefore reflects within-year feedback rather than a slope that investment alone would trace, and the estimate is best read as a conditional comovement.

Difference GMM does not resolve the Nickell bias in the within-country convergence coefficient. It is valid on its own terms, with the Sargan test at $p = 0.93$ and the AR(2) test at $p = 0.53$, but its standard errors are large enough that no coefficient separates from zero, and the convergence term returns positive at 0.68 ($p = 0.71$). With 24 countries and a panel of 24 years, the instrument set carries too little variation to identify the GMM estimates with useful precision. Both the size and direction of the Nickell bias therefore remain open, and the within estimate of -0.074 stays the basis for the conclusions on convergence, with the caveat that the true coefficient may be somewhat more negative.

Conclusion

This thesis asked whether the empirical relationships between the structural determinants of growth and GDP per capita growth changed after the global financial crisis of 2008. Across a balanced annual panel of 24 OECD economies from 1996 to 2019, the answer is that they did not measurably change. Joint tests on the post-2008 interactions fail to reject stability, and the post-2010 alternative falls further short of rejection. A sup-Wald scan across every candidate break year from 2001 to 2015 finds no significant shift anywhere in the sample. Conditional convergence and the positive investment association that anchor the baseline hold with the expected signs through every sample exclusion, the alternative labor measure, and both break dates.

Slope stability stands against an unmistakable movement in the level. Mean growth fell from 2.5% per year before the crisis to 0.6% after it. Every country in the sample grew more slowly on average after 2008, and the dispersion of annual growth widened. That shift, then, lowered the path of output without altering the relationship that maps investment, human capital, population growth, and the convergence term into annual growth. Crisis literature records the same shape. Banking crises produce permanent output losses with no return to the prior trend, and the hysteresis mechanisms behind those losses operate on the level of potential output, leaving the parameters of the growth process untouched. Part of the post-2008 shortfall also predates the crisis, since productivity growth had been decelerating across the major advanced economies since the early 2000s, so the lower post-crisis level conflates the shock with a trend already underway.

These conclusions carry the limitations established in the discussion. With 24 countries, the break tests run on few effective degrees of freedom, and the failure to reject stability is consistent with a stable relationship without proving one; the design could detect a change in the best-estimated slope only once it approached the size of the slope itself, so a smaller shift could escape detection. Null results on human capital and population growth reflect the absence of within-country variation at annual frequency rather than evidence against their long-run role. Investment enters as a conditional association that contemporaneous reverse causation can inflate, and it loses significance once investment enters with a one-year lag; difference GMM is too weakly identified on a panel of this size to resolve the Nickell bias in the convergence term. These caveats bound the strength of the claim without reversing its direction.

For long-horizon projection, the result is usable as it stands. Models that carry growth relationships forward across decades assume the stability this thesis tested, and the assumption survived the strongest shock the post-war OECD record contains. Conditional projections built from the estimates show what survives and what does not: paths projected under pre-crisis and post-crisis common growth components are nearly indistinguishable, while the crisis years themselves removed about 7% of common output that no later year recovered. A forecaster re-basing projections after a financial crisis should lower the launch point and leave the propagation mechanism alone.

Several extensions follow naturally from the design's limits. Pooled testing estimates one average break across the 24 economies, so splitting the panel by group, the crisis-hit periphery against the core economies, would show whether offsetting shifts cancel in the aggregate and hide behind the null. A wider panel of advanced economies would raise the power of the break tests beyond what 24 clusters allow. A longer post-crisis window, once the COVID-19 shock can be cleanly separated, would test whether the stability found here extends through a second major disruption. Multi-country factor models in the spirit of the forecasting literature could test slope stability and level shifts jointly instead of conditioning one on the other. On the question it was built for, this design returns a clear answer: the determinants of growth did not measurably change at the crisis, and within the limits of this panel the post-2008 growth model is statistically indistinguishable from the one that preceded it.

References

Andrews, Donald W. K. 1993. "Tests for Parameter Instability and Structural Change With Unknown Change Point." *Econometrica* 61 (4): 821–56. https://doi.org/10.2307/2951764.

