The Neoclassical Growth of China

Jesús Fernández-Villaverde
University of Pennsylvania, NBER, and CEPR

Lee Ohanian
University of California, Los Angeles

Wen Yao
Tsinghua University*

September 9, 2023

Abstract

This paper studies China's four-fold increase in per capita GDP relative to the U.S. between 1995 and 2019. We first show that China's growth pattern is very similar to that of several other East Asia economies that initially grew very quickly. We find that a minimalist Ramsey-Cass-Koopmans model, endowed with a parsimonious TFP catch-up process, captures the TFP paths of all these economies very closely and accounts for China's real GDP growth and the growth paths of the other East Asia economies. We use the model to address China's future growth prospects and find that China's growth will likely slow substantially in the future, to the point where the U.S. growth rate will likely exceed China's in 20 years. We also find that China's income per capita will converge to roughly 44% of the U.S. level around 2100 and that the size of the overall U.S. economy will ultimately overtake China's economy, reflecting faster projected U.S. population growth and much higher projected U.S. productivity.

Keywords: China, East Asia, economic growth, Ramsey-Cass-Koopmans model, TFP catch-up.

JEL codes: E10, E20, O4

*jesusfv@upenn.edu, ohanian@econ.ucla.edu, yaow@sem.tsinghua.edu.cn. We thank Kaiji Chen, Kjetil Storesletten, Zheng Song, Fabrizio Zilibotti, Tao Zha, and participants at many seminars for useful comments.

1 Introduction

China's real per capita GDP surged from 6.6% of the U.S. per capita GDP level in 1995 (when China became a middle-income country) to 25% in 2019 (the last year of data in Penn World Tables 10.0, PWT 10.0), representing a nearly four-fold catch-up to the U.S. in just 25 years.[1] China's economic growth has been one of the most transformative events in economic history, as China's economy is now the largest in the world in PPP terms.

Understanding China's growth is an active research topic, and the future growth of China is an important open question. This paper conducts a neoclassical analysis of China's economic growth between 1995-2019 and applies it to evaluate China's future growth prospects. Our first main finding is that China's real GDP catch-up is remarkably similar to those of several other East Asia countries that grew rapidly. Our second main finding is that a minimalist Ramsey-Cass-Koopmans model, endowed with a parsimonious TFP catch-up process, very accurately accounts for China's real GDP growth and the growth paths of the other East Asia economies. Our third main finding is that projecting the TFP catch-up process into the future forecasts that China's growth will slow substantially, with the U.S. growth rate overtaking China's in 20 years.

We begin by comparing China's growth experience to the experiences of other East Asian growth miracles (Korea, Taiwan, Japan, Hong Kong, and Singapore).[2] We document that China's catch-up growth pattern is remarkably similar to that of the other growth miracles at the same stage of development. All had rapid initial growth, followed by a significant decline in their growth rate. This similarity leads us to develop a minimalist Ramsey-Cass-Koopmans one-sector growth model augmented with a parsimonious TFP catch-up process governed by three parameters that specify: (i) the initial TFP level of an economy, (ii) the asymptotic catch-up TFP bound relative to the U.S., and (iii) the speed of catch-up. We calibrate standard model parameters to conventional aggregate targets and estimate the parameters of the TFP catch-up process to fit the observed patterns of TFP growth. While we are agnostic about the specific mechanisms that may be driving TFP catch-up in China or the other growth miracles, our approach provides a simple and transparent evaluation of the growth process in these economies.

We find that the model matches China's income per capita growth from 1995 to 2019 *surprisingly* well. We emphasize *surprisingly* because the model is so simple; it does not include factors studied within the literature, including sectoral reallocation, financial market developments, industrial policies, foreign direct investment, or trade policies. We do not dismiss the importance of these or other factors. Rather, this analysis suggests that many of the important effects of these and other factors influencing China's economic growth are manifested within a

[1] Unless otherwise noted, all per capita measures are from the PWT 10.0 and adjusted for purchasing power parity (PPP). We set the threshold for classifying a country as a middle-income nation at $3,000 in 2017 PPP-adjusted dollars, a commonly used benchmark.

[2] To conserve space, we refer to the Republic of Korea; China, Taiwan Region; and China, Hong Kong SAR, respectively, as Korea, China, Taiwan, and Hong Kong.

simple productivity catch-up process operating within a standard one-sector framework.

Our model has more difficulty matching China's measured investment rates, but we will make two arguments to reconcile the data with the model. First, many indications suggest that China's official investment rates are mismeasured. When we use the corrected investment series from Chen et al. (2019), our model matches investment rates extremely well. Second, we will present evidence that investment efficiency in China is low, creating a wedge between investment and output growth.

We use the same minimalist model (with some country-specific recalibration) to analyze Japan, Korea, Taiwan, Hong Kong, and Singapore and find that the model also accounts for these growth episodes. This is a key validation for the model, particularly since the analysis for these other growth miracles covers different institutional and policy environments and many more years than the analysis of China. The much longer period of study for these other four growth miracles has important implications for China's future, as TFP growth has continued to decline in all four economies over time after their growth miracle stage.

The common growth experiences of China and the other growth miracles, and the quantitative model's success, lead us to assess China's future growth prospects. We use the model and the fitted TFP catch-up process to predict that China's relative per capita GDP level will asymptote to about 41% of the U.S. level around 2050, reflecting a substantial slowdown in China's observed TFP catch-up in recent years. Moreover, given China's forecasted demographics, which we take from the UN, the model predicts that U.S. output will grow faster (2.29%) than China's (2.27%) by 2043 and that, by 2088, the U.S. economy will be larger than China's in PPP terms.[3] Slowing TFP growth and a declining population suggest a far-from-booming Chinese economic future.

This analysis builds on much previous research. Most importantly, given our neoclassical approach, Barro (2016) has argued that, from the perspective of conditional convergence, China's GDP growth rate since 1990 has been surprisingly high and that, consequently, China's growth is bound to fall back toward historical averages. We show that Barro's findings can be reconciled with neoclassical growth by introducing a process of TFP catch-up. Furthermore, our projection of China's future growth supports Barro's view that China's growth will decline and provides a quantification of China's economic future compared to that of the U.S.

Pritchett and Summers (2013) elaborate on the regression to the mean in terms of growth across many economies. Also, they argue that rapid growth episodes are frequently followed by discontinuous drop-offs in growth. The authors conjecture that the salient characteristics of China (high levels of state control, corruption, and authoritarian rule) make a discontinuous decline in growth relatively likely. We reach a similar conclusion, reflecting a rapid decline in TFP catch-up, which in turn may indeed be related to the institutional aspects emphasized by

[3]For our statement regarding 2043, there is very little demographic uncertainty left. Given current birth and death rates in China, around 85% of people who will be alive in China were already born by the end of 2022.

Pritchett and Summers (2013).

Song et al. (2011) construct a growth model consistent with China's economic transition. The key behind their dynamics is that firms that use more productive technologies must finance investments through internal savings due to financial market imperfections. The most intriguing aspect of their model is that the downsizing of financially integrated firms forces domestic savings to be invested abroad, generating a foreign surplus. Since we do not deal with an open economy, we do not explore this aspect of China's growth experience. Thus, we see our paper as complementing Song et al. (2011). Chen and Zha (2023) highlight the importance of the gradualist approach to China's reforms and compare the path under gradualism with a laissez-faire counterfactual. The authors focus much attention on the trade-offs between active government interventions and long-term growth.

Our emphasis on the low investment efficiency of China is shared by Bai et al. (2020), who focus on the distortions created by loans given by local authorities to unskilled labor-intensive firms after 2008. Similarly, Jiang et al. (2022) analyze the role of local governments using land-sale revenue to fuel infrastructure investment and how such a policy leads to a declining capital return. Xiong (2019) studies the agency problem between the central and local governments, which might lead to over-investment and unreliable economic statistics.

Our paper is also related to studies on China's business cycle fluctuations, such as Chang et al. (2019) and Yao and Zhu (2021). In particular, Chang et al. (2016) explore the policy implications of subsidized investments in capital-intensive industries facing a collateral constraint within a distorted banking system and find that these subsidies increased China's growth.

Finally, our paper is linked with the literature on the growth experience of Japan, the first East Asian growth miracle. Christiano (1989) argues that the neoclassical growth model cannot account for Japan's postwar saving experience. Interestingly, he gets much better results when he slows down convergence to the model's balanced growth path (BGP) by introducing a minimum consumption level. Our assumption of TFP catch-up with respect to the U.S. yields the exact mechanism (slower convergence to the BGP, as there is a strong incentive to accumulate further capital only once TFP has caught up). This idea is corroborated by Chen et al. (2006), who show how using actual Japanese TFP growth rates in a standard growth model generates saving rates close to those in the Japanese data between 1956 and 2000. Parente and Prescott (1994) explore how barriers to technological adoption slow down convergence to a BGP and account for large differences in income per capita.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 compares China's growth experience with that of other Asian growth miracles. Section 4 presents the model and Section 5 calibrates it. Section 6 reports our quantitative results for China. Section 7 extends our model to the other East Asian growth miracles. Section 8 concludes. An Appendix offers additional details about the data and more quantitative exercises, including growth accounting.

2 Data

This section summarizes the data used in our study. Data for China, Taiwan, Hong Kong, Japan, Korea, and Singapore are from the Penn World Tables (PWT) 10.0 (see Feenstra et al., 2015, for details). The Appendix presents a full description of the construction of these data.

Real GDP is denoted as **GDP**, and we focus on the production definition of real GDP, which is $RGDP^O$ in PWT 10.0. PWT provides two different GDP measures for this variable for China. One is based on GDP growth constructed by Wu (2014), who has produced a measure of Chinese real GDP growth by modifying the official Chinese GDP data. Wu's measure of real GDP grows more slowly than the official real GDP from China's national accounts, the other measure of production-side GDP in the PWT. We focus on the official data because they conform well to other measures of China's economic growth, including energy consumption and construction activity, and because many other researchers agree with this assessment and use the official series, including Brandt and Zhu (2010), Zhu (2012), Song et al. (2011), Fernald et al. (2013), and Chang et al. (2016). Moreover, our main findings regarding China's growth performance would be strengthened if we used Wu's measure, as described below.

The variable **population** corresponds to *pop* in PWT 10.0. The population data from PWT 10.0 are slightly lower than those from the China Statistical Yearbook 2021, published by China's National Bureau of Statistics (NBS).[4] However, this difference is not quantitatively important.

