Oxford Economic Papers, 2026, 1–29
https://doi.org/10.1093/oep/gpag006
Article

OXFORD

Does productivity growth boost savings? Causal evidence from transition episodes

Abhishek Kumar[1,2], **Sushanta Mallick**[3,*], **Kunal Sen**[4,5,]

[1]University of Southampton, Southampton, SO17 1BJ, United Kingdom
[2]Centre for Social and Economic Progress, New Delhi, India
[3]School of Business and Management, Queen Mary University of London, London, E1 4NS, United Kingdom
[4]UNU-WIDER, Finland
[5]Global Development Institute, The University of Manchester, Manchester, M13 9PL, United Kingdom
*Corresponding author. School of Business and Management, Queen Mary University of London, London, E1 4NS, United Kingdom. E-mail: s.k.mallick@qmul.ac.uk

Abstract

We introduce a novel approach to uncovering causation by exploiting transition episodes in aggregate macroeconomic data. Using a panel of 132 countries over 1961–2017, we identify twenty-six productivity-growth and forty savings transitions, episodes characterized by sudden, persistent, and unpredictable increases in growth and savings. Event-study evidence shows that productivity-growth transitions lead to sustained increases in savings, whereas savings transitions do not generate sustained growth effects. We conduct extensive robustness checks addressing staggered transitions, heterogeneous effects, and unobserved time-varying confounders. Having established the direction of causality, we estimate the effects of productivity-growth and productivity shocks (estimated separately for each country using a neoclassical model) on savings using panel Vector Auto-Regressions and local projections. We find that a one standard deviation increase in productivity growth raises the savings ratio by approximately 0.5 percentage points, highlighting the importance of policies that promote productivity growth to increase savings, rather than relying solely on savings-focused interventions.

Keywords productivity growth, savings, event study, local projection, counterfactual simulation

1. Introduction

The issue of savings has been important in the development literature, as savings lead to investment, which increases capital stock and output. According to Lewis (1954), ‘the central problem in the theory of economic development is to understand the process by which a community which was previously investing and saving 4 per cent or 5 per cent of its national income or less, converts itself into an economy where voluntary saving is running at about 12 per cent to 15 per cent of national income or more’. Therefore, it is not surprising that capital accumulation plays a central role in traditional growth theories. An increase in the savings ratio is important for development, but Lewis (1954) did not say much about the ways to achieve this. The savings ratio can increase due to an increase in income (growth preceding savings) or an exogenous increase in savings (savings preceding growth), which leads to an increase in income that further increases savings. There is thus a consensus about the benefits of higher saving ratios, but there is a lack of consensus on the mechanism to achieve this, which could be due to differences in belief and evidence about savings–growth causation. A positive correlation between savings and growth is omnipresent in cross-country data. We present a few instances from developing countries.

In the mid-1980s to late 1990s, as a result of international economic conditions, drought, falling petroleum prices, and years of corruption, mismanagement, and cronyism, Cameroon was hit by a serious economic crisis. The country turned to foreign aid, sharply cut government spending, and started privatizing industries to mobilize resources. Figure 1a shows the level of total factor productivity (TFP) for Cameroon. We explain the estimation of TFP later. As we can see, there was a sharp decline in TFP in Cameroon. The negative TFP growth in Cameroon was accompanied by a sharp fall in the savings ratio (Fig. 1b).[1] One can find several such instances in cross-country data. Figure 1 also presents the experience from

(a) TFP in Cameroon (b) Savings in Cameroon

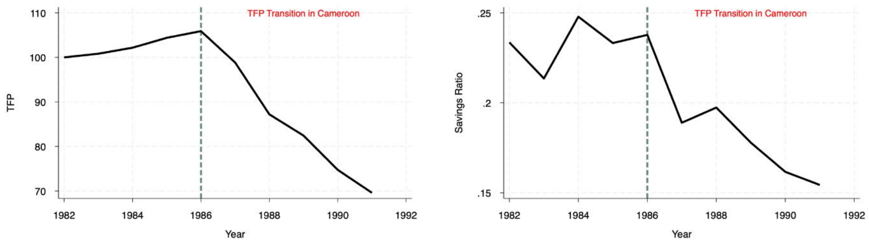

(c) TFP in Malaysia (d) Savings in Malaysia

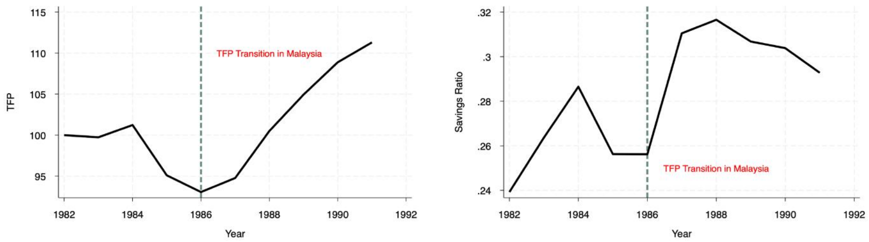

Figure 1 TFP and savings: before and after TFP transition in Cameroon and Malaysia.

Notes: The top panel shows the TFP and savings before and after a TFP transition in Cameroon. The bottom panel shows the TFP and savings before and after a TFP transition in Malaysia.

---

[1] This period was also associated with a sharp decline in the investment ratio in Cameroon. Feldstein and Horioka (1979) gave evidence that saving and investment are very tightly linked, and despite capital flows, investment is mostly explained by domestic saving. Therefore, understanding the savings-growth causation is important for understanding the drivers of investment to achieve long-run development.

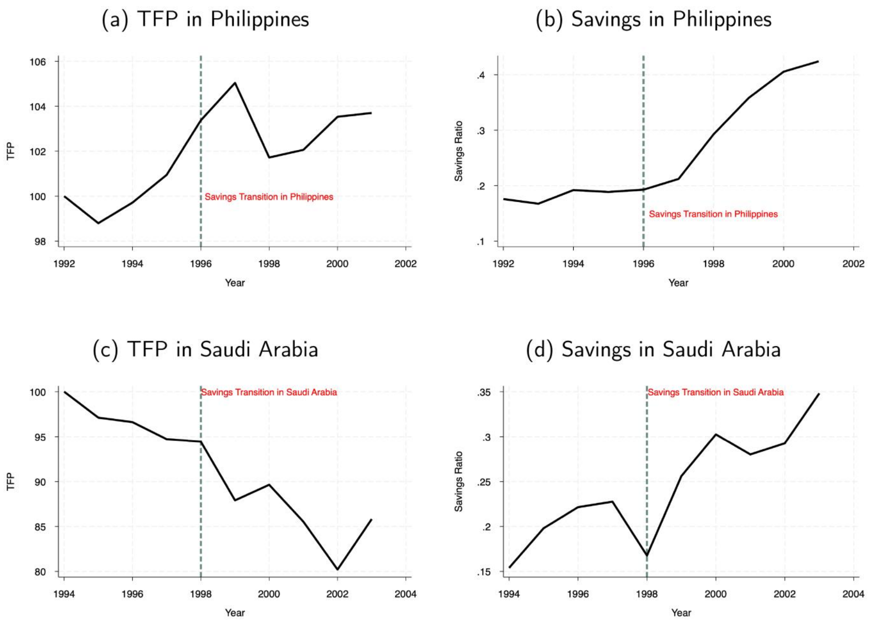

Figure 2 TFP and savings: before and after savings transition in the Philippines and Saudi Arabia.

Notes: The top panel shows the TFP and savings before and after a savings transition in the Philippines. The bottom panel shows the TFP and savings before and after a savings transition in Saudi Arabia.

Malaysia, which experienced a positive TFP transition in the same year (Fig. 1c) that increased the savings ratio in Malaysia, as shown in Fig. 1d.

Figure 2 presents the TFP and savings ratio from the Philippines, which experienced a savings transition in 1996, and the savings ratio increased by more than 20 percentage points between 1996 and 2002 as shown in Fig. 2b. However, this savings transition did not lead to a sustained increase in TFP, as evident from Fig. 2a. Figure 2 also presents the TFP and savings ratio from Saudi Arabia, which experienced a savings transition in 1998, and the savings ratio increased by almost 20 percentage points between 1998 and 2004 as shown in Fig. 2d. However, this savings transition was followed by a decline in TFP, as evident from Fig. 2c. There exists a correlation between TFP and savings in the data, but little is known about their causal relation. Understanding the causation between saving and growth is important for policymakers, as it can lead to the prioritization of policies with the correct focus. We believe that exploring the saving–growth causation is of utmost importance for understanding growth and designing the right policy mix for development.

There are important theoretical contributions in the savings preceding growth literature. Savings are essential for growth in innovation-led endogenous growth models of Grossman and Helpman (1991) where growth is preceded by investment in industrial innovation research. Aghion et al. (2016) argue that local savings are necessary for less developed countries to attract foreign capital and innovation projects, enabling growth. However, the evidence on savings preceding growth is very weak, as made evident in Figs 1 and 2 as well as in the existing literature. There are mainly two external sources that may cause an increase in savings: foreign official aid and private capital inflows, and the evidence for both is largely negative. In the 1950s and 1960s, the low saving rates in poor countries were augmented by foreign aid (Burnside and Dollar, 2000, 2004). There are studies suggesting that this aid inflow rarely worked (Easterly, 2003; Easterly et al., 2003). The recent example from Afghanistan shown in Figure S1 in the Supplementary Appendix reinforces the arguments made by Easterly (2003) and Easterly et al. (2003). In the early years, the per capita income increased with an increase in aid, but once the foreign aid started declining after 2012, the per capita income started decreasing as well. This suggests that the exogenous variation in saving through aid could not lead to sustained growth in per capita income. There was growth when inflows were increasing due to increasing demand created by aid inflows. Prasad et al. (2006) provide negative evidence about foreign inflows. According to them, countries that relied more on external finance to augment savings did not have higher long-run growth.

On the other hand, there is some positive evidence for growth preceding savings. Carroll and Weil (1994), using a purely data-driven approach, suggest that all significant increases in saving ratios across countries have been preceded by higher growth rates. For example, even after South Korea was growing rapidly, Williamson (1979) questioned the lower saving rate in the country.[2] The period of high-income growth in Japan began in the late 1940s and early 1950s, but particularly high saving rates were not established there until the 1960s and 1970s. Fernandez et al. (2019), in a simple neoclassical model of growth with government consumption, population growth, TFP growth, and capital taxation, determined a model-based savings rate and compared it with the observed savings rate.[3] They suggest that feeding the TFP growth of East Asian economies into Latin American economies increases their savings by 5 per cent. Therefore, they suggest that the lower saving rate in Latin America is also due to low TFP growth in Latin America.

In this context, we depart from the earlier literature and use productivity growth instead of GDP growth to investigate savings–growth dynamics. We believe that productivity growth is well-suited for this exercise for the following reasons. First, productivity growth is like an income shock, which we can relate to savings based on the permanent income hypothesis, and we use productivity shocks from a dynamic real business cycle model as a proxy for TFP growth in the data to estimate the effect of productivity shocks on savings. Second, by focusing on productivity growth, more precise policy initiatives can be introduced to enhance the rate of productivity growth. We estimate productivity growth using data obtained from the World Bank and Penn World Tables for 132 countries from 1961 to 2017. For these countries during the above period, we also use the savings ratio from the World Bank. The methodology for constructing TFP growth is the same as in Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999). We find that the standard deviation of productivity growth is higher in non-high-income countries, compared to high-income countries. We also obtain the productivity shocks for each country by estimating a real business cycle model based on Chari et al. (2007). Even our model-based shocks have this pattern, suggesting that productivity shocks have higher variance in non-high-income countries, in comparison to high-income countries, which substantiates the findings of Aguiar and Gopinath (2007) for a large set of countries.