Arellano, Manuel, and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *The Review of Economic Studies* 58 (2): 277. https://doi.org/10.2307/2297968.

Bai, Jushan, and Pierre Perron. 1998. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica* 66 (1): 47–78. https://doi.org/10.2307/2998540.

Barro, Robert J. 1991. "Economic Growth in a Cross Section of Countries." *The Quarterly Journal of Economics* 106 (2): 407–43. https://doi.org/10.2307/2937943.

Bell, Robert M., and Daniel F. McCaffrey. 2002. "Bias Reduction in Standard Errors for Linear Regression with Multi-Stage Samples." *Survey Methodology* 28 (2): 169–81.

Cameron, A. Colin, and Douglas L. Miller. 2015. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources* 50 (2): 317–72. https://doi.org/10.3368/jhr.50.2.317.

Caselli, Francesco, Gerardo Esquivel, and Fernando Lefort. 1996. "Reopening the Convergence Debate: A New Look at Cross-Country Growth Empirics." *Journal of Economic Growth* 1 (3): 363–89. https://doi.org/10.1007/BF00141044.

Cerra, Valerie, and Sweta Saxena. 2007. "Growth Dynamics: The Myth of Economic Recovery." *SSRN Electronic Journal*, ahead of print, January 1. https://doi.org/10.2139/ssrn.777706.

Cette, Gilbert, John Fernald, and Benoît Mojon. 2015. "The Pre-Great Recession Slowdown in Productivity." *Federal Reserve Bank of San Francisco, Working Paper Series*, March 28, 01–54. https://doi.org/10.24148/wp2016-08.

Chow, Gregory C. 1960. "Tests of Equality Between Sets of Coefficients in Two Linear Regressions." *Econometrica* 28 (3): 591. https://doi.org/10.2307/1910133.

Driscoll, John C., and Aart C. Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *The Review of Economics and Statistics* 80 (4): 549–60. https://www.jstor.org/stable/2646837.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table." *American Economic Review* 105 (10): 3150–82. https://doi.org/10.1257/aer.20130954.

Fernald, John G. n.d. *Productivity and Potential Output Before, During, and After the Great Recession.*

Fiaschi, Davide, and Paul A. Johnson. 2026. "(In)stability in the Dynamics of the Cross-Country Distribution of Income Per Capita." *Empirical Economics* 70 (2): 24. https://doi.org/10.1007/s00181-025-02853-7.

Islam, N. 1995. "Growth Empirics: A Panel Data Approach." *The Quarterly Journal of Economics* 110 (4): 1127–70. https://doi.org/10.2307/2946651.

Johnson, Paul, and Chris Papageorgiou. 2020. "What Remains of Cross-Country Convergence?" *Journal of Economic Literature* 58 (1): 129–75. https://doi.org/10.1257/jel.20181207.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (s2): 3–28. https://ideas.repec.org//a/wly/jmoncb/v45y2013is2p3-28.html.

Kremer, Michael, Jack Willis, and Yang You. 2021. "Converging to Convergence." Working Paper 29484. Pre-published November. https://doi.org/10.3386/w29484.

Levine, Ross, and David Renelt. 1992. "A Sensitivity Analysis of Cross-Country Growth Regressions." *The American Economic Review* 82 (4): 942–63. https://www.jstor.org/stable/2117352.

Mankiw, N. Gregory, David Romer, and David N. Weil. 1992. "A Contribution to the Empirics of Economic Growth*." *The Quarterly Journal of Economics* 107 (2): 407–37. https://doi.org/10.2307/2118477.

Mountford, Andrew. 2025. "Economic Growth Analysis When Balanced Growth Paths May Be Time Varying." *Oxford Bulletin of Economics and Statistics* 87 (3): 477–98. https://doi.org/10.1111/obes.12638.

Müller, Ulrich K., James H. Stock, and Mark W. Watson. 2022. "An Econometric Model of International Growth Dynamics for Long-Horizon Forecasting." *The Review of Economics and Statistics* 104 (5): 857–76. https://doi.org/10.1162/rest_a_00997.