The variable **investment rate** corresponds to the variable csh_i (or $csh_i + csh_x + csh_m$) in PWT 10.0. The investment rate is nominal investment divided by nominal output. We follow Whelan (2002) in using the nominal-to-nominal ratio, as he notes that this measure avoids some of the measurement difficulties involved in using ratios of chained (real) data and thus may be considerably more accurate. The nominal investment rate from China's National Bureau of Statistics is higher than that in PWT 10.0. These differences are plotted in Figure B.2 in Appendix B. We discuss the implications of this difference in the analysis below.

The variable **labor share** corresponds to the variable *labsh*. Finally, our variable **TFP** is constructed using variables $RTFP^{NA}$ and $CTFP$ in PWT 10.0. Due to data limitations, and although we use $RGDP^O$ as our GDP measure, the TFP data come from using $RGDP^{NA}$.[5] However, given the dynamics of the growth model that we will present later on, the quantitative differences for the arguments we want to make about the main patterns of the Chinese growth experience are likely to be small. We present a comparison of different measures of output in Figure B.1 in Appendix B.

[4] http://www.stats.gov.cn/sj/ndsj/2021/indexeh.htm

[5] PWT 10.0 constructs TFP only using $RGDP^{NA}$, the real GDP at constant national prices. The real capital stock reported by PWT 10.0 is also at constant national prices, which prevents us from estimating TFP directly using $RGDP^O$.

3 GDP Growth and the Investment Rate

This section compares China's GDP per capita growth and investment rate to those of other East Asian economies when they were at similar stages of development to provide a benchmark for interpreting China's growth performance. The data will show that China's per capita growth and investment rate have been similar to those of the other economies, suggesting a commonality among the factors driving growth in East Asia.

3.1 China's GDP Growth, 1995-2019

We analyze China's growth performance for the 25 years from 1995 to 2019. We begin in 1995 because China's per capita GDP of $2,825 (2017 PPP-adjusted dollars) at that time places it at the threshold of a middle-income country ($3,000). This per capita income level is roughly consistent with the World Bank's decision to classify China as a middle-income country in 1997.[6] Moreover, China's income per capita in 1995 was close to Japan's income per capita in 1950, the first year for which Japan's data are available in the PWT 10.0.

To compare China's performance to that of other East Asian economies that grew quickly, we find the country-specific year in which Japan, Korea, and Taiwan had approximately the same level of per capita income ($2,825 in 2017 PPP dollars). We then examine the evolution of GDP and investment dynamics for each of those economies from that time onward. In Appendix C, we extend the analysis to include Hong Kong and Singapore, which also had rapid growth. We do not include them in the main text to avoid cluttering the graphs and because their dynamics could reflect some factors that may be specific to the fact that they are city-states.

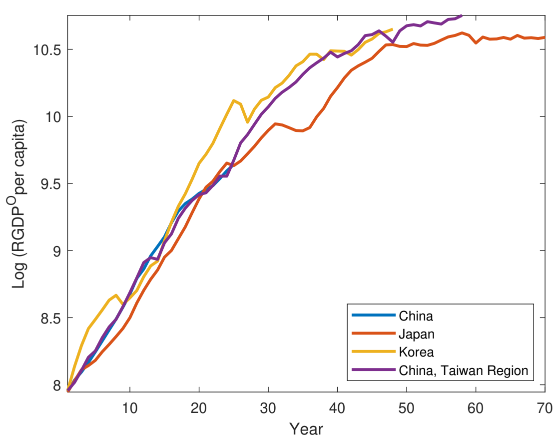

Figure 3.1: GDP per capita, normalized at China's GDP per capita in 1995

[6]We are less interested in the growth pattern of countries below the threshold of $3,000 because growth in low-income countries tends to be heavily dependent on issues that are specific to the very early stages of development, such as agricultural policies (e.g., collective farms vs. private/privatized farms), and are beyond the scope of our investigation. Nonetheless, the main findings are fairly robust to choosing a lower threshold income level.

Figure 3.1 shows real GDP for China, Japan, South Korea, and Taiwan, in which each economy begins at approximately the same per capita income level. The most striking aspect of these data is that the time path of China's per capita GDP is remarkably similar to those of the three other East Asian miracle economies. After 25 years of becoming middle-income economies, China, Japan, and Taiwan achieved nearly the same per capita real GDP level. The one exception is Korea, which was somewhat ahead of the other three. These data show nothing unusual about China's real GDP growth rate relative to the experiences of these other economies.

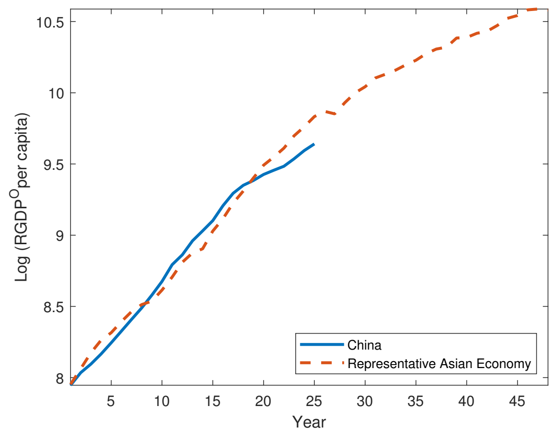

Figure 3.2: GDP per capita, normalized at China's GDP per capita in 1995

Figure 3.2 compares China's GDP per capita to the average of the other three economies (a "representative Asian economy"). China's real GDP path is marginally lower than the average of Japan, Korea, and Taiwan.

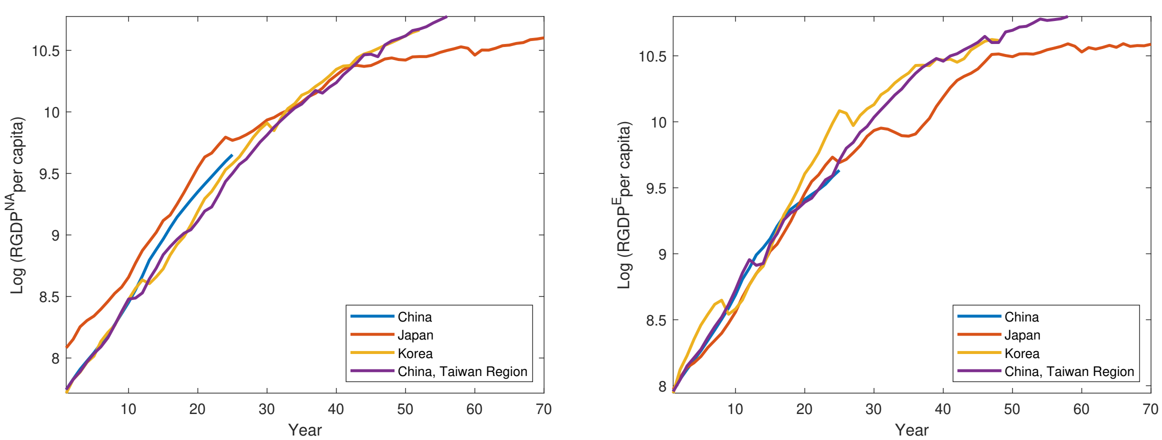

Figure 3.3: GDP per capita, alternative definitions

Figure 3.3 shows a very similar pattern between China and the other East Asian economies when per capita real GDP is measured using two alternative measures of GDP: the national accounts approach GDP ($RGDP^{NA}$) and the expenditure approach GDP ($RGDP^{E}$).[7] To check

[7]We note that Japan's first year (1950) $RGDP^{NA}$ per capita is \$3237, which was around 15% higher than

for robustness, Figure 3.4 uses 2000 as the base year for a normalization point in time since that was the year when China had GDP ($RGDP^{NA}$) per capita ($3350) that was closest to Japan's in 1950.

These figures suggest that China's future growth pattern will slow substantially, just as growth slowed substantially in Japan, Korea, and Taiwan after their first 25 years of becoming middle-income economies.

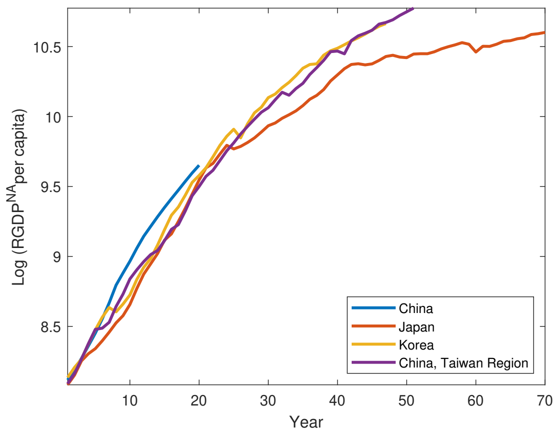

Figure 3.4: GDP per capita, alternative definitions, alternative base year

A possible shortcoming of the previous figures is that the world technological frontier was at a different point in time when China had (roughly) the same levels of income as the other East Asian economies. To control for this mechanism, we examine the evolution of GDP per capita relative to the U.S. level for each of those economies. This exercise assumes that the U.S. GDP per capita is a fair proxy for the GDP per capita achievable given the world technological frontier and prevailing social institutions.

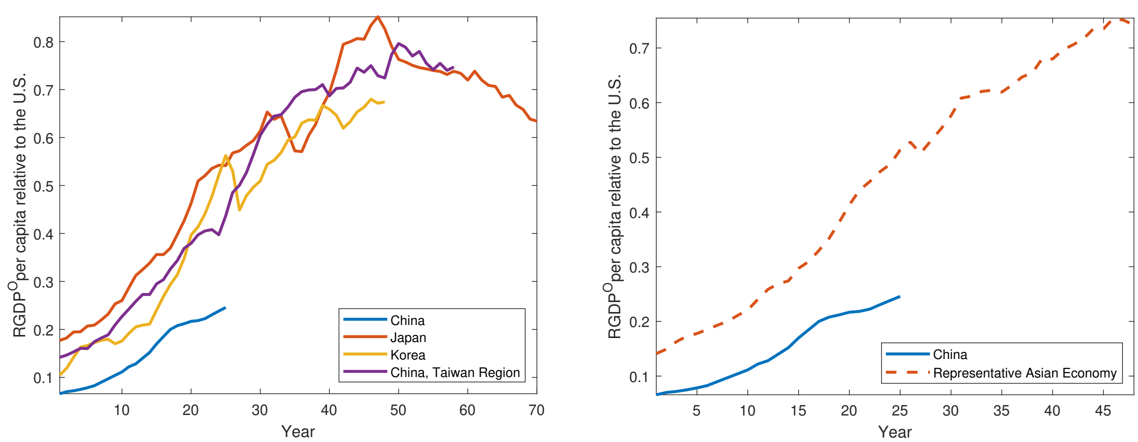

Figure 3.5: GDP per capita relative to the U.S., normalized at China's GDP per capita in 1995

China's $RGDP^{NA}$ in 1995. Japan started its industrialization in the late 19th century, much ahead of all the other East Asian economies, and the destruction of World War II did not erase this early advantage. See Fukao and Settsu (2021) for a summary of Japan's growth experience.