We estimate contemporaneous correlation and lagged correlation between savings and TFP growth. Although the lagged effect of productivity growth on savings is higher, these reduced-form correlations in the data are inconclusive. To understand this savings–growth dynamics better, we create artificial experiments in our dataset by creating episodes of TFP growth and savings transitions. These transitions are, by design, sudden and persistent increases in TFP growth and savings rate for countries experiencing transitions and occur at different points in time. We also find that these transitions are not predictable using a large

---

[2] Korea, as used in this article, refers to South Korea.
[3] In this article, we use TFP and productivity interchangeably.

set of macroeconomic variables widely used in the literature. We identify twenty-six TFP growth and forty savings transitions, displaying sudden, persistent, and unpredictable increases in growth and savings. We treat these transitions as an event or treatment and estimate the effect of TFP growth transitions on savings, and savings transitions on TFP growth. We start with a widely used two-way fixed-effect model (TWFE) in which we include only countries experiencing transition and compare them before and after transitions. This design is called a timing-based data structure according to Miller (2023) because the variation arises due to differences in transition years for different countries. The crucial assumption for identification is the constant effect of these transitions on outcomes and the absence of time-varying confounders. These results suggest that the increase in TFP growth is followed by a sustained increase in the savings ratio, whereas the savings transitions are not followed by an increase in TFP growth. These results also do not show any significant deviation from parallel trend before transition, especially for TFP growth transitions. This makes these results credible as parallel trends suggest the absence of time-varying confounders (Miller 2023).

It can be argued that the parallel trend obtained using TWFE is not conclusive evidence for the absence of time-varying confounders for two reasons. First, the parallel trend is obtained using the assumption of a homogeneous transition effect, which is hard to justify. This is because these countries experience transition at different points in time and are heterogeneous. Second, there are multiple periods involved, and time-varying confounders need not be trending. These two make the parallel trend less credible evidence in this setting. To mitigate concerns related to homogeneous transition effect and unobserved time-varying confounders, we estimate the generalized version of the interactive fixed effects model using matrix completion (MC) methods (Xu 2017; Athey et al. 2021; Liu et al. 2024). These models allow heterogeneous transition effects and control for time-varying confounders by incorporating common time-varying factors that affect countries differently. As a result, concerns related to homogeneous transition effects are resolved, and concerns related to time-varying confounders are largely mitigated. Moreover, the presence of parallel trends—if observed—provides even stronger evidence, when we explicitly allow for both heterogeneous treatment effects and time-varying confounders, which are two factors that could undermine the credibility of the parallel trends assumption. The MC method also shows no deviation from parallel trend and suggests that an increase in TFP growth is followed by a sustained increase in the savings ratio up to ten years, whereas savings transitions are not followed by a sustained increase in TFP growth. We do several exercises in this article for robustness and use alternative TFP growth and savings transitions, event windows, and Placebo tests to show that these results are robust.

The results discussed above suggest that the impact of TFP growth on savings is persistent, but these effects are hard to quantify in an easy-to-understand way, that is, percentage point change in savings ratio due to percentage point change in growth. This is because the countries experiencing transitions have different changes in TFP growth. To address this, we estimate the effect of TFP growth on the savings ratio using the local projection framework of Jorda (2005) where we include per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, foreign direct investment as percent of GDP, credit as percent of GDP, young dependency ratio, old dependency ratio, country fixed effects, year fixed effects, and four lags of TFP growth and savings ratio. These local projection regressions suggest that a one standard deviation increase in TFP growth increases the savings ratio by 0.5 percentage points, and the effect is significant for up to 6 years. We also do robustness using a range of Panel VAR models with different lags and different endogenous variables and estimate the response of the savings ratio due to TFP growth using Cholesky decomposition.

The most comprehensive model with eight endogenous variables gives a very similar effect of TFP growth on the savings ratio.

Further, to eliminate the concerns related to unobserved omitted variable bias, we use the productivity shock from the neoclassical growth model (real business cycle model) for forty-nine countries in our sample. Productivity shock has been widely used in macroeconomic research to understand the business cycle (see Gali 1999; Basu and Fernald 2002; Gali et al. 2003; Basu et al. 2006; Gortz et al. 2022). These shocks may have further advantages than reduced-form TFP growth, as these shocks are estimated conditional on investment wedges in the model. The investment wedge in the model captures investment-specific technology shocks and financial friction (Brinca et al. 2016). We estimate the local projection regressions using these productivity shocks and find qualitatively similar responses of the savings ratio due to these shocks, but the effect is slightly higher and less persistent. Based on these findings, we argue that countries should focus on promoting policies to boost TFP growth and thereby achieve higher savings instead of focusing on savings-induced policies alone.[4]

Our results are also related to a large literature on the determinants of cross-country variation in saving rates, encompassing both theoretical models and empirical studies of saving behaviour. Modigliani (1970) showed, based on a large cross-section of countries, that both the ratio of old-age population (aged 65 and over) and of young population (below 20) to the working-age population (20–65) had a strong and highly significant negative effect on the saving ratio.[5] Loayza et al. (2000) found that real interest rate, per capita income and its growth, ratio of money supply to gross national disposable income, old-age dependency ratio, young-age dependency ratio, terms of trade, inflation rate, and financial sector development are important determinants of saving. One can also argue that the effect of saving-enhancing policies on saving depends upon the use of savings in respective countries. We control for the terms of trade to capture this channel in our regressions. This allows us to control for large increases in savings associated with the accumulation of foreign exchange reserves in developing countries. Buera and Shin (2017), using a model, show that an increase in TFP growth increases the savings rate, as found in this article, but they rely on a theoretical model and simulation to argue that higher TFP growth is associated with a higher savings rate. The empirical results in this article validate the theoretical results of Buera and Shin (2017). Carroll (1994) provided evidence that consumption does not respond to any predictable change in income, but it responds to future income uncertainty, in line with the buffer stock model of saving. Based on the findings in Carroll (1994), one can argue that unexpected changes in income, such as shocks to productivity growth, may increase savings. Aizenman and Noy (2013), on the other hand, bring another dimension to the determinants of household savings and argue that households that have experienced a higher number of adverse shocks in the past tend to save more. Mody et al. (2012) use a panel of Organisation for Economic Co-operation and Development countries and argue that at least two-fifths of the increase in households' saving rates between 2007 and 2009 was due to increased uncertainty about labour–income prospects. Zabot and Gomes (2025) show that higher permanent income leads to greater savings and that saving patterns differ systematically across wealth levels.

---

[4] This is no coincidence that policy discourses are more about reforms in low-income countries to promote growth rather than aid. Moreover, the recent literature suggests more specific ways of enhancing TFP, such as reducing misallocation as in Hsieh and Klenow (2009), and David et al. (2016), and improving law enforcement as in Erosa and Cabrillana (2008). Hence results in this article are of interest for informed growth-enhancing policies.

[5] In recent times, the literature on types of savings and its determinants has also emerged. Chen et al. (2017) explore the rise in corporate savings in the last three decades and argue that it has been accompanied by a decline in labour share and stable dividends.

The plan of the article is as follows. The section " Data and TFP growth" explains the data and estimation of TFP from the data and the model, and shows some stylized facts about savings and TFP growth. Section "Effect on saving and TFP …" explains saving and growth transitions and the empirical framework. "Results" section presents results from event studies and results from counterfactual estimation, local projection, and counterfactual simulation. "Concluding remarks" concludes and suggests policies resulting from the empirical evidence in this article.

2. Data and TFP growth

2.1 Data

We estimate a human capital-adjusted measure of TFP from data. Our main source of data is the World Bank and Penn World Table. The gross savings ratio (% of GDP), is obtained from the World Bank. Our main variable of interest is productivity (TFP) growth. Estimation of TFP requires data on per-capita income and physical and human capital. Per capita income is obtained from the World Bank. Capital is compiled from the Penn World Table and transformed into per capita terms. We use two measures of mean years of schooling for the labour force: (1) mean years of schooling for the population aged 15+, obtained from the Barro-Lee data set, Barro and Lee (2013)[6]; and (2) mean years of schooling are obtained after correcting for differences in labour force participation for males and females. Average years of schooling for females have increased across the globe, but their labour force participation has fallen. This means that the average years of schooling are biased in capturing the schooling years of the labour force. We use mean years of schooling for females aged 15+ obtained from the Barro-Lee data set, and female population aged 65 and above, male population aged 65 and above, total population aged 65 and above, the female population aged 15–64, male population aged 15–64, total population aged 15–64, labour force, females as percentage of labour force, labour force participation rate for females, males, and total population to obtain the second measure of human capital.[7]

Average years of schooling and human capital are related through:

$$h = e^{\phi(s)} \tag{1}$$

where s is the average years of schooling. We assume $\phi(s)$ to be piece-wise linear as in Hall and Jones (1999). The slope of $\phi(s)$ is 0.13 for $s \leq 4$, 0.10 for $4 < s \leq 8$, and 0.07 for $s > 8$. Using average years of schooling in a country in a particular year, we can obtain human capital for the country–year pair.

We also use old dependency ratio, young dependency ratio, domestic credit to private sector (% of GDP), trade (% of GDP), net barter terms of trade index (2,000 = 100), FDI (net

---

[6] The Barro-Lee data set is till 2010, observed at five-year intervals beginning in 1960. We interpolate this data to obtain values between observed data points. Years of schooling change slowly and should not lead to any issues.

[7] The Barro–Lee data set contains average years of schooling for the fifteen plus population and the fifteen plus female population. Using the expression below, we calculate the average years of schooling for fifteen plus males: $\text{Average years of schooling}_{15+} = \frac{\text{Male population}_{15+}\,\text{Male average} + \text{Female population}_{15+}\,\text{Female average}}{\text{Male population}_{15+} + \text{Female population}_{15+}}$ Once we have the average years of schooling for fifteen plus males, we calculate the average years of schooling for those participating in the labour market using the expression below: $\text{Average years of schooling}_{\text{LF}} = \frac{\text{Male labour}_{15+}\,\text{Male average} + \text{Female labour}_{15+}\,\text{Female average}}{\text{Male population}_{15+} + \text{Female population}_{15+}}$

inflows as % of GDP), Per capita income as control variables which are obtained from the World Development Indicators (World Bank). These variables have been used widely in the literature (see Loayza et al. 2000; Mody et al. 2012). Figures S2–S4 in the Supplementary Appendix give the distribution of the main variables used in this article.200020002

2.2 TFP growth from data

The basic Solow model is given by:

$$Y_t = A_t K_t^{\alpha} (L_t h_t)^{1-\alpha}$$

where L_t is the number of employed person, K_t is capital stock, and h_t is average human capital of employed person. Productivity is assumed to be Hicks-neutral. Dividing by L_t, we get

$$y_t = A_t k_t^{\alpha} h_t^{1-\alpha} \tag{1}$$

where y_t and k_t are output and capital per worker. The two measures of human capital used in this article are highly correlated with the measure of human capital in the Penn World Table (PWT) table. Once we know y_t, k_t, and h_t, it is straightforward to obtain TFP using Equation (1). It is important to mention that Equation (1) is based on level accounting to derive TFP and calculate TFP (productivity) growth. Another way of calculating TFP growth is based on growth accounting.[8] Also, one can represent level accounting in terms of the capital–output ratio as done in Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999).

$$A_t = \frac{1}{\left(\frac{k_t}{y_t}\right)^{\alpha} h_t^{1-\alpha} y_t^{\alpha-1}} \tag{2}$$

where y_t and k_t are per capita output and capital. Knowing y_t, k_t, and h_t from Equation 1, we can obtain estimates of A_t. Thus, we obtain three estimates of TFP from the data: (1) Hicks-neutral productivity growth using average years of schooling for the population aged fifteen plus, which is called TFP growth in the article; (2) Hicks-neutral productivity growth using average years of schooling for the population aged 15+, corrected for differences in labour force participation; (3) Hicks-neutral productivity growth without human capital using $y_t = A_t k_t^{\alpha}$. The last two measures of productivity growth are highly correlated with the first measure, with a correlation 0.99. This is because of the very low variation in human capital in the data. Hence, we use only the first measure to obtain the results in the article.[9] The other two measures give similar results.[10] We obtain TFP estimates for 132 countries during

---

[8] This is also called Solow Residual and is the component of output growth that cannot be explained by the growth in labour and capital alone, implying it is due to technological progress or improvements in the efficiency of production (Solow 1957). Young (1994) is another well-cited work on Solow residual, which argues that the East Asian Miracle was predominantly a factor accumulation phenomenon and not an illustration of technological progress alone.