Müller, Ulrich K., and Mark W. Watson. 2016. "Measuring Uncertainty about Long-Run Predictions." *The Review of Economic Studies* 83 (4): 1711–40. https://doi.org/10.1093/restud/rdw003.

Nickell, Stephen. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49 (6): 1417–26. https://doi.org/10.2307/1911408.

Patel, Dev, Justin Sandefur, and Arvind Subramanian. 2021. "The New Era of Unconditional Convergence." *Journal of Development Economics* 152 (September): 102687. https://doi.org/10.1016/j.jdeveco.2021.102687.

Pesaran, M. Hashem. 2021. "General Diagnostic Tests for Cross-Sectional Dependence in Panels." *Empirical Economics* 60 (1): 13–50. https://doi.org/10.1007/s00181-020-01875-7.

Pustejovsky, James E., and Elizabeth Tipton. 2018. "Small-Sample Methods for Cluster-Robust Variance Estimation and Hypothesis Testing in Fixed Effects Models." *Journal of Business & Economic Statistics* 36 (4): 672–83. https://doi.org/10.1080/07350015.2016.1247004.

Reinhart, Carmen, and Kenneth Rogoff. 2008. *This Time Is Different: A Panoramic View of Eight Centuries of Financial Crises.* In *Annals of Economics and Finance*, vol. 15, 15. https://doi.org/10.3386/w13882.

Sala-I-Martin, Xavier X. 1997. "I Just Ran Two Million Regressions." *The American Economic Review* 87 (2): 178–83. https://www.jstor.org/stable/2950909.

Smith, Ron. 2024. "Econometric Aspects of Convergence: A Survey." *Open Economies Review* 35 (February): 701–21. https://doi.org/10.1007/s11079-024-09753-w.

Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *The Quarterly Journal of Economics* 70 (1): 65. https://doi.org/10.2307/1884513.

Startz, Richard. 2020. "The Next Hundred Years of Growth and Convergence." *Journal of Applied Econometrics* 35 (1): 99–113. https://ideas.repec.org//a/wly/japmet/v35y2020i1p99-113.html.

Summers, Lawrence H. 2014. "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound." *Business Economics* 49 (2): 65–73. https://doi.org/10.1057/be.2014.13.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data.* 2nd ed. MIT Press.