The left panel of Figure 3.5 shows how China's performance is less impressive than the performance of the other economies. This result is not a surprise: Figure 3.1 shows that China's growth performance was approximately the same as those of other East Asian economies. However, China's growth started much later, when U.S. GDP per capita was much higher. This means China benefited from an even higher technological frontier for its catch-up. This notion of distance to the technological frontier will play a key role in our model in Section 4. The right panel of Figure 3.5 replicates Figure 3.5 but now compares China's GDP per capita relative to the U.S. to the average GDP per capita relative to the U.S. of the "representative Asian economy" defined above.

3.2 China's Investment Rate

Figures 3.6 and 3.7 plot the investment rate for China, Japan, Korea, and Taiwan for their respective 25-year periods after becoming middle-income economies. The left panel presents gross capital formation (GCF). In contrast, the right panel of the figure presents an alternative concept of investment, which is the sum of gross capital formation and net exports. Variables are all in nominal terms in the figures. These data show that the nominal investment rate of China is qualitatively very similar to the investment rate in the other Asian miracle economies. Quantitatively, China's investment rate is somewhat higher.

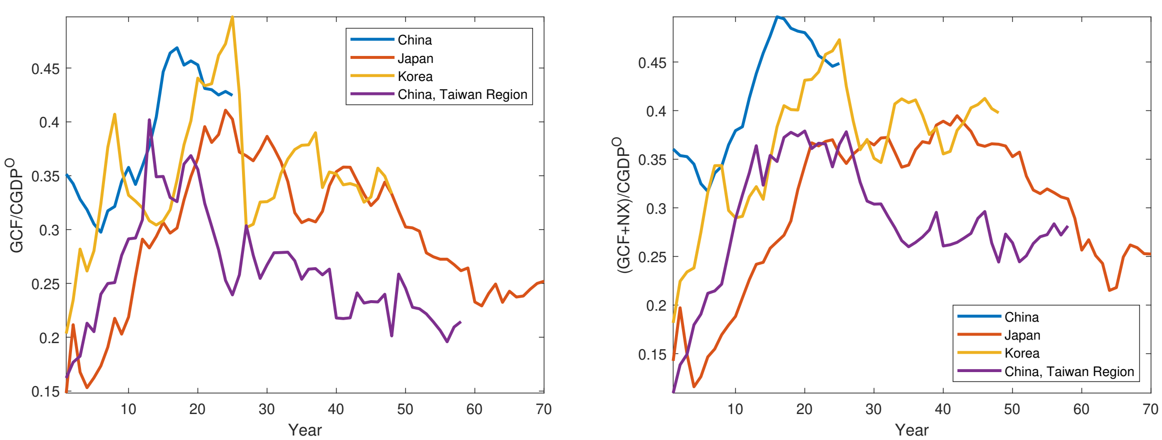

Figure 3.6: Investment rates

Figure 3.8 plots the return to capital across economies. Bai et al. (2006) construct the return to capital in China from 1978 to 2005, and Chen et al. (2019) extend their results to 2016. According to Chen et al. (2019), China's return to capital fell sharply after 2008 and became substantially lower than that of the other economies, suggesting low investment efficiency.[8] This observation will play an important role when we discuss how our model matches the observed

[8] PWT 10.0 also reports rates of return to capital for China. We believe the data from Chen et al. (2019) are superior. In any case, the results with data from PWT 10.0 are very similar and available upon request.

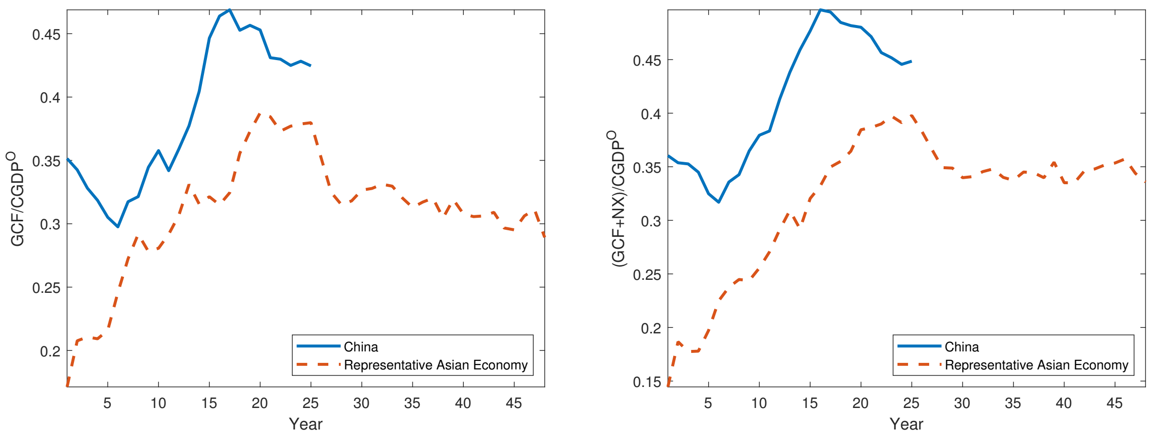

Figure 3.7: Investment rates

investment data.

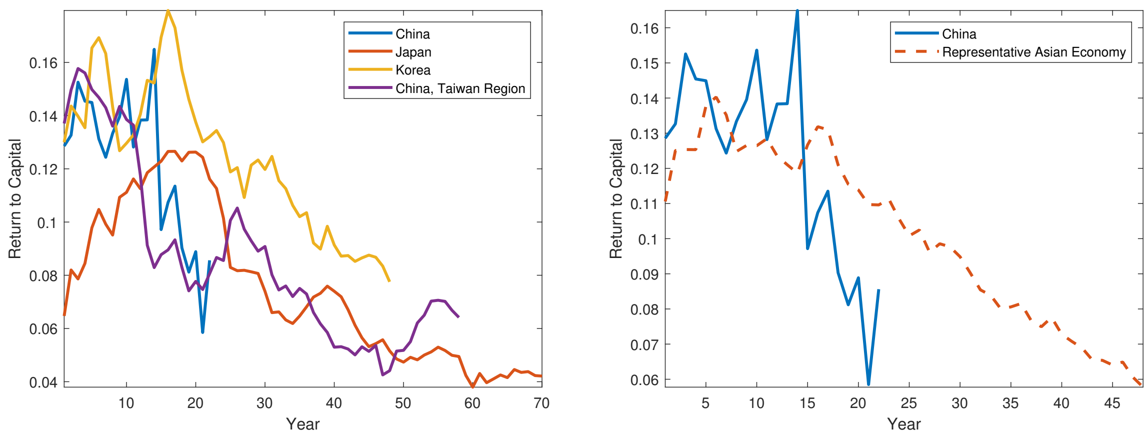

Figure 3.8: Return to capital

3.3 Taking Stock

We have documented two significant similarities between the growth experiences of China and those of Japan, Korea, and Taiwan:

1. China has grown at about the same rate as these other East Asian economies. The difference between China and the other economies is not in its economic growth rate but rather in the size of China's population compared to the other economies.

2. China's investment rate has been qualitatively similar to those of the other economies and quantitatively has been somewhat higher, suggesting capital accumulation contributed more to growth in China. Also, we documented a lower rate of return on capital in China.

The following section uses a minimalist growth model to interpret these facts.

4 A Minimalist Neoclassical Growth Model

This section presents a deterministic discrete-time Ramsey-Cass-Koopmans framework with two countries. One country is at the technology frontier, and its technology grows at a constant rate. The other country has a lower level of technology but catches up deterministically to some fraction of the technological frontier. There is no connection between the two countries other than the technological diffusion from the frontier to the catch-up country. These are all the elements needed to make our points and thus allow us to abstract from complicated specifications of the evolution of technology, uncertainty, multiple sectors, and other complexities. More broadly, we will see that our results suggest that to the extent that these and other complexities impacted China's growth, they may well be reasonably captured by our TFP catch-up process.

4.1 Preferences and Technology

The economy consists of two countries. The first country (China) is populated by an infinitely lived representative household of varying size N_t and whose preferences over per capita consumption are:

$$\max_{C_t/N_t} \sum_{t=0}^{\infty} \beta^t N_t \log\left(\frac{C_t}{N_t}\right),$$

where β is the discount factor and C_t is aggregate consumption.

China's output is given by $Y_t = K_t^{\theta}(A_t L_t)^{1-\theta}$ where K_t is capital, A_t is China's level of labor-augmenting technology, and L_t is employment (which we will assume is an exogenously given but time-varying fraction of China's population). Recall, when we present results in the next sections, that TFP is equal to $A_t^{1-\theta}$.

Output is used for consumption or investment I_t, which induces a law of motion for capital $K_{t+1} = I_t + (1-\delta)K_t$, where δ is China's depreciation rate. China's resource constraint is given by $C_t + I_t = Y_t$. Finally, we assume that China's population N_t grows at a time-varying rate n_t, so that $N_t = \prod_{i=1}^{t}(1+n_i)$, given $N_0 = 1$.

The second country, which we will call the U.S., is also populated by an infinitely lived representative household of varying size N_t^* (when a variable or a parameter does not have a superscript star, it denotes either a parameter for China or a parameter common to the U.S. and China) with the same preferences:

$$\max_{C_t^*/N_t^*} \sum_{t=0}^{\infty} \beta^t N_t^* \log\left(\frac{C_t^*}{N_t^*}\right),$$

and technology $Y_t^* = (K_t^*)^{\theta^*}(A_t^* L_t^*)^{1-\theta^*}$. This is the same technology as China's, except for four elements.

First, the capital share in the U.S. is θ^*, and, second, the depreciation rate in the U.S. is δ^*. The differences between China and the U.S. in the elasticity of output with respect to capital and the depreciation rate might be due to differences in industry composition, production techniques, or maintenance decisions by firms that we do not model.[9] Third, the U.S. population grows at a time-varying rate of n_t^*: $N_t^* = \prod_{i=1}^{t}(1+n_i^*)$, given $N_0^* = 1$. We will only compare China and the U.S. in terms of relative per capita income. Thus, population size does not matter in our exercise, and normalizing $N_0 = N_0^* = 1$ is irrelevant for our purposes. The fourth is the U.S. technology A_t^*, which we consider as the world technological frontier, and which grows at the exogenously given rate g, that is, $A_t^* = (1+g)^t$.