[9] The low variation in the estimated human capital could be driven by two factors. First, years of schooling may be an inadequate proxy for human capital. Second, the use of a homogeneous mapping across countries using the same h may mechanically compress variation. But the main results in this article should not be driven by these measurement errors in TFP due to human capital. This is because we use a sudden and persistent large increase in TFP growth to identify a TFP growth transition, as shown in the next section. It is very unlikely that changes in human capital measures should cause these TFP growth transitions

[10] A Harrod-neutral productivity process is given by $y_t = k_t^{\alpha}(Z_t h_t)^{1-\alpha}$ where Z_t is the TFP. Comparing the Hicks neutral and Harrod neutral productivity processes, we can see that Z_t=$A_t^{1/(1-\alpha)}$. Hence, the Harrod neutral TFP is a non-linear multiple of Hicks neutral productivity and hence the direction of the relationship obtained between savings and TFP using Hicks neutral TFP will also hold with Harrod neutral TFP, but the coefficient may vary in magnitude. We only use Hicks neutral TFP in this article.

1961–2017. We have data from very few countries before 1980. Burnside et al. (1996), and Fernald (2014) argue for the use of capacity utilization to estimate the more precise value of Solow residual, that is, productivity growth for the US economy. In a recent paper, Comin et al. (2020) extend the estimation of capacity utilization-adjusted productivity growth for a set of European countries. Due to the paucity of capacity utilization data for a large number of countries being used in this study, we are unable to adjust the measure of productivity growth for capacity utilization. However, we do some robustness exercises using TFP adjusted for capacity utilization for a set of countries for which data has been provided by Comin et al. (2020).

2.3 TFP growth across income levels

Figures 3a and 3b give the distribution of TFP growth in non-high and high-income countries respectively. Before the 2000s, the average and median TFP growth in non-high-income countries was very low and often negative compared to mostly positive TFP growth in high-income countries. However, TFP growth in non-high-income countries surged after the 2000s. Both high and non-high-income countries experienced a sharp decline in TFP growth during the financial crisis. Figure S5a in the Supplementary Appendix gives the mean standard deviation of TFP growth over time in our sample of high- and non-high-income countries. We use the World Bank classification to select high-income countries; the remaining countries in our sample are categorized as non-high-income countries. With productivity growth and country classification, we construct the standard deviation of TFP growth on a recursive basis across these two sets of countries and then take the average over time and across groups. As we can see, the standard deviation of productivity growth is higher in non-high-income countries than in high-income countries. This substantiates the findings of Aguiar and Gopinath (2007) for a large set of countries where productivity growth is more volatile than in non-high-income countries. Figure S5b in the Supplementary Appendix gives the group and year-wise average of the 10-year rolling standard deviation of TFP growth. Even in this case, we find the TFP growth is more volatile in non-high-income countries, but in recent years, the TFP growth in high-income countries has become more volatile than in non-high-income countries.

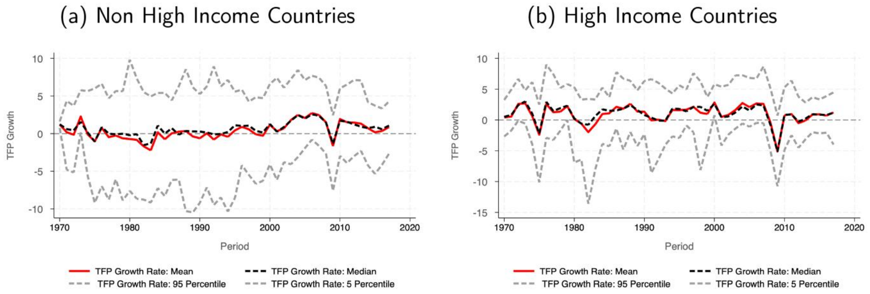

Figure 3 TFP growth distribution.

Notes: (a) Hicks neutral productivity growth is obtained with average years of schooling for the population aged 15+.

Figure S6 in the Supplementary Appendix gives a contemporaneous correlation between these two measures of TFP growth and savings ratio. These correlations are around 0.13 and statistically significant. Figures S7 and S8 in the Supplementary Appendix give the correlation between the saving ratio and TFP growth for all countries in the sample at lag 1 and lag 2. The correlation between the saving ratio and lag 1 of productivity growth in our sample is 0.14, whereas the correlation between growth rate I and lag 1 of the saving ratio is 0.1. The correlation between the saving ratio and lag 2 of productivity growth I in our sample is 0.14, whereas the correlation between growth rate and lag 2 of the saving ratio is 0.08. With both lags, the correlations between the saving ratio and lagged productivity growth, and the correlations between productivity growth and lagged savings ratio are significant. These simple correlations are not useful in giving us the direction of causality. If lagged productivity growth had significant correlations with savings, but lagged savings did not have a significant correlation with productivity growth, that would have suggested likely causation from productivity growth to savings. But that is not the case, and deciding the direction of causality is an empirical issue that we address in the next section. But these correlations suggest that the effect of productivity growth on saving is higher and more persistent, which is also evident from Figs 1 and 2. Throughout this article, our estimates using different definitions of productivity growth suggest that the effect of productivity growth on savings is indeed persistent and significant.

2.4 TFP from neoclassical growth model

In this subsection, we use a neoclassical growth model to obtain TFP shocks. The TFP in the previous section was obtained using results by Hall and Jones (1999), to transform years of schooling into a composite measure of human capital. The assumption is that the technology mapping education into effective human capital is stable both across time and different regions of the world, which may not be true. Thus, we estimate TFP shocks (productivity shocks) from a model that does not use this assumption and also helps mitigate other concerns related to reduced-form estimates from the previous section, as we explain below. We estimate the model for forty-nine countries (Table S1 in the Supplementary Appendix) in our sample for which at least 40 years of observation are available individually, using their data, and obtain the model-based productivity shocks. We use these productivity shocks in our local projection regressions along with TFP from the data. We use the simplest neoclassical growth model following Chari et al. (2007). The representative household utility is given by:

$$U(c_t, n_t) = \log(c_t) + \chi \log(1 - n_t)$$

The household maximizes lifetime utility given by:

$$\sum_{t=0}^{\infty} \beta^t [\log(c_t) + \chi \log(1 - n_t)]$$

Their period-by-period budget constraint is given by:

$$c_t + (1 + \tau_{i,t}) I_t = w_t n_t + r_t k_{t-1} + T_t$$

where $(1 + \tau_{i,t})$ is the investment wedge as in Chari et al. (2007). $\tau_{i,t}$ is the stationary zero mean process given by:

$$\tau_{i,t} = \rho_i \tau_{i,t-1} + \varepsilon_{i,t}$$

There is a growing literature that suggests that investment-specific technological change plays an important role in business cycles (Greenwood et al. 1997). Investment-specific technological shock influences the relative price of investment, and therefore affects investment. Our model is parsimonious, with few frictions, but Brinca et al. (2016) show that the investment-specific technology shock maps into our economy with investment wedges. Financial frictions are also important for investment, and the financial frictions studied in Kiyotaki and Moore (1997) and Gertler and Kiyotaki (2009) also map investment wedges. Therefore, we are not modeling investment-specific technology shocks and financial friction explicitly, but our investment wedge should capture them. A household accumulates capital, and the law of motion for this is given by:

$$k_t = (1-\delta)k_{t-1} + I_t$$

The representative firm uses labour and capital to produce output using the Cobb–Douglas production function, which is given by:

$$y_t = A_t^{1-\alpha} k_{t-1}^{\alpha} n_t^{1-\alpha}$$

$$\log(A_t) = \rho \log(A_{t-1}) + \varepsilon_{A,t}$$

There is government expenditure that fluctuates around a steady-state level of the government-expenditure-output ratio:

$$G_t = \left(1 - \frac{1}{g_t}\right) y_t$$

$$\log(g_t) = (1-\rho_g)\log(g) + \rho_g \log(g_{t-1}) + \varepsilon_{g,t}$$

2.4.1 Model estimation

We find the steady state of the above model after obtaining the first-order condition, and the model is log-linearized. Section S15 in the Supplementary Appendix gives first-order conditions, steady state, and log-linearized equations. We estimate the model using maximum likelihood for each of the forty-nine countries one by one using the approach outlined in Ireland (2011) using Dynare. Table S1 in the Supplementary Appendix gives the list of these countries and also the period used in the estimation. We calibrate g as the average value of the government expenditure to GDP for each country, and this varies across countries. We calibrate the same value of $\beta = 0.97$, $\delta = 0.1$, and $\alpha = n = 1/3$ across countries. Cyclical components of log per capita income, and private and government consumption are used to estimate the model.[11] We use two methods to obtain cyclical components: linear trend and

---

[11] We use general government final consumption expenditure, household final consumption expenditure, gross domestic savings ratio, total population, and GDP per capita from the World Bank Development Indicators (WDI) dataset. To make the model results comparable with the savings rate in the data, we use GDP and savings ratio from the World Bank to obtain total savings. Further, the sum of total consumption and savings is output (GDP). We keep countries with at least 40 years of observations, and all variables are transformed into per-capita terms. Table S1 in the Supplementary Appendix gives the list of countries with the period used in the estimation.

the Hamilton filter. Hamilton (2018) suggests eight forward lags and four backward lags to obtain a cyclical component using quarterly data. Since we use annual data, we take two forward lags and one backward lag to obtain a cyclical component. Thus, we obtain two measures of productivity shocks $\varepsilon_{A,t}$—using linear trend and Hamilton filter—for each country from the estimation. Figures S9 and S10 in the Supplementary Appendix give the standard deviation of TFP shocks across income levels, which have patterns similar to TFP growth. We understand that exclusion of human capital and homogeneity in discount rate and labour share in these estimates imply that these are not necessarily better than the estimates obtained from data using human capital, and are used for robustness.

3. TFP growth and savings transitions and empirical framework

3.1 TFP growth and savings transitions

In our data set, we have many countries that experienced a significant change in their savings ratios as well as TFP growth in a very short period. We showed some of these examples in Figs 1 and 2. We use this to identify savings and growth transitions. Figure S11 in the Supplementary Appendix gives a schematic representation of savings and TFP growth transitions. Saving transition is defined as follows: at each point in time, we calculate the 3-year forward-moving average of the saving ratio (the current year and the 2 years after that) and the 5-year backward-moving average of the saving ratio (the current year and the four years before that). A savings transition is said to have occurred in a year for a country if the forward-moving average in that year and up to 6 years ahead is greater than 5 per cent of the backward-moving average one year before. The top decile of the annual change is 4 per cent (Table S2 in the Supplementary Appendix) in the savings rate in the sample, and hence 5 per cent is a comparatively large and sudden change in the savings rate. We also impose that the forward-moving average in that particular year and up to 6 years ahead must be greater than 21.5 per cent. The median of the average savings rate of countries in the sample is 21 per cent (Table S2 in the Supplementary Appendix). We retain data for all countries covering 5 years before and 5 years after the transition for describing transitions and the baseline estimates. We also use data up to 10 years after the transition for the main estimations.