Appendices

Appendix A: Growth Paths by Country

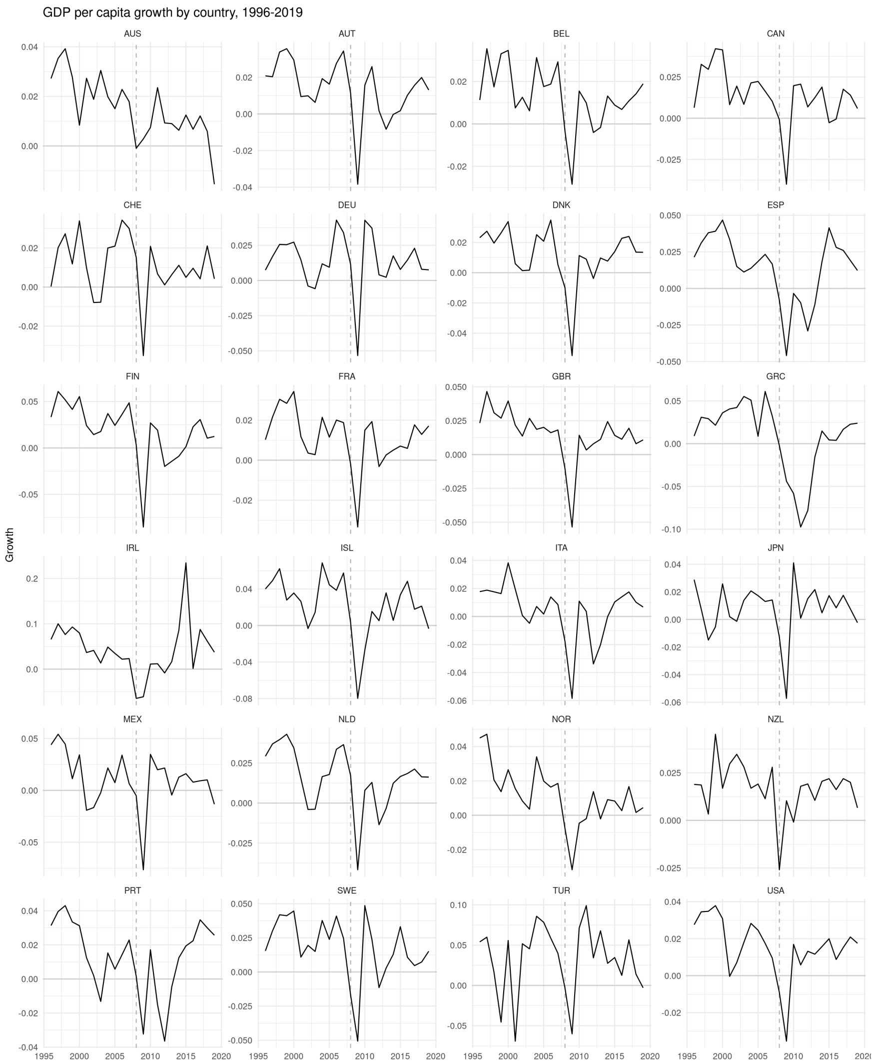

Figure 5: GDP per capita growth by country, 1996–2019.

Figure 5 reports the annual GDP per capita growth path of each sample economy over the full window, with the 2008 crisis year marked. The common contraction of 2009 is visible in every panel, and each economy records a lower average growth rate after 2008 than before it, the pattern summarized in the Results chapter.

Appendix B: Unknown-Date Sup-Wald Scan

Table 9: Unknown-date sup-Wald scan. Joint CR2 Wald statistic on the four break interactions at each candidate year, 2001–2015. The maximum falls at 2001 ($F = 2.77$, naive $p = 0.108$), short of the 10% level before any correction for the maximum selection.

Break year	Wald F	p (naive)
2001	2.77	0.108
2002	1.32	0.346
2003	1.18	0.392
2004	0.59	0.679
2005	0.57	0.691
2006	0.65	0.646
2007	1.36	0.332
2008	1.30	0.351
2009	1.34	0.336
2010	1.03	0.450
2011	0.96	0.478
2012	0.92	0.499
2013	1.46	0.301
2014	1.57	0.273
2015	2.16	0.166

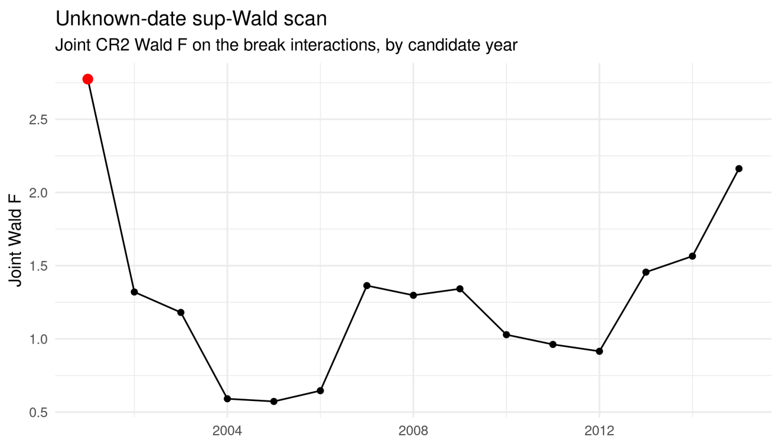

Figure 6: Joint CR2 Wald statistic on the break interactions at each candidate year, 2001–2019.

Table 9 and Figure 6 report the unknown-date break search underlying the Results chapter. The joint Wald statistic on the four interactions is recorded at each candidate break year, and the maximum falls at 2001 at $F = 2.77$ (naive $p = 0.108$). No candidate year reaches significance, and the crisis years themselves carry small statistics, so the data locates no structural break anywhere in the sample.