4.2 Specifying Technological Catch-Up

There are two regimes for China's technology, A_t. In the first regime, $t \leq 0$, A_t is a small constant fraction $0 < \gamma_0 \ll 1$ of the U.S. technology. This regime stands in for China's pre-reform economy:

$$A_t = \gamma_0 A_t^* = \gamma_0 (1+g)^t.$$

In the second regime, $t > 0$, China's technology is governed by a catch-up process. This regime stands in for China's post-reform economy:

$$A_t = \gamma \alpha_t A_t^* = \gamma \alpha_t (1+g)^t.$$

The catch-up process is a simple one, but it captures important common features that countries experience during the process of catch-up growth, including (i) the level to which they ultimately catch up, (ii) the speed of their catch-up, and (iii) how their catch-up speed changes over time.[10] Three parameters characterize the process we use. There is a constant, γ ($0 < \gamma_0 \ll \gamma < 1$), the asymptotic fraction of the U.S. technology that China can achieve. We interpret γ as a measure of institutional quality (and, therefore, the ability to adopt new technologies). However, other interpretations are possible (e.g., it can be interpreted as an index of human capital or allocative efficiency). The second term, $0 < \alpha_t < 1$, measures China's distance to the technology frontier γ at any point in time.

Given an initial value α_1 such that $\gamma_0 < \alpha_1 \gamma$ (i.e., technology is better after economic reform begins), we assume that $\ln(\alpha_t)$ is declining at rate η, which is the third parameter:

$$\ln(\alpha_t) = (1-\eta)\ln(\alpha_{t-1}),$$

[9] See, on maintenance and repair, McGrattan and Schmitz (1999). Below, we will use data from PWT 10.0 to discipline our choices for θ and δ.

[10] Appendix D documents, using growth accounting, that TFP growth accounts for 75% of China's output growth from 1995 to 2019. Thus, a process of TFP catch-up is key for understanding China's aggregate dynamics.

(i.e., α_t is growing in levels). This produces a concave path for α_t that asymptotes to 1. This captures the common experience among low-income economies that grow very quickly during the early stages of modernization, with a rapid adoption rate (e.g., "low-hanging fruit"). Still, the rate slows down as it approaches the frontier when improvements become harder.

4.3 The Stationary Problem

Given the model has continuous growth, we follow standard practice and normalize the variables. We use the growth rates of U.S. technology and China's population to make the problem stationary in China. We use the U.S. growth rate of technology as a scaling factor because before reform $A_t = \gamma_0 A_t^*$ and after reform $A_t \to \gamma A_t^*$. Specifically, let:

$$
\begin{aligned}
c_t &= \frac{C_t}{A_t^* N_t}, \\
k_t &= \frac{K_t}{A_t^* N_t}, \\
i_t &= \frac{I_t}{A_t^* N_t}, \\
y_t = \frac{Y_t}{A_t^* N_t} &= \left(\frac{K_t}{A_t^* N_t}\right)^{\theta} \left(\frac{A_t L_t}{A_t^* N_t}\right)^{1-\theta} = k_t^{\theta} (\gamma \alpha_t l_t)^{1-\theta},
\end{aligned}
$$

where l_t denotes the exogenously given employment-to-population ratio L_t/N_t.

Therefore, we rewrite the social planner's problem for post-reform China as follows:

$$
\max_{c_t} \sum_{t=0}^{\infty} \beta^t N_t \log c_t
$$

subject to:

$$
\begin{aligned}
y_t &= k_t^{\theta} (\gamma \alpha_t l_t)^{1-\theta}, \\
c_t + i_t &= y_t, \\
i_t &= (1+g)(1+n_{t+1})k_{t+1} - (1-\delta)k_t.
\end{aligned}
$$

A standard Euler equation characterizes the solution to the optimization problem:

$$
c_t^{-1}(1+g) = \beta c_{t+1}^{-1} \left(\theta k_{t+1}^{\theta-1} (\gamma \alpha_{t+1} l_{t+1})^{1-\theta} + 1 - \delta\right).
$$

There is an equivalent stationary problem for the social planner in the U.S., where we use the U.S. growth rate of technology and population to scale the variables, but we skip this in the interest of space. The Euler equation for this normalized problem is given by:

$$(c_t^*)^{-1}(1+g) = \beta(c_{t+1}^*)^{-1}\left(\theta^*\left(k_{t+1}^*\right)^{\theta^*-1}\left(l_{t+1}^*\right)^{1-\theta^*} + 1 - \delta^*\right).$$

5 Calibration

We calibrate the nine parameters of our model on an annual base using data for 1995-2019 (the same period as in Section 3). For the first six parameters, we follow commonly used values and targets. See Table 1 for a summary of the calibration, and, in Subsection 6.2, we will discuss how robust our quantitative results are to changes in our calibrated values.

Table 1: Calibration

Parameter		Value
Discount factor	β	0.948
Capital share, China	θ	0.5
Capital share, U.S.	θ^*	0.39
Depreciation rate, China	δ	0.054
Depreciation rate, U.S.	δ^*	0.040
Technology growth rate, U.S.	g	0.02
Initial technology, China	γ_0	0.03
Catch-up bound	γ	0.22
Catch-up rate	η	0.08

We pick a discount factor β of 0.948 to replicate a 7.6% annual rate of return to capital reported by PWT 10.0 for the U.S. between 1995 and 2019 (notice that our model matches the return on all capital goods, not on bonds or other financial assets). We choose the capital share $\theta^* = 0.39$ for the U.S. to match the average shares from PWT 10.0. For China, in comparison, we pick $\theta = 0.5$ instead of using the capital income share from the PWT (0.42). Bai et al. (2006) defend this higher value by using data on provincial capital income share from China's National Bureau of Statistics. Indeed, the share of 0.5 is widely used in the literature on China, such as Song et al. (2011) and Brandt et al. (2008). Nonetheless, in Appendix F, we report our results with $\theta = 0.42$.

Similarly, the depreciation rate is the average depreciation rate from PWT 10.0: $\delta = 0.054$ for China and $\delta^* = 0.040$ for the U.S. We calibrate $g = 0.02$ to match the long-run productivity growth of the U.S. economy during the 20th century. The population growth rates match the observed data year by year in China and the U.S. These parameters yield a steady-state K/Y ratio in the U.S. of 3.39, which is close to the average value of 3.24 (including government capital stocks) over the 1995-2019 period.

To specify the three catch-up parameters for China, we assume that the U.S. was on its BGP in 1995 or, in the transformed stationary problem, that the U.S. was at its steady state with capital k_0^*. More concretely, k_0^* satisfies the steady-state Euler equation:

$$1+g=\beta\left(\theta^*\left(k_0^*\right)^{\theta^*-1}\left(l_0^*\right)^{1-\theta^*}+1-\delta^*\right), \tag{1}$$

where l_0^* is the employment-to-population ratio in the U.S. in 1995. Solving this equation with the calibrated parameters above yields $k_0^* = 3.59$. To compute $i_0^* = (1+g)(1+n_0^*)k_0^* - (1-\delta)k_0$, we assume that population growth, n_0^*, is the population growth observed in 1995.

Similarly, we assume that in 1995 ($t = 0$), China was at the initial steady state implied by γ_0. Thus, we have the Euler equation:

$$1+g=\beta\left(\theta(k_0)^{\theta-1}(\gamma_0 l_0)^{1-\theta}+1-\delta\right). \tag{2}$$

where l_0 is China's employment-to-population ratio in 1995. To compute investment in this steady state $i_0 = (1+g)(1+n_0)k_0 - (1-\delta)k_0$, we assume that n_0 is equal to the population growth rate in China in 1995.

From PWT 10.0, we know that, in 1995, China's PPP-adjusted per capita output was 0.066 of the U.S. per capita output:

$$\frac{Y_0/N_0}{Y_0^*/N_0^*}=0.066.$$

That is,

$$\frac{y_0}{y_0^*}=\frac{(k_0)^\theta}{(k_0^*)^{\theta^*}}\frac{l_0^{1-\theta}}{(l_0^*)^{1-\theta^*}}\gamma_0^{1-\theta}=0.066. \tag{3}$$

Equations (2) and (3) are a system of two equations in two unknowns: γ_0 and k_0 (given k_0^*). This system yields $\gamma_0 = 0.03$ and $k_0 = 0.27$. These numbers imply that the ratio in per capita output between the U.S. and China in 1995 (around 15) is accounted for by a relative TFP of around 5.8 ($\approx 1/(\gamma_0^{1-\theta}) = 1/(0.03^{0.5})$) and a ratio of capital of around 13.4, but which is partially offset by a Chinese employment-to-population ratio that is around 17% higher.

Next, we assume that an institutional reform occurs in 1996 ($t = 1$), in which γ_0 jumps to γ: $\frac{A_0}{A_0^*} = \gamma_0$, $\frac{A_1}{A_1^*} = \alpha_1\gamma$, ..., $\frac{A_{T-1}}{A_{T-1}^*} = \alpha_{T-1}\gamma$, $\frac{A_T}{A_T^*} = \gamma$. The last equation only holds approximately for a large T, which we pick to be 105, as α_T asymptotes to 1. As we will see momentarily, T is given by the furthest UN population projections into the future, which we use to render the problem stationary.

Finally, we choose the institutional level γ and convergence rate η to minimize the distance (using the least squares metric) between the generated TFP for China in the model and the TFP in the data. This process yields $\gamma = 0.22$ and $\eta = 0.08$. The parameter α_1 therefore evolves as $\ln(\alpha_1\gamma) = (1-\eta)\ln(\gamma_0)$.

We pick this particular form of TFP convergence process because it is simple, easy to replicate and use in other applications, has just three parameters, and captures the spirit of the underlying mechanisms that the large literature on convergence has discussed, including relatively fast growth in the early stages of catch-up and a limit to convergence (e.g., conditional convergence), which may reflect institutional factors. Our analysis below requires forecasting the future path of TFP, and our process captures past data extremely well based on just a few parameters.[11]

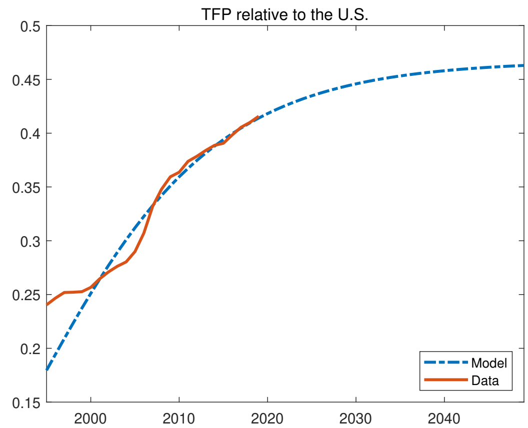

Figure 5.1: China's TFP relative to the U.S.