Technically, the transition is defined as follows: Let s_f and s_b be the forward and backward moving averages as mentioned. Also, $l.j$ and $f.j$ denote the j period lag and forward, respectively. Technically speaking, our transition condition is $s_f > l.s_b + 5$ and $f.s_f > l.s_b + 5$ and $f2.s_f > l.s_b + 5$ and $f3.s_f > l.s_b + 5$ and $f4.s_f > l.s_b + 5$ and $f5.s_f > l.s_b + 5$ and $f6.s_f > l.s_b + 5$ and $s_f > 21.5$ and $f.s_f > 21.5$ and $f2.s_f > 21.5$ and $f3.s_f > 21.5$ and $f4.s_f > 21.5$ and $f5.s_f > 21.5$, and $f6.s_f > 21.5$. For persistence, we also impose $f1.s_f > s_f$. Effectively, these conditions imply a large, sudden, and persistent increase in the savings rate that transitions a country to have a savings rate above the median savings rate in our sample. For example, a country having a backward-moving average of 16 per cent needs to have a sudden 5.5 per cent increase in the forward-moving average of the savings rate to satisfy the savings transitions. A country with a lower than 16 per cent rate will have to increase savings of more than 5 per cent to satisfy the transition. Countries with higher savings rates can also satisfy the transition, but will have higher savings rates (>21.5 per cent) after the transition. The transition is ensured by the condition on the forward moving average at time $t = 1$, and further conditions on future moving averages make this transition persistent. As we can see from Fig. 4a, the transitions lead to a 10 percentage point increase in the savings rate around the transition and a higher

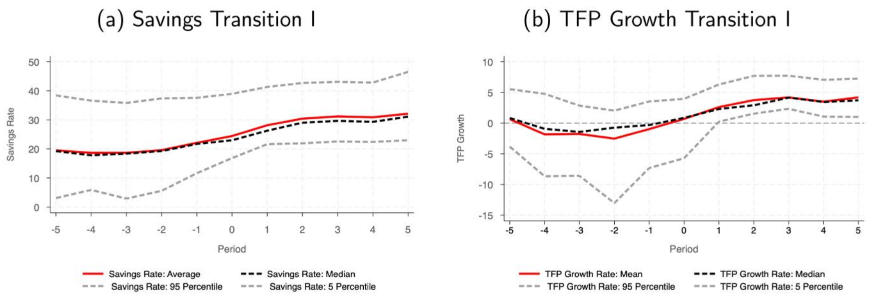

(a) Savings Transition I

(b) TFP Growth Transition I

Figure 4 Savings and TFP growth transitions.

Notes: Savings transition to the median of the savings rate: TFP Growth transition to the third quartile of average TFP growth.

increase in the savings rate for countries having lower savings before the transition. We define additional measures of savings transitions, for robustness, which are provided in the Supplementary Appendix. Table S3 in the Supplementary Appendix gives the list of countries experiencing these savings transitions, with the year of transition. Table S4 in the Supplementary Appendix gives results from a Logit and linear probability model, suggesting that these savings transitions cannot be predicted using the country-specific variables widely used in the literature. Since the results from savings transition II and III are available in the Supplementary Appendix, in the rest of the article, we refer to this savings transition I as the savings transition.

Similarly, we define transition in TFP growth. Technically, this is defined as: Let g_f and g_b be the forward and backward moving averages of TFP growth I, respectively. Technically speaking, our transition condition is $g_f > l.g_b + 1.89$ and $f.g_f > l.g_b + 1.89$ and $f2.g_f > l.g_b + 1.89$ and $f3.g_f > l.g_b + 1.89$ and $f4.g_f > l.g_b + 1.89$ and $f5.g_f > l.g_b + 1.89$ and $f6.g_f > l.g_b + 1.89$ and $g_f > 1.22$ and $f.g_f > 1.22$ and $f2.g_f > 1.22$ and $f3.g_f > 1.22$ and $f4.g_f > 1.22$ and $f5.g_f > 1.22$, and $f6.g_f > 1.22$. 1.22 and 1.89 are the third quartile and top decile of the country-wise average of TFP growth in the sample (Table S5 in the Supplementary Appendix). For persistence, we also impose $f1.g_f > g_f$. Effectively, these conditions imply a sudden and persistent increase in the TFP growth rate of 1.89, bringing the country's TFP growth above the third quartile of the average TFP growth. For example, a country with a backward moving average of −0.67 per cent needs to have a sudden 1.89 per cent increase in growth rate to satisfy the growth transitions. A country having a lower than −0.67 per cent growth rate will have to have an increase in growth of more than 1.89 per cent to satisfy the transition. Countries with higher growth rates can also satisfy the transition, but will have higher growth rates (>1.22 per cent) after the transition. As we can see, the transitions lead to a 4 per cent point increase in the growth rate around the transition, Fig. 4b. We define additional measures of TFP growth transitions, for robustness, which are provided in the Supplementary Appendix. Table S6 in the Supplementary Appendix gives the list of countries experiencing these transitions with the year of transition. Tables S7 and S8 in the Supplementary Appendix give results from a Logit and linear probability model, suggesting that these growth transitions cannot be predicted using the country-specific variables widely used in the literature. TFP growth transitions II, III, IV, and V are used for robustness, and all the results for these transitions are available in the Supplementary

Appendix. Hence, in the rest of the article, we refer to TFP growth transition I as TFP growth transition.

For defining transition in TFP growth, we choose a smaller difference, as the growth rate in TFP is not likely to be big enough. Then we keep data for all countries before and after 5 years of transition. In all these transitions, a higher increase in TFP growth is observed in the transitioning countries having low TFP growth rates before transition, which is expected because we use the criterion that all transitioning countries have a forward-moving average greater than the median of the country-wise average of respective TFP growth rates. Our savings transitions are, by design, going to bring a significant increase in savings rate, as we can see from Fig. 4a. As we can see from Fig. 4a, the median savings rate of transitioning countries increases by 10 per cent, and this increase is even higher for countries having lower growth rates before the transition. Similarly, the TFP growth transition is going to bring a significant increase in TFP growth, Fig. 4b. As we can see from Fig. 4b, the median growth rate of transitioning countries increases by more than 4 per cent and this increase is even higher for countries having lower growth rates before the transition. Our design allows us to compare these two episodes as they lead to sudden and persistent changes in savings and growth rates. An event study design is suitable as these transitions are sudden, and are not predictable using country-specific variables, as we show in Tables S4, S7, and S8 of the Supplementary Appendix. If savings cause growth, then an increase in savings should be associated with an increase in TFP growth, whereas if TFP growth causes savings, then an increase in TFP growth would lead to an increase in savings.

3.2 Empirical framework

The previous section explains the TFP growth and savings transitions. In this section, we describe the empirical strategy to estimate the effect of these transitions by treating these transitions as a treatment/event. We start with a classic TWFE in an event study setting. The event study setting is better than the difference-in-differences (DID) for two reasons. First, it gives the dynamic effects of TFP growth and savings transition rather than the average effect given by DID, and second, it helps in detecting pre-trend or violations from parallel trend, which can lead to biased inference. The baseline model is given by Equation 3.

$$y_{it} = \sum_{-5}^{5} \gamma_j I_j + \phi_i + \phi_t + e_{it} \tag{3}$$

where y_{it} is either TFP growth or savings in country i at time t, and I_j is an indicator variable depicting the transition window around either savings or TFP growth transitions at $j = 0$. ϕ_i captures the country-specific time-invariant heterogeneity and ϕ_t captures the time-varying common factors. We have two options to estimate this model. First, using only data from countries experiencing TFP growth and savings transition, and second, using data from all countries. These are called timing-based and hybrid data structures, respectively, according to Miller (2023). The timing-based structure can also be treated as pure event studies, as we compare the countries experiencing transition before and after transition, and we start with that. The underlying assumption is that the timing of the transition is as good as random, and hence those countries experiencing transition earlier or later can serve as controls for one another. As mentioned before, we find that these transitions are unpredictable using the range of macroeconomic variables used in the literature. The identifying assumption for the above model is a constant effect of transition and the absence of time-varying confounders. Time-varying confounders are variables correlated with treatment and potential outcomes of interest. In a two-period and two-group setting, a constant treatment effect is satisfied by design, and the presence of parallel trend before treatment satisfies the assumption of the absence of time-varying confounders, leading to the causal effect.

In more than two-group and two-period settings, such as the one used in this article, the inference from Equation 3 could be problematic for two reasons. First, the constant transition effect is hard to justify, as we have more than twenty (up to sixty-three) transitions occurring at different times in different countries. Second, the parallel trend before transitions is not conclusive evidence for the absence of time-varying confounders because the estimates leading to parallel trend are obtained using constant transition effects, which could not be true, and a time-varying confounder need not be trending (Liu et al. 2024). Both these violations could lead to biased estimates, which we attempt to address by estimating additional models. First, we include several variables that could be related to savings and TFP growth based on the literature, and estimate the model given by Equation 4.

$$y_{it} = \sum_{-5}^{5} \gamma_j I_j + \theta' z_{it} + \phi_i + \phi_t + e_{it} \tag{4}$$

where z_{it} includes per-capita income and squared per-capita income, terms of trade, trade to GDP ratio, foreign direct investment as % of GDP, credit as % of GDP, young dependency ratio, and old dependency ratio. Many of these control variables have missing values, and we replace the missing values with country-specific median values of these control variables. We understand that imputing some of these variables, such as foreign direct investment as % of GDP, with a median may have issues. Hence, we estimate a model with country-specific linear and quadratic trends to capture confounders that vary across countries and over time. Lovenheim and Willén (2019) also argue that the main threat to identification is the existence of secular trends. Lovenheim and Willén (2019), we estimate models with country-specific linear trends given by Equation 5. The problems of multicollinearity also compound when we add country-specific trends. There are two options suggested by Miller (2023)—either we impose explicit restrictions for the parameters to be identified, or drop one of these country-specific trends, as done in Stata, and we go for the second option.

$$y_{it} = \sum_{-5}^{5} \gamma_j I_j + \phi_i \times \text{Time Trend} + \phi_i + e_{it} \tag{5}$$

Further, we include a country-specific quadratic time trend. This is relevant given the fact that many of the countries transitioning to high TFP growth and savings rates in Fig. 6 show a parabolic pattern that could be driven by the country-specific quadratic trend. We estimate Equation 6 given below:

$$y_{it} = \sum_{-5}^{5} \gamma_j I_j + \phi_i \times \text{Time Trend} + \phi_i \times \text{Time Trend}^2 + \phi_i + e_{it} \tag{6}$$

Although the above models help in addressing issues related to time-varying confounders, these models are estimated assuming constant transition effects, which remains a concern, and the absence of pre-trend in these models is not conclusive evidence of the absence of time-varying confounders. Hence, we use counterfactual estimators to mitigate these concerns.

3.2.1 Counterfactual estimators: hybrid data structure

All counterfactual estimators used in this article estimate the potential outcomes for transitioning countries using data before transition and from non-transitioning countries, and consider heterogeneous treatment effects because the assumption of constant transition effects is not required. By design, these models are estimated using a hybrid data structure.[12] The estimates from these models are more robust compared to the previous ones, as this allows for a heterogeneous transition effect. Also, the absence of pre-trend is more conclusive evidence of the absence of time-varying confounders, as these are obtained assuming heterogeneous transition effects. The first counterfactual estimator is given by Equation 7.

$$y_{it} = \delta_{it} D_{it} + \theta' z_{it} + \phi_i + \phi_t + e_{it} \tag{7}$$

D_{it} is the transition indicator that equals 1 if country i has experienced either TFP growth or savings transition at time t and equals 0 otherwise; δ_{it} is the effect of transition on country i at time t. The above Equation (7) is different from a canonical two-way fixed effects model given by (1) because we allow the transition effect to be heterogeneous both across countries and over time (Liu et al. 2024). This is called a fixed effect counterfactual estimator (FEct) by Liu et al. (2024). The two-way fixed effect estimator given by (3) is a special case of the model (7). The above model also differs from models 3–6 as it includes both sources of identification: the comparison of transitioning and non-transitioning countries and the timing of the event. Jacobson et al. (1993) employ a hybrid data structure.

The estimation is done in three steps. First, we estimate a TWFE using non-transitioning observations and obtain the parameters. In the second step, using these parameters, we obtain the counterfactual for transitioning countries, and in the third step, we obtain δ_{it} as the difference between these two and take an average of that to obtain the average treatment effect on the treated (ATT), i.e. on transitioning countries. Note that δ_{it}s are not individually identifiable, while their averages are. We estimate Equation (7) with country-level covariates z_{it}. We estimate two variants of the above model with quadratic and cubic time trends. The identification from the above model relies on three assumptions: (1) Past outcomes do not directly affect current transition (no feedback), (2) past transitions do not directly affect the current outcome, and (3) absence of time-varying confounders, Liu et al. (2024). We argue that assumption (1) is plausibly satisfied, as we find that these transitions are unpredictable using the range of macroeconomic variables used in the literature. Assumption (2) is satisfied by design as we do not consider multiple transitions[13] or rollback from a transition in these estimates, as we assume that these transitions are almost permanent, hence carryover concerns are not applicable in this setting. Although the above model addresses heterogeneous treatment, and pre-trend is a more credible evidence of the absence of time-varying confounders, the concerns related to time-varying confounders are not eliminated. This is because the time-varying confounders need not be trending as mentioned above. Hence, we estimate an additional model to mitigate the concerns related to time-varying confounders.

Several authors (Gobillon and Magnac 2016; Xu 2017) have proposed the use of factor-augmented models to address this issue when the time-varying confounders can be decomposed into time-specific factors and interact with country-specific factor loading, as given by:

---

[12] One can estimate a two-way fixed effect model with constant transition effect given by (1) using a hybrid data structure as well, but we do not pursue that given its limitations.

[13] We drop only a few transitions to ensure this.

$$y_{it} = \delta_{it} D_{it} + \theta' z_{it} + \lambda_i' f_t + \phi_i + \phi_t + e_{it} \tag{8}$$

Where f_t is an $r \times 1$ vector of unobserved factors and λ_i is a $r \times 1$ factor loading. Intuitively, factors can be understood as time-varying trends and other common factors that affect each country differently, and factor loadings capture their heterogeneous impacts on countries. Because $\lambda_i' f_t$ captures such differential country-specific time-varying factors, the no-time-varying-confounder assumption will be relaxed to a large extent and is the best available option for mitigating these concerns to a certain extent. We need to impose restrictions on λ_i and f_t to achieve identification. Equation 8 implies that

$$y_{it}(0) = \theta' z_{it} + \lambda_i' f_t + \phi_i + \phi_t + e_{it} \tag{8.1}$$

$$y_{it}(1) = y_{it}(0) + \delta_{it} D_{it} \tag{8.2}$$

where 1 and 0 stand for with and without transitions, respectively. The above model is called the interactive FEct by Liu et al. (2024). Another way to address the time-varying confounders is the MC method by Athey et al. (2021). Mathematically, MC assumes that the $(N \times T)$ matrix of Y(0) can be approximated by a low-rank matrix L and the model is given by:

$$Y(0) = z\theta + L + \varepsilon \tag{9}$$

In $Y(0)$, the outcomes are missing for transitioning countries post-transition. Equations (9) and (8.2) are similar as the matrix L can be factorized as $L = \Lambda F$, where Λ and F are conforming matrices. The only difference between Equation (8.2) and (9) is the estimation approach. IFEct obtains Λ and F whereas MC obtains L directly (see Liu et al. 2024; Athey et al. 2021 for details). Athey et al. (2021) propose an iterative algorithm to solve for L and demonstrate that the resulting estimator is asymptotically unbiased for the untreated outcomes. The estimator is dependent on a hyperparameter λ_L. Liu et al. (2024) compare these models and provide a practitioner guide to choose among these models, and we follow that.[14] Results suggest that MC is superior to IFEct based on Mean Squared Prediction Error (MSPE), and hence we choose MC and obtain the optimal tuning parameter λ_L using cross-validation and report all results based on MC.

4. Results

4.1 Event studies: timing-based data structure

In this section, we present the event study results using a timing-based structure, and we keep only treated units. The year in which growth and savings transitions occur is defined as the event year (0) and is the base year in all these results. The assumption of constant transition effect and absence of time-varying confounders is crucial for the identification. Tables S4, S7, and S8 in the Supplementary Appendix show that these transitions cannot be predicted using the country-specific variables widely used in the literature. Figure 5a gives the γ_j during growth transitions when the dependent variable y_{it} is the savings ratio using Equation 3.[15] The γ_j represents the DIDs (difference between average outcomes of transitioning and non-transitioning countries at time t relative to the difference at time $t = 0$ where

---

[14] We use the Stata package 'FEct'. https://yiqingxu.org/packages/fect/stata/fect_md.html

[15] Here, we do not report the changes in TFP growth during TFP transition, as these are going to increase during the transitions.

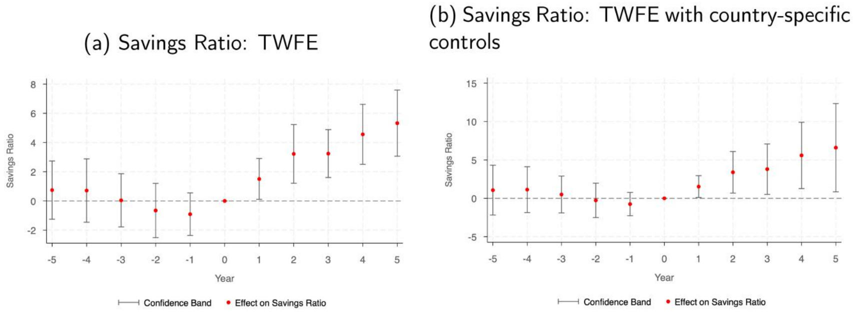
(a) Savings Ratio: TWFE

(b) Savings Ratio: TWFE with country-specific controls

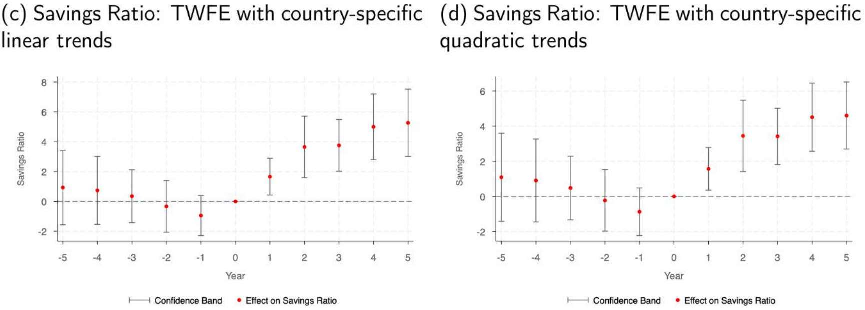
(c) Savings Ratio: TWFE with country-specific linear trends

(d) Savings Ratio: TWFE with country-specific quadratic trends

Figure 5 Change in savings ratio before and after a productivity growth transition.

Notes: (a) is γ_j from $\text{Savings Ratio}_{it} = \sum_{-5}^{5} \gamma_j I_j + \phi_i + \phi_t + e_{it}$. (b) is γ_j from $\text{Savings Ratio}_{it} = \sum_{-5}^{5} \gamma_j I_j + \theta' z_{it} + \phi_i + \phi_t + e_{it}$. (c) is γ_j from $\text{Savings Ratio}_{it} = \sum_{-5}^{5} \gamma_j I_j + \phi_i \times \text{Time Trend} + \phi_i + e_{it}$. (d) is γ_j from $\text{Savings Ratio}_{it} = \sum_{-5}^{5} \gamma_j I_j + \phi_i \times \text{Time Trend} + \phi_i \times \text{Time Trend}^2 + \phi_i + e_{it}$. I_j denotes event windows based on productivity growth transitions. z_{it} includes per-capita income, squared per-capita income, terms of trade, trade-to-GDP ratio, credit as % of GDP, FDI as % of GDP, Young Dependency Ratio (YDR), and Old Dependency Ratio (ODR). The red dot represents the DIDs estimate at time t, calculated as the difference in average outcomes between transitioning and non-transitioning countries at time t, relative to their difference at time $t = 0$, when the transition occurs. If these estimates are statistically insignificant in the periods before the transition ($t < 0$), it implies that the differences in outcomes between the two groups were stable before the transition, which is the same as parallel trends before the transition. Vertical lines represent a 95% confidence band for corresponding point estimates.

the transition occurs) estimates. Statistical insignificance of these estimates before transition at time $t = 0$ implies the difference between transitioning and non-transitioning countries was constant, which is equivalent to parallel trend or absence of pre-trend. Statistical significance of these estimates after transition gives the effect of TFP growth transitions on the savings ratio. As we can see, these TFP growth transitions lead to a sustained increase in the savings ratio, Fig. 5a. Also, we do not find any pre-trend in the savings ratio, and hence we argue that these event studies satisfy the parallel trend assumption. Figure 5b gives the γ_j during growth transitions when the dependent variable y_{it} is the savings ratio using Equation 4. These regressions are better compared to the previous one as they include country-specific time-varying variables, and hence, the time-varying confounding factors are mitigated. Again, we do not find any pre-trend in the data, and the results are very similar to Fig. 5a. This makes sense as these variables are unable to predict the TFP growth transitions.

Figure 5c gives the γ_j during growth transitions when the dependent variable y_{it} is the savings ratio using Equation 5. Equation 5 includes the country-specific time trend to control for time-varying confounders. These models effectively control for a linear trend in the savings

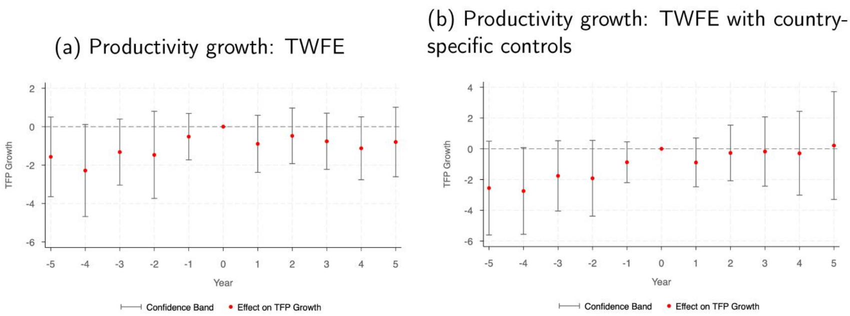

(a) Productivity growth: TWFE

(b) Productivity growth: TWFE with country-specific controls

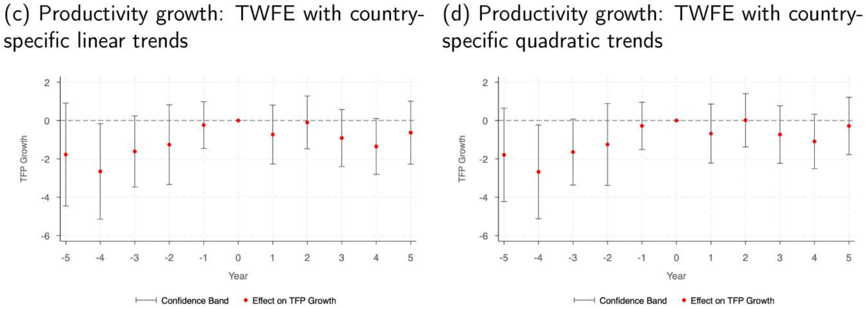

(c) Productivity growth: TWFE with country-specific linear trends

(d) Productivity growth: TWFE with country-specific quadratic trends

Figure 6 Change in productivity growth before and after a savings transition.

Notes: (a) is γ_j from $\text{TFP Growth}_{it} = \sum_{-5}^{5}\gamma_j I_j + \phi_i + \phi_t + e_{it}$. (b) is γ_j from $\text{TFP Growth}_{it} = \sum_{-5}^{5}\gamma_j I_j + \theta' z_{it} + \phi_i + \phi_t + e_{it}$. (c) is γ_j from $\text{TFP Growth}_{it} = \sum_{-5}^{5}\gamma_j I_j + \phi_i \times \text{Time Trend} + \phi_i + e_{it}$. (d) is γ_j from $\text{TFP Growth}_{it} = \sum_{-5}^{5}\gamma_j I_j + \phi_i \times \text{Time Trend} + \phi_i \times \text{Trend}^2 + \phi_i + e_{it}$. I_j denotes event windows based on productivity growth transitions. z_{it} includes per-capita income, squared per-capita income, terms of trade, trade-to-GDP ratio, credit as % of GDP, FDI as % of GDP, Young Dependency Ratio, and Old Dependency Ratio. The red dot represents the DIDs estimate at time *t*, calculated as the difference in average outcomes between transitioning and non-transitioning countries at time *t*, relative to their difference at time $t = 0$, when the transition occurs. If these estimates are statistically insignificant in the periods before the transition ($t < 0$), it implies that the differences in outcomes between the two groups were stable before the transition, which is the same as parallel trends before the transition. Vertical lines represent a 95% confidence band for corresponding point estimates.

ratio in each country. These linear trend models overcome the limitations of missing country-specific variables, which we impute using country-specific median values of the variable. Even with these country-specific linear trends, we find similar results to before. Figure 5d gives the γ_j during growth transitions when the dependent variable y_{it} is the savings ratio using Equation 6. Equation 6 includes the country-specific quadratic time trend to control for time-varying confounders. These models effectively control for a quadratic trend in the savings ratio in each country. These are relevant as the patterns of many low-growth and savings ratio countries show a parabolic pattern during the transitions, Fig. 4b, and a quadratic country-specific unobserved factor is the biggest threat to identification. Even with these country-specific quadratic trends, we find similar results to those before.