Figure 5.1 plots China's TFP relative to the U.S. in the data and our calibration to illustrate how the former matches the latter. The estimated process starts somewhat lower but fits the data very closely for 2000 to 2019, inducing the "right" neoclassical growth dynamics for most of the observed series, as shown below. Choosing a value of $\gamma = 0.22$ implies a final TFP of 46.7% of the U.S. level, indicating serious structural/institutional limitations in China's long-run productivity level. While our simple process abstracts from those limitations and their specific effects, it captures the implicit impact of those factors on TFP up through 2019 well.

We stress that specifying that the institutional reform began in 1996 is immaterial for the numerical results. It is simply a convenient benchmark. The assumptions we make are that China was on the BGP implied by γ_0 in 1995 and that the technological gap with the U.S. began to close at a constant rate from that moment. An analogous interpretation is that γ_0 is already the product of a reduction in the technological gap between the U.S. and China that started with the first wave of reforms that followed the Third Plenary Session of the 11th Central Committee of the Chinese Communist Party in December 1978.[12] Given the low level of γ_0 that we infer, 0.03,

[11]The alternative of directly feeding China's observed TFP into our model does not allow us to extrapolate future values for China's TFP. Our fitted TFP curve makes forecasting future TFP transparent and simple.

[12]See Naughton (2018, Ch. 5) for a chronology of economic reform in China since 1978. The main economic reforms came in two waves: from 1979 to 1989 and 1993 to 1999. For example, China unified its dual foreign exchange regime at the end of 1993, there was a large fiscal reform in 1994, and inflation was tackled in 1996. Thus, our benchmark year of 1996 is in the middle of the second reform wave, making it a reasonable point for our analyses and its focus on China's growth after China became a middle-income country. Chen and Zha (2023) emphasize the gradualist approach of China's reforms, which fits well with our idea of a smooth TFP catch-up.

whether China was at its BGP in 1995 or still converging to it, given the first wave of reforms from 1979 to 1989, has minimal effects on our simulations below.

The final steady-state capital of the U.S. k^*_T is given by:

$$(1+g) = \beta \left(\theta^* \left(k^*_T\right)^{\theta^*-1} \left(l^*_T\right)^{1-\theta^*} + 1 - \delta^* \right),$$

where l^*_T is the U.S. employment-to-population ratio in 2019. To compute investment at T, we assume that population growth in the limit, $n^*_T = 0.024\%$, is the total population growth rate predicted by the UN in 2100, the furthest projection as of early 2023.

Similarly, for China, the final steady-state capital k_T is given by:

$$(1+g) = \beta \left(\theta k_T^{\theta-1} (\gamma l_T)^{1-\theta} + 1 - \delta \right),$$

where l_T is the employment-to-population ratio in China in 2019. To compute investment at T, we assume that population growth in the limit, $n_T = -1.19\%$, is the growth rate of the total population in 2100 predicted by the UN.

To compute the planner's solution for the U.S., we take the initial steady state in 1995 and its final steady state in 2100 and compute the transition path between the two using the Euler equation and the investment and resource constraint equations using a nonlinear equation solver. For China, we take the initial steady state in 1995, assume that the reform started in 1996, and compute the transition path toward the final steady state in 2100, where $A_{2100} = \gamma A^*_{2100}$ using the same set of equations and the same nonlinear equation solver.

6 Quantitative Results

This section presents the quantitative findings. First, we report our benchmark results and discuss how they compare with the data. Second, we explore the robustness of the results to alternative calibrations and assumptions regarding how investment is determined in the model.

6.1 Benchmark Results

Figure 6.1 compares the model results to the data on the level of income per capita and its growth rate. The left panel shows that the model closely matches the evolution of China's income per capita relative to the U.S. income per capita from 1995 to 2019. This finding suggests that the minimalist growth model captures the main mechanisms underlying China's growth experience. Put differently, China's growth during this period conforms to that generated by a simple, deterministic, one-sector optimal growth model with a smoothed, exogenous TFP path.

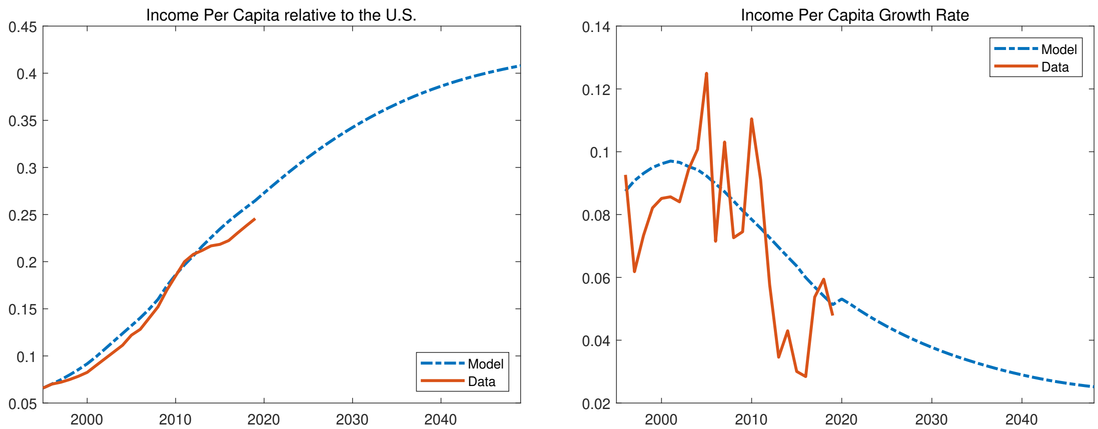

Figure 6.1: Transitional dynamics (model and data): Income per capita

Catch-up. The first main implication of the analysis is that China's catch-up process will slow substantially as China's technology level approaches its catch-up bound ($\gamma = 0.22$). Since China has differences in its capital share, depreciation rate, population growth, and employment-to-population ratio, its relative income per capita asymptotes to 0.44, which is lower than its relative TFP of 47%. The UN forecasts that, by 2100, China's population will be around 1.95 times larger than that of the U.S. (766.7 million vs. 394 million). Given this forecast, the model implies that the U.S. economy will be around 17% larger than China's by the end of this century.

The sharp slowdown in income per capita growth is seen in the right panel of Figure 6.1, which compares income per capita growth rates in the model and the Chinese data. The model replicates the hump shape of these patterns and their slowdown over time. Since our model has no shocks, we do not track the year-to-year fluctuations. Nonetheless, there is a clear similarity between the model and the data in terms of longer-run movements.

Investment. The left panel of Figure 6.2 draws investment as the model's forecast against two different investment series in the data: PWT is the investment rate reported by PWT 10.0. NBS is the investment series from China's National Bureau of Statistics (the investment rates are in nominal ratios, which avoids the need to compute relative prices). The right panel of Figure 6.2 shows the rate of return on capital from Chen et al. (2019) and compares it with the model.

Our model undershoots in terms of investment at the end of the period. But this raises the question of why China has not grown more, given this higher investment. Four factors can help reconcile the model with the data. The first factor is that much of China's investment has had a low marginal productivity due to distortions. For example, König et al. (2022) have documented that the return to the productivity of R&D investments is lower in China than in Taiwan due to pervasive output wedges. This explanation fits well with the right panel of Figure 6.2, which shows that the rate of return to capital in the model is higher than in the data. This argument also reconciles the data with the three stylized facts from Section 3: while investment rates in

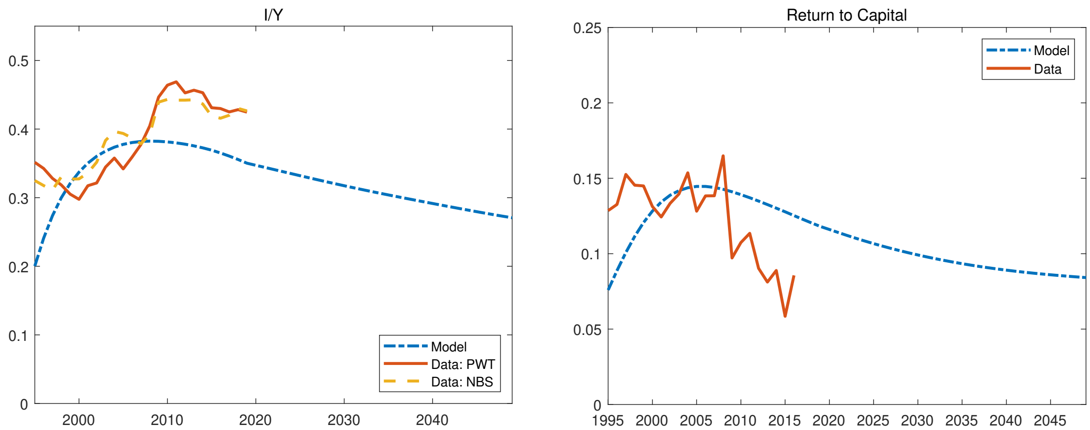

Figure 6.2: Transitional dynamics (model and data): Investment, and return to capital

China have been somewhat higher than in other East Asian economies, the growth rate has been about the same and the rate of return on capital lower.

The second factor is the mismeasurement of investment. Chen et al. (2019) have argued that the investment rate in 2016 was about 7 percentage points lower than officially reported (see also Xiong, 2019, for a model of how flawed statistics come about in China due to the agency problem between central and local governments). Figure 6.3 plots the results of our model against the nominal investment rate in Chen et al. (2019, p. 106, Figure 10).[13] With the series from Chen et al. (2019), the model does an excellent job of replicating investment. This explanation has the additional advantage that, by lowering the total amount of investment, it increases the rate of return on capital, helping with that aspect of the model as well.

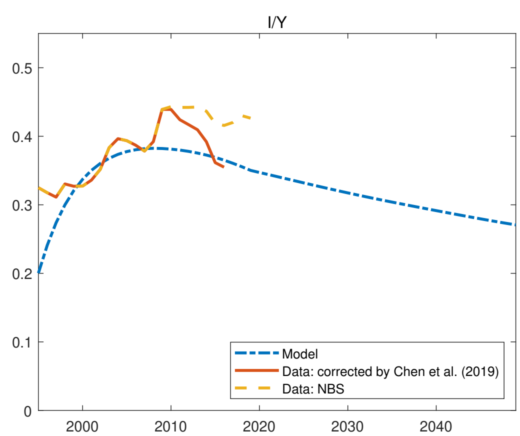

Figure 6.3: Corrected investment by Chen et al. (2019)

The third factor is that much of the investment in China has reflected political priorities. Herd (2020) has shown that business investment has the same hump shape we document. Compare

[13]For this comparison to be precise, we would need to recalibrate China's TFP process to be consistent with the output data in Chen et al. (2019). Since we aim to point out that the official investment series may mismeasure the extent of capital formation, we skip this additional step.