Figure 6a gives the γ_j during the savings transition when the dependent variable y_{it} is TFP growth using Equation 3. As we can see during savings transitions, there is no change in TFP growth. We observe slight deviations from the pre-trend, but in one period only. Figure 6b gives the γ_j during the savings transition when the dependent variable y_{it} is TFP growth using Equation 4. These results are also qualitatively similar to the results reported in Fig. 6a. Figure 6c gives the γ_j during the savings transition when the dependent variable y_{it} is TFP growth using Equation 5. Figure 6d gives the γ_j during the savings transition when the dependent variable y_{it} is TFP growth using Equation 6. These results are also qualitatively similar. Although these results suggest some deviation from the absence of pre-trend, all these results suggest the negligible effect of savings transitions on TFP growth.

Results presented in this section, based on a timing-based data structure, give overwhelming evidence that the causation is from TFP growth to savings and not the other way around. We also find strong evidence of a parallel trend before the transition, especially for the TFP growth transition. Parallel trend before transition suggests that these results are not driven by time-varying confounders. But it is difficult to claim these results as causal. As mentioned before, the parallel trend is conclusive evidence only in a two-country, two-period setting. In the context of this article, the parallel trend is not conclusive because it is obtained assuming a constant transition effect, and time-varying confounders need not have a trend. Technically, there are concerns related to staggered treatment, heterogeneous treatment, and time-varying confounders (Goodman-Bacon 2021; Borusyak et al. 2022; Borusyak et al. 2024) that persist. In the next sections, we do further estimates to mitigate these concerns.

4.2 Counterfactual estimator: hybrid data structure

The event studies in the previous section give overwhelming evidence that TFP growth transitions lead to sustained growth in the savings ratio, whereas the savings transition does not lead to any significant increase in TFP growth. As mentioned before, assumption of constant transition effect and time-varying confounders is still a threat to the identification. We first estimate a FEct estimator that relaxes the constant transition effect, and these results are given in the Supplementary Appendix (Figures S12 and S13). Although this does not address the time-varying confounders concerns, but makes the parallel trend more credible evidence for the absence of time-varying confounders as explained before. These results also satisfy the parallel trend and give qualitatively similar results as reported in Figs 5 and 6. Then we estimate the model given by Equations 8 and 9, which mitigates the concerns related to these time-varying confounders explicitly and allows heterogeneous transition effects. Following Liu et al. (2024), we compare the estimates from Equations 8 and 9, and choose MC based on mean squared prediction error (MSPE). Hence, in this section, we present results from Equation 9. We choose the hyperparameter for MC using cross-validation.

Figure 7 gives the effect on savings during the TFP growth transitions, and we find a persistent effect on the savings ratio up to 10 years after the transition. These results do not show any deviation from a parallel trend up to 5 years before the transition. Figure 8 gives the effect on TFP growth during the savings transitions, and we find qualitatively similar results as before, although the TFP growth turns out to be significant for a few years, but these are not persistent. We do not find any deviation from the parallel trend up to 5 years before the transition in Fig. 8. Also, the effect of savings transitions on TFP growth is not significant if we use another savings transition, as shown in S9 of the Supplementary Appendix. But the effect of TFP growth on savings transition remains very similar to the ones reported here (Supplementary Appendix 8), and hence we argue that the effect of TFP growth on savings is causal, as these results have very minimal concerns related to time-varying confounders for two reasons. First, the MC method explicitly addresses the time-varying confounders by including the country-specific time-varying factors in the estimation. Second, since the MC method considers heterogeneous treatment effects, the absence of pre-trend in the data is credible evidence of the absence of time-varying confounders.

Finally, we do a Placebo/falsification test to establish the validity of our empirical design. The results explained before suggest the sustained effect of TFP growth transitions on the

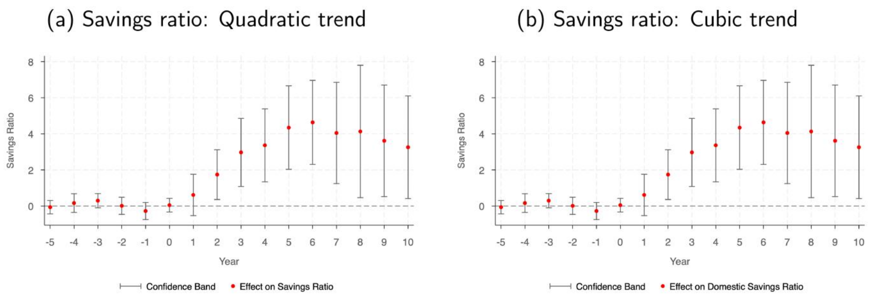
(a) Savings ratio: Quadratic trend

(b) Savings ratio: Cubic trend

Figure 7 Change in savings before and after a productivity growth transition.

Notes: MC method. We include per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, credit as % of GDP, FDI as % of GDP, Young Dependency Ratio, and Old Dependency Ratio. (a) is with a quadratic trend, and (b) is with a cubic trend. The red dot represents the DID estimate at time t, calculated as the difference in average outcomes between transitioning and non-transitioning countries at time t, relative to their difference at time $t = 0$, when the transition occurs. If these estimates are statistically insignificant in the periods before the transition ($t < 0$), it implies that the differences in outcomes between the two groups were stable before the transition, which is the same as parallel trends before the transition. Vertical lines represent a 95% confidence band for corresponding point estimates.

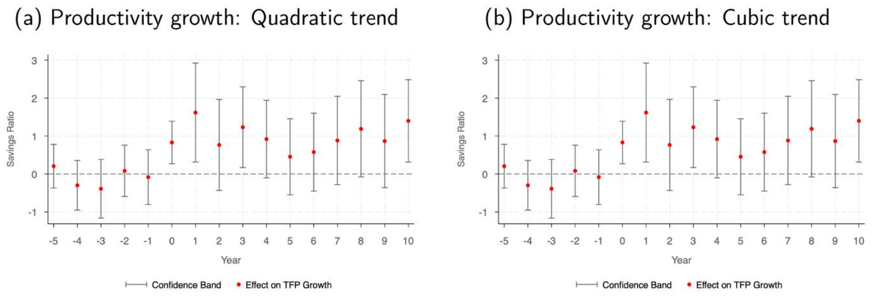
(a) Productivity growth: Quadratic trend

(b) Productivity growth: Cubic trend

Figure 8 Change in productivity growth after savings transitions.

Notes: MC method. We include per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, credit as % of GDP, FDI as % of GDP, Young Dependency Ratio, and Old Dependency Ratio. (a) is with a quadratic trend, and (b) is with a cubic trend. The red dot represents the DID estimate at time t, calculated as the difference in average outcomes between transitioning and non-transitioning countries at time t, relative to their difference at time $t = 0$, when the transition occurs. If these estimates are statistically insignificant in the periods before the transition ($t < 0$), it implies that the differences in outcomes between the two groups were stable before the transition, which is the same as parallel trends before the transition. Vertical lines represent a 95% confidence band for corresponding point estimates.

savings ratio. We know that these transitions occur at $t = 0$ by design. If these are indeed the effects of transition, then if we implement a false TFP growth transition, we should not find any effect on the savings ratio. This can also be considered as a falsification test. There are several ways of estimating the models with false transitions. We estimate a simple model by assigning transitions to three periods before the actual TFP growth transition. These also help in ruling out the presence of a pre-trend. This is because if there is a pre-trend, then assigning the transition early will give rise to a statistically significant effect on the savings ratio. As we can see from Fig. 9, assigning transition three periods before the actual TFP growth transition does not lead to any significant change in the savings ratio after that false TFP growth transition. Hence, we conclude that the empirical design passes the necessary

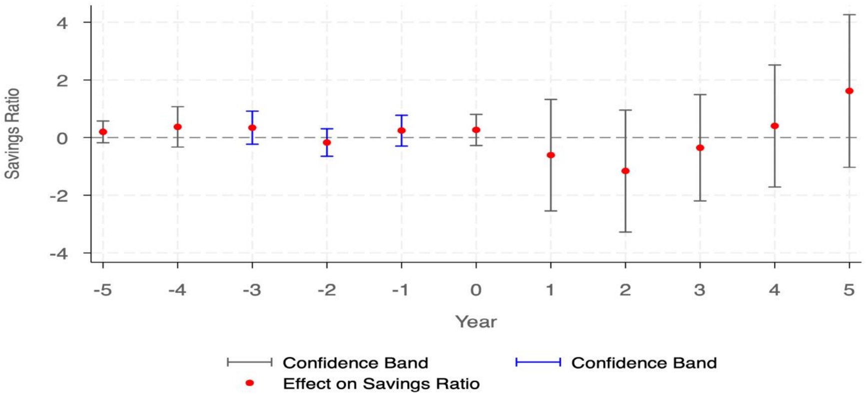

Figure 9 Placebo; change in savings before and after a productivity growth transition.

Notes: MC method. We include per-capita income, squared per-capita income, terms of trade, trade-to-GDP ratio, credit as % of GDP, FDI as % of GDP, Young Dependency Ratio, and Old Dependency Ratio. We include a cubic trend. We assign treatment three periods before the actual TFP growth transition, and vertical blue lines reflect that. The red dot represents the DID estimate at time *t*, calculated as the difference in average outcomes between transitioning and non-transitioning countries at time *t*, relative to their difference at time $t = 0$, when the transition occurs. If these estimates are statistically insignificant in the periods before the transition ($t < 0$), it implies that the differences in outcomes between the two groups were stable before the transition, which is the same as parallel trends before the transition. Vertical lines represent a 95% confidence band for corresponding point estimates.

falsification test, and the parallel trend reported in Fig. 9 is credible evidence of the absence of time-varying confounders.

4.2.1 Robustness

We do various robustness exercises to validate these results. As mentioned before, we identify alternative savings and TFP growth transitions. We also identify two sets of TFP growth transitions using country-specific cut-offs. We estimate the effect of savings ratio and TFP growth during these alternative TFP growth and savings transitions using MC. Results related to TFP growth transitions II, III, IV, and V are given in S8 of the Supplementary Appendix. All these results obtained using different combinations of TFP growth transitions are similar to the results reported here and give confidence that TFP growth transitions indeed lead to a persistent rise in savings. Results related to savings Transition II and III are given in S9 of the Supplementary Appendix. We do not find any evidence of the sustained effect of savings transition on TFP growth due to Savings Transition II and III. Hence, we argue that this evidence suggests that TFP growth transitions lead to persistent increases in savings, whereas savings transitions fail to bring such an increase in TFP growth.

4.3 Dynamic response of savings due to TFP growth and productivity shocks

The results presented above suggest that the impact of TFP growth on savings is persistent, whereas the savings transitions do not lead to a persistent increase in TFP growth. But this does not give us an easy-to-interpret effect of the TFP growth transition on the savings ratio. This is because the countries experiencing TFP growth transitions have different changes in TFP growth rates during the transition. It is impossible to identify many countries experiencing transition by imposing a uniform increase in the TFP growth rate during the transition. Hence, we estimate a local projection regression following Jorda (2005, 2023) given by Equations 10 and 11.

$$S_{i,t+j} = \beta_1^j f_{it} + \theta'^j z_{it} + \phi_i^j + \phi_t^j + \varepsilon_{it}^j \text{for} j = 0, 1, \ldots 8 \tag{10}$$

$$S_{i,t+j} = \beta_1^j f_{it} + \theta'^j z_{it} + \phi_i^j + \phi_i^j \times \text{Time Trend} + \varepsilon_{it}^j \text{for} j = 0, 1, \ldots 8 \tag{11}$$

Where f_{it} is either TFP growth or productivity shocks obtained from the data (section "TFP growth from data") and model (section "TFP from neoclassical growth model"), respectively. To ensure comparability across these measures, we standardize different measures of f_{it}. z_{it} includes per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, foreign direct investment as % of GDP, credit as % of GDP, young dependency ratio, old dependency ratio, and four lags of TFP growth and savings ratio. The above model identifies innovation in TFP growth conditions on four lags of TFP growth, savings, and a large number of control variables. In this way, the innovation of TFP growth and productivity shocks used in Equations 10 and 11 are even better than the ones we can obtain using panel vector autoregression with limited number of variables and Cholesky decomposition, as we condition on many variables apart from the lag of the productivity growth and savings, including country-specific linear trend to capture time-varying unobserved confounders.[16] Without establishing the direction of causality, it would be hard to argue for exogeneity of TFP growth innovations used in Equations 10 and 11 to obtain the response of savings due to the potential simultaneity between these two variables.

Figure 10a shows the impact on the saving ratio due to a one-standard-deviation increase in TFP growth with year-fixed effects. One standard deviation increase in these two measures increases the savings ratio by 0.4 percentage points, and this is very persistent. These responses are significant for up to 6 years. The model with country-specific time trends helps in controlling for confounders that vary across countries and time, which is a threat to identification. With a linear trend, the impact on savings is slightly higher but less persistent, and the effect is only significant for up to 3 years, Fig. 10b. Figures 11a and 11b plots the impact on the savings ratio due to a one standard deviation increase in these two measures of productivity shocks. These shocks have a higher impact on savings and are significant for up to 6 years, but are less persistent compared to the response of the savings ratio due to the TFP growth. With the inclusion of a country specific linear trend, the effect is significant for only up to 3 years, which is similar to the response of savings due to TFP growth with country-specific linear trends (Figs 12a and 12b).

Many of the control variables used in the above local projection regressions have missing values, which we impute by the country-specific median of that variable. To ensure that the results are not driven by this imputation, we estimate the local projection regressions without these control variables. These results are given in Figures S30–S31 in the Supplementary Appendix. These are almost identical to the responses presented here, and hence, that gives us further confidence that these responses are not driven by imputed variables or omitted variables. This is because these are widely used variables in the savings literature, and if these are not significantly influencing the response of the savings ratio due to TFP growth, then it is hard to argue that other variables would do so. Figures S32–S34 in the

---

[16] The use of the panel VAR with Cholesky decomposition also requires some understanding related to the exogeneity of variables to decide the identification, or else we require instrumental variables to identify the exogeneity. Results in the previous section help us in understanding the direction of causality, and hence, we estimate the effect of TFP growth on savings using local projection.

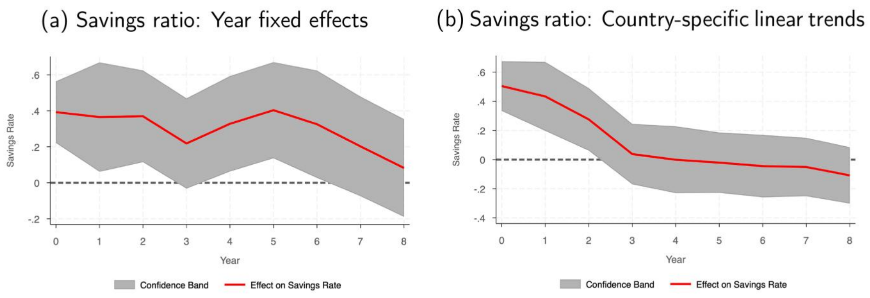

Figure 10 Response of the saving ratio due to productivity growth.

Notes: (a) β_1^j from $S_{i,t+j} = \beta_1^j TFP_{it} + \theta'^j z_{it} + \phi_i^j + \phi_t^j + \varepsilon_{it}^j \text{for} j = 0, 1, \ldots 8$. (b) is β_1^j from $S_{i,t+j} = \beta_1^j TFP_{it} + \theta'^j z_{it} + \phi_i^j + \phi_i^j \times \text{Time Trend} + \varepsilon_{it}^j \text{for} j = 0, 1, \ldots 8$. z_{it} includes per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, foreign direct investment as % of GDP, credit as % of GDP, young dependency ratio, old dependency ratio, and four lags of TFP growth and savings ratio. The shaded area represents a 95% confidence band for corresponding point estimates.

Source: authors' construction.

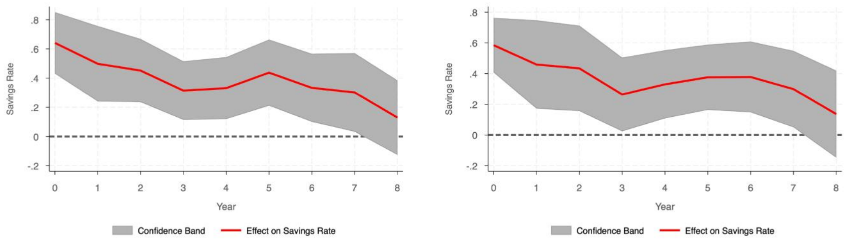

(a) Productivity Shock: Linearly Detrended Model

(b) Productivity Shock: Detrended With Hamilton Filter

Figure 11 Response of saving ratio due to productivity shocks.

Notes: β_1^j from $S_{i,t+j} = \beta_1^j TFPShock_{it} + \theta'^j z_{it} + \phi_i^j + \phi_t^j + \varepsilon_{it}^j \text{for} j = 0, 1, \ldots 8$. z_{it} includes per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, foreign direct investment as % of GDP, credit as % of GDP, young dependency ratio, old dependency ratio, and four lags of productivity shocks and savings ratio. The shaded area represents a 95% confidence band for corresponding point estimates.

Supplementary Appendix provide the response of domestic savings due to one standard deviation innovation in these two measures of TFP growth and productivity shocks. These responses are very similar to the gross savings response reported here. As mentioned before, in the absence of capacity utilization, TFP growth could be biased. Hence, we estimate the response of savings due to capacity utilization-adjusted TFP growth for a small set of countries for which the data is made available by Comin et al. (2020). We find that even the capacity utilization-adjusted TFP growth leads to an increase in the savings rate. This is provided in Figure S35 in the Supplementary Appendix. Moreover, the savings response due to the TFP growth, adjusted and unadjusted for capacity utilization, is statistically indistinguishable. Hence, we believe that the results reported in this section are robust and unlikely to be driven by the biases in TFP growth due to unobserved capacity utilization for a large set of countries.

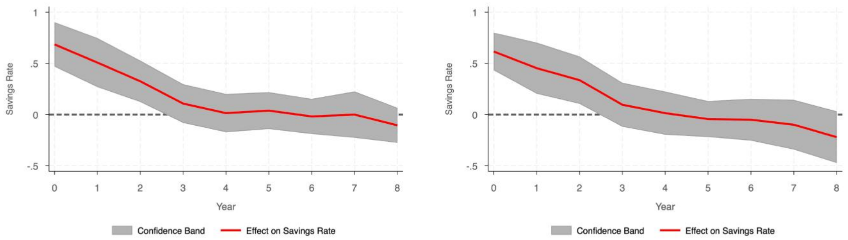
(a) Productivity Shock: Linearly Detrended Model

(b) Productivity Shock: Detrended With Hamilton Filter

Figure 12 Response of saving ratio due to productivity shocks: with country specific linear trend.

Notes: β_1^j from $S_{i,t+j} = \beta_1^j TFPShock_{it} + \theta'^j z_{it} + \phi_i^j + \phi_i^j \times \text{Time Trend} + \varepsilon_{it}^j$ for $j = 0, 1, \ldots 8$. z_{it} includes per-capita income, squared per-capita income, terms of trade, trade to GDP ratio, foreign direct investment as % of GDP, credit as % of GDP, young dependency ratio, old dependency ratio, and four lags of productivity shocks and savings ratio. The shaded area represents a 95% confidence band for corresponding point estimates.

4.3.1 Robustness

Although the results in the previous section provide overwhelming evidence of the persistent impact of TFP growth on savings, concerns related to omitted variables can only be dealt with in a comprehensive panel VAR model. Given data limitations and the absence of instrumental variables, we try our best to mitigate these concerns and establish the credibility of results via estimating Panel VARs using different endogenous variables and lag lengths, and estimate the response of savings due to TFP growth innovations using Cholesky decomposition. In our most comprehensive model, we include seven other endogenous variables identified in the growth-savings literature apart from TFP growth (see Modigliani 1970; Loayza et al. 2000). We order TFP growth first in all these panel VARs, given the credible evidence in Sections S4.1 and S4.2 about the direction of causation. We estimate five models with different sets of endogenous variables, exogenous variables, and lag lengths. In all of these models, we find a significant effect of TFP growth on the savings ratio. The response of the savings ratio due to TFP growth shock in all these estimates is qualitatively similar to the responses presented here using local projection, and the most comprehensive model with eight endogenous variables suggests that a one standard deviation shock to TFP growth leads to 0.5 percentage point increase in the savings rate. These results are given in Section S13 of the Supplementary Appendix. We also perform a counterfactual simulation for the savings ratio during the TFP transition for Cameroon and Malaysia using the model explained above, and these results are given in Section S14 of the Supplementary Appendix.

5. Concluding remarks and policy recommendation

Increasing per capita income growth in poor countries has always been a concern for policymakers around the world. The literature on this issue is broadly divided into two categories. The first strand of the literature suggests that increasing the savings ratio is essential for increasing per capita income growth. In other words, this literature believes in causation running from saving to growth. The second strand of literature suggests that growth causes savings, and growth can be achieved via simple reforms and allocating resources more wisely. These two strands of literature have not been able to reconcile the issues and reach a conclusion, which is of utmost importance, as shown in this article.

Using aggregate data from WDI and PWT, we provide evidence using different approaches that it is productivity growth that causes savings and not the other way around. We find savings and growth transitions in the data that lead to sudden and sustained increases in the savings ratio and TFP growth. We show that these transitions are exogenous and implement event studies and counterfactual estimators on savings and productivity growth transitions. Our results suggest that savings transitions, that is, increases in savings, do not change productivity growth, whereas productivity growth transitions, that is, an increase in productivity growth, significantly increases savings, which remain at an elevated level even after 10 years. Having established the direction of causality, we estimate the local projection regressions and find that a one standard deviation increase in TFP growth rate leads to 0.5 percentage point increase in the saving ratio, and these are significant for up to 6 years.

Then we obtain technological shocks from a neoclassical growth model by estimating the model for forty-nine countries (countries with at least 40 years of observations). We then use these productivity shocks instead of TFP growth in local projection regressions, and they give a similar effect on savings as the TFP growth. Finally, we do a counterfactual simulation that suggests that a large decline in productivity growth is associated with a large decline in the savings ratio and vice versa. Evidence provided in this article suggests that the saving–growth causation runs from growth to savings and not the other way around. Policymakers should design policies that can increase productivity growth rather than look for ways to increase savings, which do not seem to be effective enough to boost savings.

Acknowledgements

The authors are grateful to the editor, Francesco Zanetti, and three anonymous reviewers of this journal for their constructive comments and suggestions on earlier versions of this article. The authors are grateful to the participants at several conferences, including CUFE-OEP conference at the School of Innovation and Development, Central University of Finance and Economics, Beijing (November 2025); 1st International Conference on Applied Economics and Financial Issues, Aristotle University of Thessaloniki (August 2025); Qingdao University (April 2025); Conference on Enabling Sustained Growth in Emerging Asia, University of Hyderabad (February 2025); the EEFS annual conference (June 2021); Trade and Development Conference in memory of Helena Marques (April 2021); Indian Statistical Institute, Delhi (February 2021). Also, comments received from Harald Uhlig, Parantap Basu, Knut Arild Thonstad, Abhiroop Mukhopadhyay, and Monisankar Bishnu are gratefully acknowledged. The usual disclaimer applies.