Figure 6.B in his paper with our right panel in Figure 6.2; the difference between the model and the data can be attributed to the increase in investment in housing and infrastructure. Bai et al. (2020) and Jiang et al. (2022) make related points based on the policies of local authorities, which have taken advantage of the local government financing vehicles to borrow and invest in low-rate-of-return projects. However, notice that investment in infrastructure in China has fluctuated between 7% and 10% of output since the early 2000s (Jiang et al., 2022). While this is a high rate compared to other middle-income countries, investment in infrastructure can only account for part of the difference in total investment.

To explore this third factor more thoroughly, in Subsection 6.2, we will document that if we fix investment exogenously, à la Solow, at a higher level, the model cannot match the data on per capita income and per capita income growth. This factor fits particularly well with Figure 6.2: the rate of return on capital in the data is lower than in the model, while the investment rate is higher than in the model, especially after 2008, because many investment decisions during this period were made for reasons other than maximizing profits. Therefore, endogenous saving decisions must be a key part of the analysis.

Finally, a fourth factor is that our calibration might yield too little investment. We will revisit and dismiss this explanation in Subsection 6.2.

Future income growth rates. Our second main result is that, according to our model, by 2043, U.S. output will grow faster (2.29%) than China's (2.27%) (recall the model forecast that, by 2100, the U.S. economy will be larger than China's in PPP terms; in 2019 according to PWT 10.0, China's economy is larger than the U.S. economy). This can be seen in the left panel of Figure 6.4, which plots the (total) income growth rates for China and the U.S. This sharp prediction is based on three elements. One is China's unfavorable demographics, with a falling population.[14] Second is the stability of the employment-to-population ratios (right panel of Figure 6.4). The third is the slowdown of technological catch-up (Figure 5.1).

The decomposition of these three elements helps us gauge possible caveats to the model's forecast. Let us start with demographics. This is the least uncertain of the three elements. Absent a particularly deadly epidemic or massive migration flows, the population as far out as 2043 is roughly already determined: around 85% of the Chinese people who will be alive in 2043 have already been born (and nearly all workers: a person born in 2023 will only be 20 years old in 2043, and thus either studying or just entering the labor force). Demographic (and social norms) changes might have a greater impact on the employment-to-population ratio but by only a few percentage points up or down. If anything, the right panel of Figure 6.4 suggests that China's employment-to-population ratio will fall due to aging, accelerating the date at which the U.S.

[14] According to China's National Bureau of Statistics, China's total population started falling in 2022.

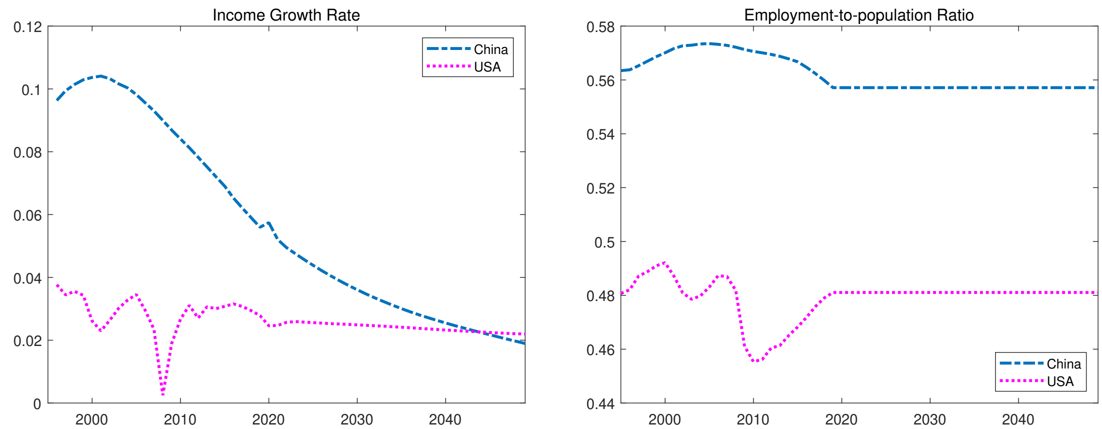

Figure 6.4: Transitional dynamics: Total income and employment-to-population ratio

would overtake China's output growth rate.[15]

The main source of uncertainty in the model forecast of per capita income growth rates is the evolution of relative TFP in Figure 5.1. Our calibrated process matches the data, and the evidence of other East Asian economies presented in Section 3 suggests that avoiding a slowdown in the convergence process is a tall order. Even if TFP growth falls in the U.S. (lowering total U.S. income growth), this reduction in TFP might hurt China through international technological diffusion or just postpone the takeover by the U.S. in terms of the income growth rate for a decade.

The logic of the neoclassical growth model is straightforward and powerful in terms of its implications. With a population growth rate in the long run of -1.19% projected by the UN (which assumes a recovery in fertility from the current 1.08 children per woman to 1.48), TFP needs to grow 1.19% a year just to keep total income constant (given constant employment-to-population and capital-worker ratios). Delivering a growth rate of 3% would require a TFP growth rate of 4.19% in an economy with a middle-high income already. While not out of the realm of possibility, given the upside potential of AI and automation, the historical record is not sanguine about the likelihood of a long-run TFP growth rate of 4.19%. A shrinking population and slowing TFP catch-up are substantial headwinds facing China.[16]

[15] Nonetheless, there might be a counterbalancing effect as the average human capital of Chinese workers increases as younger, more educated cohorts enter the job market. We will explore this possibility in future research. Still, the experience of southern European countries, such as Italy and Spain, that have experienced a similar increase in average human capital over the last 25 years, is not very optimistic. In the absence of fast TFP growth, many highly educated workers are allocated to jobs for which they are overqualified, thus substantially reducing the potential contribution of higher human capital.

[16] There are many other possible headwinds (financial crises, adverse political developments, etc.), but since those are absent from our model, we remain silent about them. Suffice it to say that considering them will only make our main argument stronger.

6.2 Robustness

We conduct an extensive battery of robustness exercises with our model by varying the parameter values. In the interest of space, we will report the three most interesting experiments. In these exercises, we recompute the institutional levels γ_0 and γ and the convergence rate η to minimize the distance between TFP in the model and the data to make all results easily comparable.

Changing the elasticity of intertemporal substitution. In our first experiment, we change the utility function from a log to a general CRRA form. If, for instance, we lower the elasticity of intertemporal substitution to $1/2$, we deliver a model-implied investment rate much lower than that in the data (bottom right panel of Figure 6.5) and, with it, we miss the growth of income per capita (top left panel of Figure 6.5).

The converse exercise of increasing the elasticity of intertemporal substitution over one delivers a better match of investment to the data (which we argue may reflect significant mismeasurement), but too much income growth relative to the data.

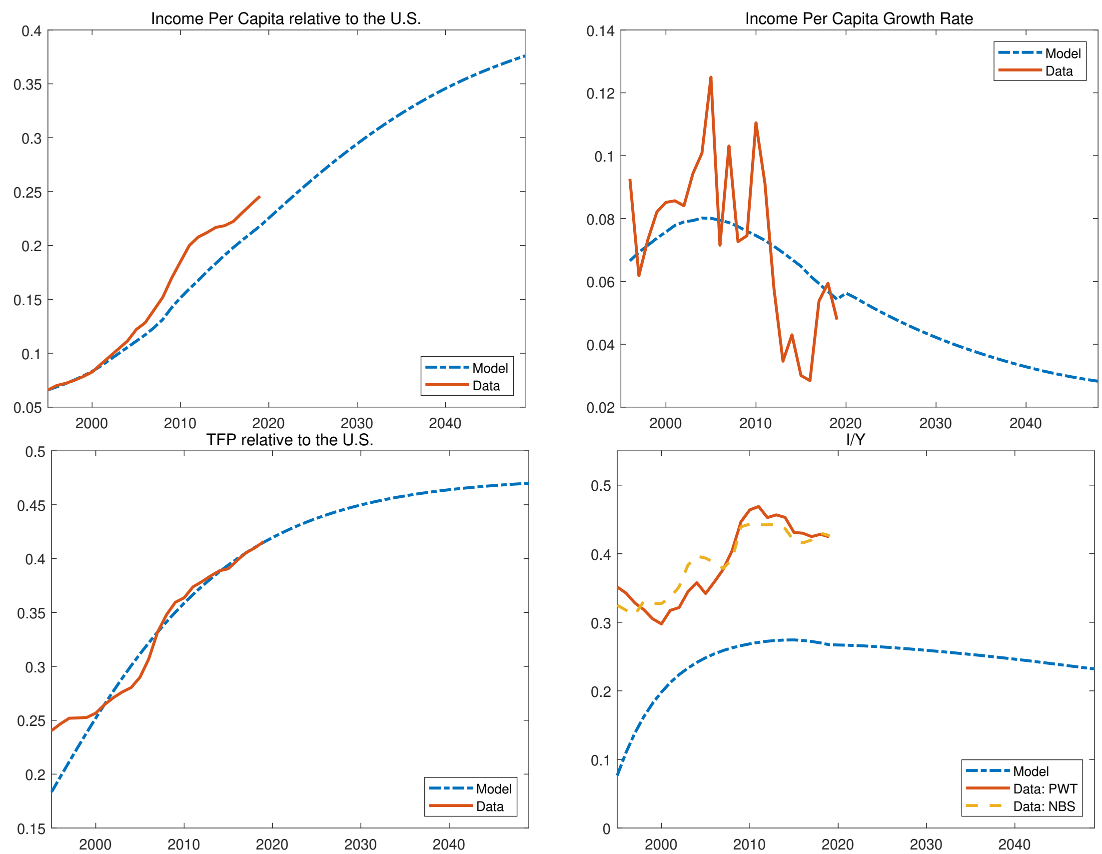

Figure 6.5: Transitional dynamics (model and data): IES $= 1/2$

From Ramsey-Cass-Koopmans to Solow. In our second experiment, we fix the investment rate exogenously at its 2009 level (0.45) to capture the evidence in the left panel of Figure 6.2. We can think of this exercise as moving from a Ramsey-Cass-Koopmans model, where consumption-saving decisions are chosen endogenously, to a Solow model, where consumption-saving is fixed exogenously. We interpret this experiment as addressing the possibility that our model misses key mechanisms behind investment or because investment is determined by political forces outside the model.

Figure 6.6 reports these results. Here China continues growing much faster for a longer period. However, the model can no longer match the recent per capita income and income growth observations. While a simple Solow model can deliver fast economic growth from the high investment rates of China for decades to come, the data do not support the hypothesis that China is getting much out of this high investment rate. Our previous explanation –a combination of low investment efficiency and mismeasurement– matches the data better.