Supplementary material

Supplementary material is available at *Oxford Economic Papers Journal* online. The folder contains an online appendix, data, and codes that replicate the results in the paper and online appendix.

Conflicts of interest

The authors declare that they have no conflicts of interest.

Funding

An early version of this study was conducted as part of the Domestic Revenue Mobilization programme of UNU-WIDER in Helsinki.

References

Aghion, P., Comin, D., Howitt, P., and Tecu, I. (2016) 'When Does Domestic Savings Matter for Economic Growth?', *IMF Economic Review*, 64: 381–407. https://doi.org/10.1057/imfer.2015.41

Aguiar, M., and Gopinath, G. (2007) 'Emerging Market Business Cycles: The Cycle Is The Trend', *Journal of Political Economy*, 115: 69–102. https://doi.org/10.1086/511283

Aizenman, J., and Noy, I. (2015) 'Saving and the Long Shadow of Macroeconomic Shocks', *Journal of Macroeconomics*, 46: 147–59. https://doi.org/10.1016/j.jmacro.2015.08.007

Athey, S., Bayati, M., Doudchenko, N., Imbens, G., and Khosravi, K. (2021) 'Matrix Completion Methods for Causal Panel Data Models', *Journal of the American Statistical Association*, 116: 1716–30. https://doi.org/10.1080/01621459.2021.1891924

Barro, R., and Lee, J. (2013) 'A New Data Set of Educational Attainment in the World, 1950-2010', *Journal of Development Economics*, 104: 184–98. https://doi.org/10.1016/j.jdeveco.2012.10.001

Basu, S., and Fernald, J. G. (2002) 'Aggregate Productivity and Aggregate Technology', *European Economic Review*, 46: 963–91. https://doi.org/10.1016/S0014-2921(02)00161-7

Basu, S., Fernald, J. G., and Kimball, M. S. (2006) 'Are Technology Improvements Contractionary?', *American Economic Review*, 96: 1418–48. https://doi.org/10.1257/aer.96.5.1418

Borusyak, K., Hull, P., and Jaravel, X. (2022) 'Quasi-Experimental Shift-Share Research Designs', *The Review of Economic Studies*, 89: 181–213. https://doi.org/10.1093/restud/rdab030

Borusyak, K., Jaravel, X., and Spiess, J. (2024) 'Revisiting Event-Study Designs: Robust and Efficient Estimation', *Review of Economic Studies*, 91: 3253–85. https://doi.org/10.1093/restud/rdae007

Brinca, P., Chari, V. V., Kehoe, P. J., and McGrattan, E. (2016) 'Accounting for Business Cycles', in J. B. Taylor and H. Uhlig (eds) *Handbook of Macroeconomics*, volume 2, pp. 1013–1063. Amsterdam: Elsevier.

Buera, F. J., and Shin, Y. (2017) 'Productivity Growth and Capital Flows: The Dynamics of Reforms', *American Economic Journal: Macroeconomics*, 9: 147–85. https://doi.org/10.1257/mac.20160307

Burnside, A. C., Eichenbaum, M. S., and Rebelo, S. T. (1996) 'Sectoral Solow Residuals', *European Economic Review*, 40: 861–9. https://doi.org/10.1016/0014-2921(95)00095-X

Burnside, C., and Dollar, D. (2000) 'Aid, Policies, and Growth', *American Economic Review*, 90: 847–68. https://doi.org/10.1257/aer.90.4.847

Burnside, C., and Dollar, D. (2004) 'Aid, Policies, and Growth: Reply', *American Economic Review*, 94: 781–4. https://doi.org/10.1257/0002828041464524

Carroll, C. D. (1994) 'How Does Future Income Affect Current Consumption?', *Quarterly Journal of Economics*, 109: 111–47. https://doi.org/10.2307/2118430

Carroll, C. D., and Weil, D. N. (1994) 'Saving and Growth: A Reinterpretation', *Carnegie-Rochester Conference Series on Public Policy*, 40: 133–92. https://doi.org/10.1016/0167-2231(94)90006-X

Chari, V. V., Kehoe, P. J., and McGrattan, E. R. (2007) 'Business Cycle Accounting', *Econometrica*, 75: 781–836. https://doi.org/10.1111/j.1468-0262.2007.00768.x

Chen, P., Karabarbounis, L., and Neiman, B. (2017) 'The Global Rise of Corporate Saving', *Journal of Monetary Economics*, 89: 1–19. https://doi.org/10.1016/j.jmoneco.2017.03.004

Comin, D., Quintana, J., Schmitz, T., and Trigari, A. (2020) 'A New Measure of Utilization-Adjusted TFP Growth for Europe and the United States', *NBER Working Paper* w28008. https://www.nber.org/papers/w28008

David, J. M., Hopenhayn, H. A., and Venkateswaran, V. (2016) 'Information, Misallocation, and Aggregate Productivity', *The Quarterly Journal of Economics*, 131: 943–1005. https://doi.org/10.1093/qje/qjw006

Easterly, W. (2003) 'Can Foreign Aid Buy Growth?', *Journal of Economic Perspectives*, 17: 23–48. https://doi.org/10.1257/089533003769204344

Easterly, W., Levine, R., and Roodman, D. (2003) 'New Data, New Doubts: A Comment on Burnside and Dollar's "Aid, Policies, and Growth" (2000)', Working Paper 9846. Cambridge, MA: NBER. https://doi.org/10.3386/w9846

Erosa, A., and Hidalgo Cabrillana, A. (2008) 'On Finance as a Theory of TFP, Cross-Industry Productivity Differences, and Economic Rents', *International Economic Review*, 49: 437–73. https://doi.org/10.1111/j.1468-2354.2008.00486.x

Feldstein, M., and Horioka, C. (1979) 'Domestic Savings and International Capital Flows', Working Paper 0310. Cambridge, MA: NBER. https://doi.org/10.3386/w0310

Fernald, J. 2014 *A Quarterly, Utilization-Adjusted Series on Total Factor Productivity*. San Francisco, CA: Federal Reserve Bank of San Francisco.

Fernandez, A., Imrohoroglu, A., and Tamayo, C. E. (2019) 'Saving Rates in Latin America: A Neoclassical Perspective', *IMF Economic Review*, 67: 791–823. https://doi.org/10.1057/s41308-019-00093-1

Gali, J. (1999) 'Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?', *American Economic Review*, 89: 249–71. https://doi.org/10.1257/aer.89.1.249

Gali, J., Lopez-Salido, J. D., and Valles, J. (2003) 'Technology Shocks and Monetary Policy: Assessing the Fed's Performance', *Journal of Monetary Economics*, 50: 723–43.

Gertler, M., and Kiyotaki, N. (2009) 'Financial Intermediation and Monetary Policy in the Business Cycle', *American Economic Review*, 92: 739–64.

Gobillon, L., and Magnac, T. (2016) 'Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls', *Review of Economics and Statistics*, 98: 535–51. https://doi.org/10.1162/REST_a_00537

Goodman-Bacon, A. (2021) 'Difference-in-Differences with Variation in Treatment Timing', *Journal of Econometrics*, 225: 254–77. https://doi.org/10.1016/j.jeconom.2021.03.014

Gortz, C., Tsoukalas, J. D., and Zanetti, F. (2022) 'News Shocks under Financial Frictions', *American Economic Journal: Macroeconomics*, 14: 210–43. https://doi.org/10.1257/mac.20170066

Greenwood, J., Hercowitz, Z., and Krusell, P. (1997) 'Long-Run Implications of Investment-Specific Technological Change', *American Economic Review*, 87: 342–62.

Grossman, G. M., and Helpman, E. (1991) *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.

Hall, R. E., and Jones, C. I. (1999) 'Why Do Some Countries Produce So Much More Output Per Worker than Others?', *Quarterly Journal of Economics*, 114: 83–116. https://doi.org/10.1162/003355399555954

Hamilton, J. D. (2018) 'Why You Should Never Use the Hodrick-Prescott Filter', *Review of Economics and Statistics*, 100: 831–43. https://doi.org/10.1162/rest_a_00706

Hsieh, C. T., and Klenow, P. J. (2009) ‘Misallocation and Manufacturing TFP in China and India’, *The Quarterly Journal of Economics*, 124: 1403–48. https://doi.org/10.1162/qjec.2009.124.4.1403

Ireland, P.N. (2011) ‘A New Keynesian Perspective on the Great Recession’, *Journal of Money, Credit and Banking*, 43: 31–54. https://doi.org/10.1111/j.1538-4616.2010.00364.x

Jacobson, L.S., LaLonde, R.J., and Sullivan, D.G. (1993) ‘Earnings Losses of Displaced Workers’, *The American Economic Review*, 83: 685–709.

Jorda, O. (2005) ‘Estimation and Inference of Impulse Responses by Local Projections’, *American Economic Review*, 95: 161–82. https://doi.org/10.1257/0002828053828518

Jorda, O. (2023) ‘Local Projections for Applied Economics’, *Annual Review of Economics*, 15: 607–31.

Kiyotaki, N., and Moore, J. (1997) ‘Credit cycles’, *Journal of Political Economy*, 105: 211–48. https://doi.org/10.1086/262072

Klenow, P. J., and Rodriguez-Clare, A. (1997) ‘The Neoclassical Revival in Growth Economics: Has it Gone Too Far?’, *NBER Macroeconomics Annual*, 12: 73–103. https://doi.org/10.1086/654324

Lewis, W.A. (1954) ‘Economic Development with Unlimited Supplies of Labor’, *Manchester School of Economic and Social Studies*, 22: 139–91. https://doi.org/10.1111/j.1467-9957.1954.tb00021.x

Liu, L., Wang, Y., and Xu, Y. (2024) ‘A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data’, *American Journal of Political Science*, 68: 160–76. https://doi.org/10.1111/ajps.12723

Loayza, N., Schmidt-Hebbel., , K., and Serven, L. (2000) ‘What Drives Private Saving Across the World?’, *The Review of Economics and Statistics*, 82: 165–81.

Lovenheim, M.F., and Willén, A. (2019) ‘The long-run effects of teacher collective bargaining’, *American Economic Journal: Economic Policy*, 11: 292–324. https://doi.org/10.1257/pol.20170570

Miller, D.L. (2023) ‘An introductory guide to event study models’, *Journal of Economic Perspectives*, 37: 203–30. https://doi.org/10.1257/jep.37.2.203

Modigliani, F. (1970) ‘The Life Cycle Hypothesis of Saving and Intercountry Differences in the Saving Ratio’, in *Induction, Growth and Trade*. Oxford: Clarendon Press.

Mody, A., Ohnsorge, F., and Sandri, D. (2012) ‘Precautionary Savings in the Great Recession’, *IMF Economic Review*, 60: 114–38. https://doi.org/10.1057/imfer.2012.5

Prasad, E., Rajan R. G., and Subramanian A. (2006) ‘Patterns of International Capital Flows and their Implications for Economic Development’, in *Proceedings of the 2006 Jackson Hole Symposium*, pp. 119–158. Kansas, MO: Federal Reserve Bank of Kansas City.

Solow, R. M. (1957) ‘Technical Change and the Aggregate Production Function’, *The Review of Economics and Statistics*, 39: 312–20. https://doi.org/10.2307/1926047

Williamson, J. G. (1979) ‘Why Do Koreans Save So Little?’, *Journal of Development Economics*, 6: 343–62. https://doi.org/10.1016/0304-3878(79)90021-X

Xu, Y. (2017) ‘Generalized Synthetic Control Method: Causal Inference With Interactive Fixed Effects Models’, *Political Analysis*, 25: 57–76. https://doi.org/10.1017/pan.2016.2

Young, A. (1994) ‘Lessons from the East Asian NICs: A Contrarian View’, *European Economic Review*, 38: 964–73.

Zabot, U. C., and Gomes, F. A. R. (2025) ‘Do the Rich Save More? Evidence from Brazil’, *Oxford Economic Papers*, 77: 1080–105.

Oxford Economic Papers, 2026, 00, 1–29
https://doi.org/10.1093/oep/gpag006
Article