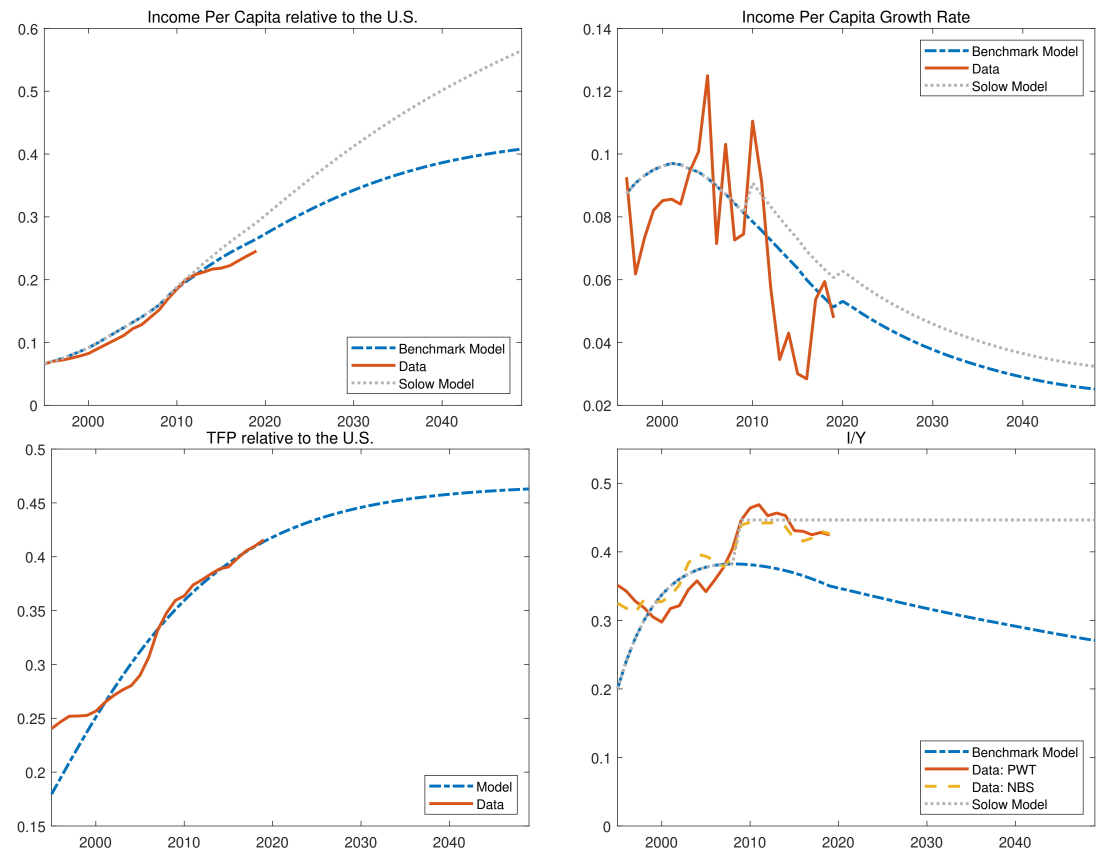

Figure 6.6: Transitional dynamics (model and data): Investment rate fixed at 0.45

Changing the depreciation rate. In our third experiment, we experiment with alternative depreciation rates. In Figure 6.7, we set the depreciation rate to 0.1 for both China and the U.S. The simulation results illustrate that the model's prediction is robust to this alternative calibration. Remember that, in our benchmark calibration, the depreciation rate is 0.04 for the U.S. according to PWT 10.0. Hence, we also experiment with setting both countries' depreciation rates to 0.04. Given that the model's predictions are similar, we do not report them.

Figure 6.7: Transitional dynamics (model and data): $\delta = \delta^* = 0.10$

7 The Asian Miracle Economies

We have shown that a minimalist Ramsey-Cass-Koopmans model can account for the growth experience of China since 1995 and that it projects a substantial slowdown in China's middle-run growth prospects. But can this model account for the experience of other Asian miracles? In this section, we will argue that the answer is yes, and that the results for these other economies will provide important support for our assertion that the model captures important aspects of China's catch-up experience. This validation exercise is particularly valuable because, as we explained in the introduction, the model should account for longer experiences (for example, in the case of Japan, since the late 1950s) than China's, pushing the model to its explanatory limits. This provides additional confidence for projecting China's TFP process into the future.

More concretely, we examine the experiences of Korea, Japan, and Taiwan.[17] In all cases, we calibrate the model as in Section 5. The discount factor and the three U.S. parameters remain unchanged; we pick the capital share and depreciation rates to match the PWT 10.0 data, the

[17] We also show the growth experience of Hong Kong and Singapore in Appendix F.

parameter values for TFP to match the observed TFP using a quadratic metric, and the reported population growth rates. We will make an exception for Taiwan regarding θ. PWT 10.0 imputes self-employed labor compensation, assuming the self-employed earn the same average wage as employees, which greatly overestimates labor share and leads to a θ as low as 0.27 for Taiwan. Instead, we choose an alternative labor share for Taiwan from PWT 10.0, which only includes the labor compensation of employees; this gives $\theta = 0.47$. Nonetheless, we report the results for Taiwan with $\theta = 0.27$ in Appendix F.[18]

As an initial year for our calibration, we pick the same year we use for each economy in Figure 3.1. If TFP data are unavailable from the initial year, we start from the first year when TFP data are available. The long-run relative TFP (γ) ranges between 0.45 for Korea and 0.67 for Taiwan. Table 2 summarizes the calibration.

Table 2: Calibration: The Asian miracle economies

Economy	δ	γ	η	θ	Initial Year	TFP Available From
Japan	0.037	0.51	0.17	0.40	1950	1954
Korea	0.045	0.45	0.10	0.45	1972	1972
Taiwan	0.050	0.67	0.06	0.47	1962	1962
China	0.054	0.22	0.08	0.50	1995	1995

Japan. First, we consider Japan. In the top left panel of Figure 7.1, we see that the model does a fair job of tracking Japan's experience until the late 1980s, when output grows for several years well above the predictions of the model (top right panel of Figure 7.1). This divergence between the model's predictions and the data is compatible with the notion that Japan went through an asset price bubble (typically dated from 1986 to 1991) that pushed the economy above its neoclassical transitional dynamics. Furthermore, the return of per capita income to the level predicted by the model at the end of the top left panel of Figure 7.1 accounts for the poor performance of Japan's economy since 1992, as the effects of the previous decade were undone.

The plot for investment (middle right panel of Figure 7.1) reinforces this view, as investment rates were substantially higher in the 1980s and 1990s than the model's predictions. The evolution of TFP (middle left panel of Figure 7.1) suggests that Japan's convergence process with the U.S. has finalized with a per capita income of around 60% of the U.S. level. Finally, the bottom panel of Figure 7.1 shows that the observed rate of return on capital is below what the model predicts by the end of the period. Nonetheless, the model captures the hump shape of this rate of return, reflecting the interaction between productivity, investment, and labor.

[18] See `https://www.stlouisfed.org/on-the-economy/2018/january/measuring-labor-share-asian-tigers` and Appendix E about the difficulties of measuring labor shares in East Asian economies.

Figure 7.1: Japan: Transitional dynamics

Korea. Figure 7.2 plots the results for Korea. We can see in the top left panel that the model has done quite well in accounting for growth at the start (from around 1970 to 1985) and the end of the sample (from around 2000 to 2019). The model misses, however, the very high growth rates between 1985 and 1996 that pushed observed output well above the predicted level and the great Asian financial crisis of 1997. Interestingly, and in contrast with Japan, Korea has not returned to the path forecast by the model. However, the data at the end of the sample in the top right panel of Figure 7.2, where growth seems to underperform the model's forecast, suggest that such a return might start appearing.

Nonetheless, the central message of the model, namely, that growth slows down after a country catches up, is vindicated by the data: Korea has been growing recently exactly as the Ramsey-Cass-Koopmans model predicts, if not slightly more slowly (top right panel of Figure 7.2). The model also matches TFP relative to the U.S. (middle left panel), the hump shape of investment (middle right panel), and the rate of return on capital (bottom panel).

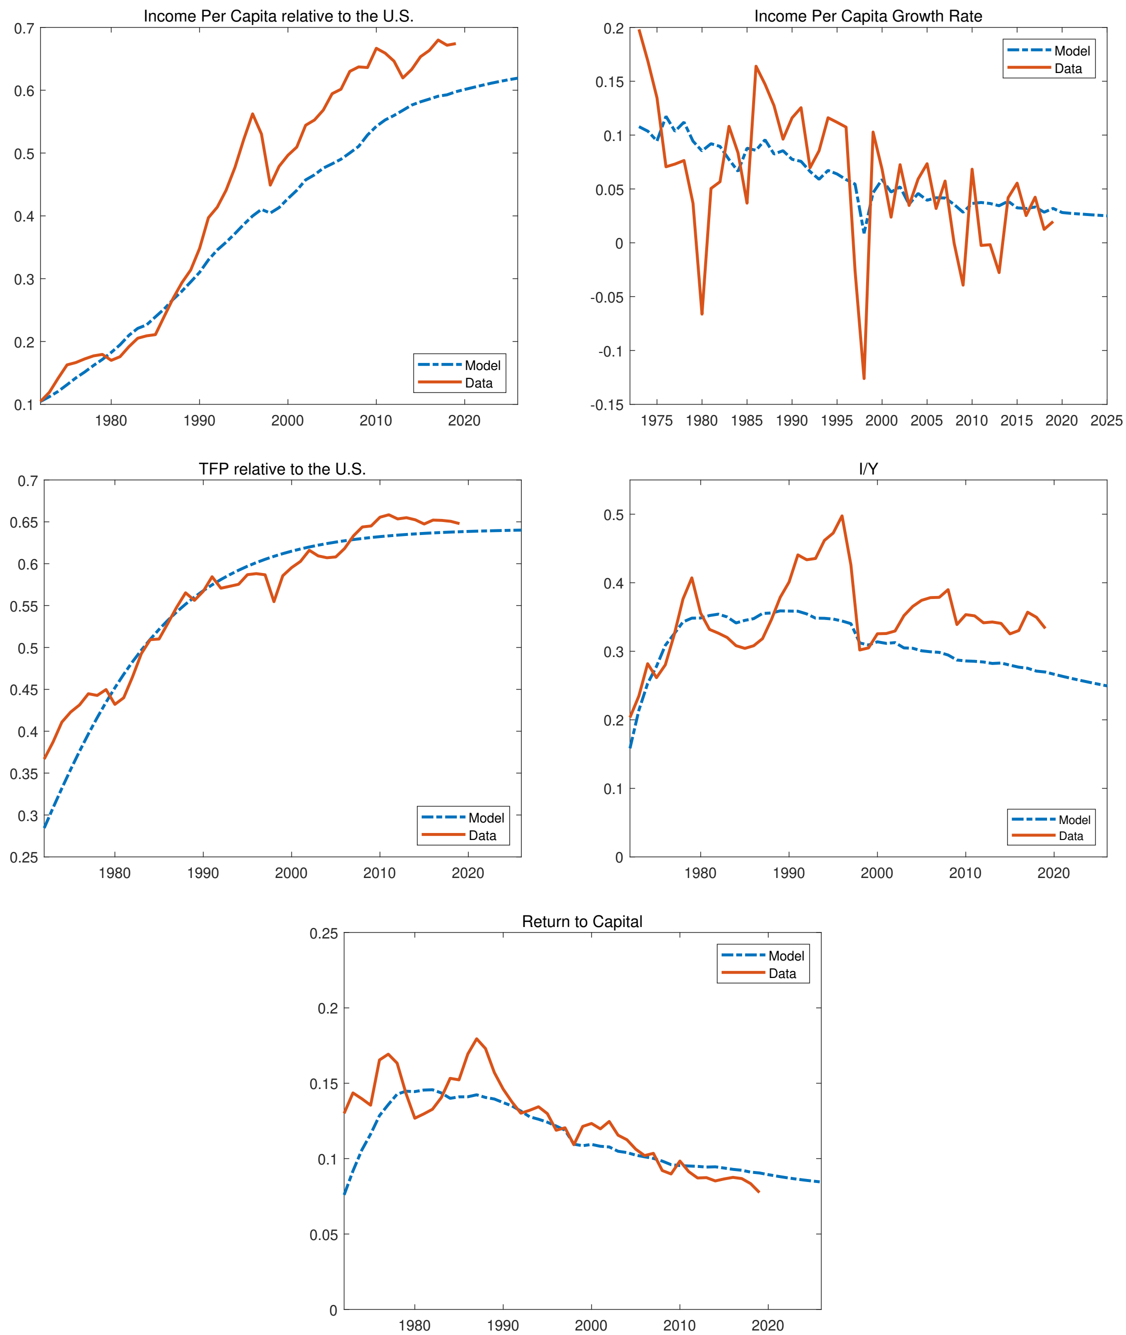

Figure 7.2: Korea: Transitional dynamics

Taiwan. Our next case study is Taiwan. The top left panel of Figure 7.3 shows that per capita income growth is faster in the data than in the model until the early 2000s, when convergence with the U.S. stops. After several decades of fast growth, Taiwan's growth has returned to levels below those forecasted by the model (top right panel of Figure 7.3). Interestingly, the model does a good job with investment without underestimating it (middle right panel of Figure 7.3).

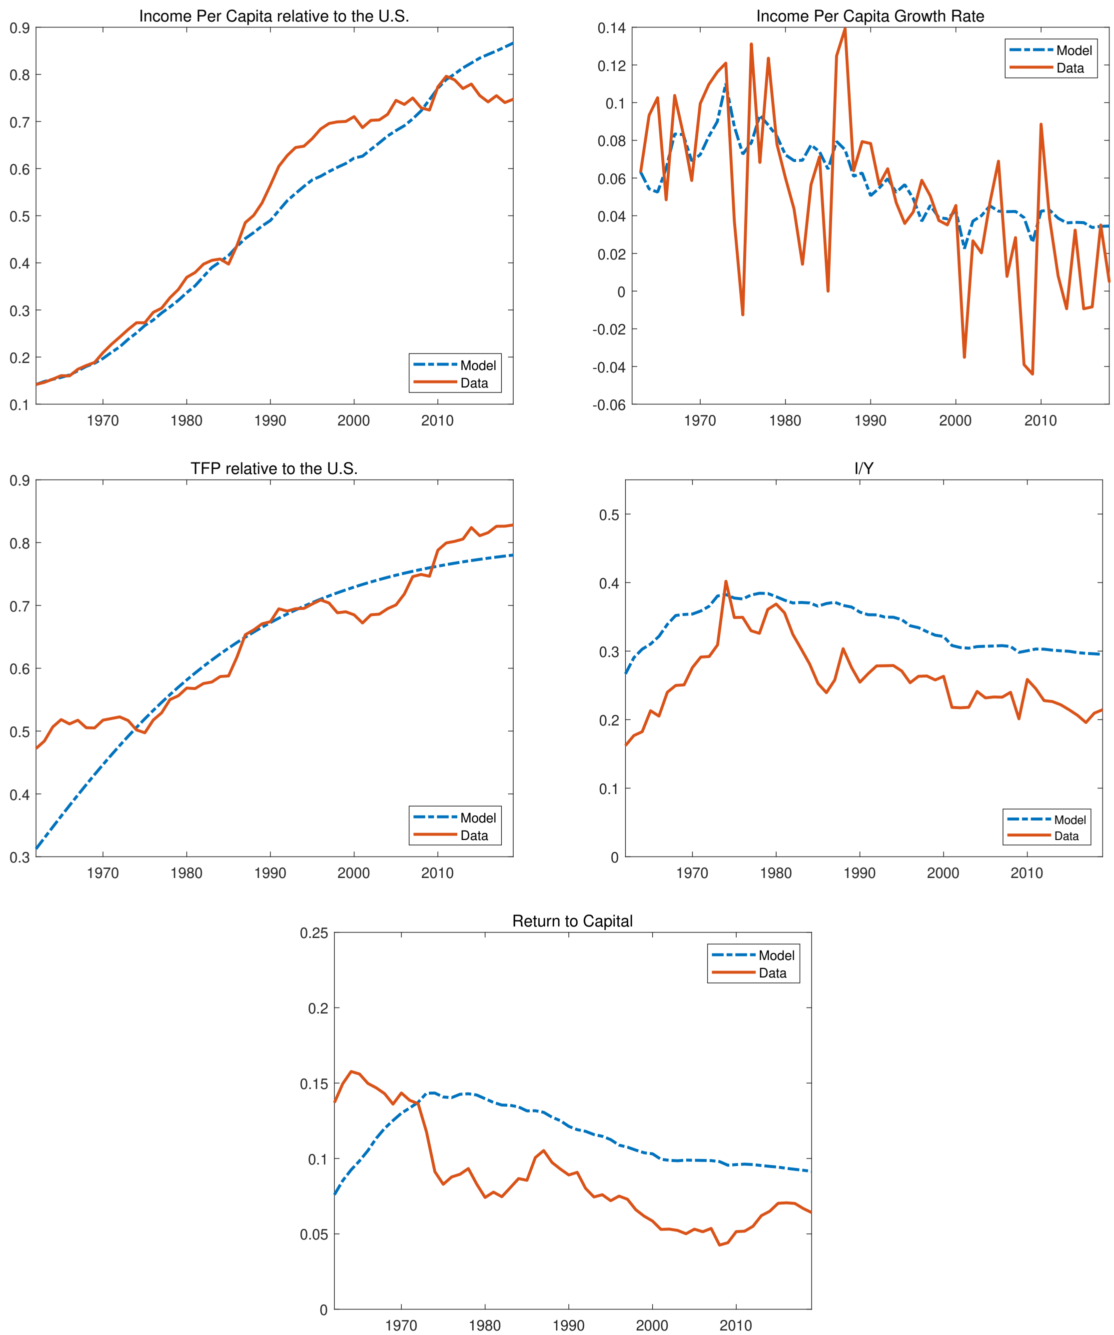

Figure 7.3: Taiwan: Transitional dynamics

Taking stock Our minimalist model accounts reasonably well for the experience of Japan, Korea, and Taiwan, particularly for the first 25 years of data in each economy. Later, wider divergences appear, with exceptionally fast growth in the late 1980s/early 1990s, a decade that some economists have referenced as reflecting a bubble, both in terms of growth exceeding the neoclassical benchmark, followed by a severe financial crisis. Interestingly, the growth above the model's predictions has been undone in the case of Japan and Taiwan, though it has yet to be undone in the case of Korea.

8 Summary and Conclusion

Since China became a middle-income country in 1995, its growth performance has been remarkably similar to that of other East Asian growth miracles at the same stage of development. We find that a neoclassical framework with a simple three-parameter TFP catch-up process accounts for the growth patterns in these economies. More broadly, we find that the many factors viewed as central to understanding these growth miracles, ranging from industrial policies to financial market modernization to trade and FDI policies, manifest themselves within a simple TFP catch-up process in a one-sector framework. We, therefore, view this paper as a complement to those papers in the literature that have developed detailed models with these mechanisms. We interpret our model's minimalism as a fundamental virtue; increased modeling complexity comes with empirical challenges, including the very difficult task of trying to project these mechanisms' future course, which is needed to forecast China's future growth.

China has been growing because it is accumulating capital and is catching up with the world's technological frontier. But China's technology catch-up is slowing considerably, both in absolute terms and relative to the other East Asian growth miracles. This fact, combined with a declining population, means that China's growth can only slow down, perhaps significantly. Our future forecast for China's economy, which assumes a continued comparatively high savings rate and employment-to-population ratio, might be too optimistic should either of these factors decline. On the other hand, our forecast may be too pessimistic if China implements new institutional reforms that improve economic efficiency to reverse its declining TFP catch-up rate.

References

Bai, C.-E., C.-T. Hsieh, and Y. Qian (2006): "The return to capital in China," *Brookings Papers on Economic Activity*, 2, 61–101.

Bai, C.-E., Q. Liu, and W. Yao (2020): "Earnings inequality and China's preferential lending policy," *Journal of Development Economics*, 145, 102477.

Barro, R. J. (2016): "Economic growth and convergence, applied to China," *China & World Economy*, 24, 5–19.

Brandt, L., C.-t. Hsieh, and X. Zhu (2008): "Growth and Structural Transformation in China," in *China's Great Economic Transformation*, ed. by L. Brandt and T. G. Rawski, Cambridge University Press, 683–728.

Brandt, L. and X. Zhu (2010): "Accounting for China's growth," Tech. rep., IZA.

Chang, C., K. Chen, D. F. Waggoner, and T. Zha (2016): "Trends and cycles in China's macroeconomy," *NBER Macroeconomics Annual*, 30, 1–84.

Chang, C., Z. Liu, M. M. Spiegel, and J. Zhang (2019): "Reserve requirements and optimal Chinese stabilization policy," *Journal of Monetary Economics*, 103, 33–51.

Chen, K., A. İmrohoroğlu, and S. İmrohoroğlu (2006): "The Japanese saving rate," *American Economic Review*, 96, 1850–1858.

Chen, K. and T. Zha (2023): "China's Macroeconomic Development: The Role of Gradualist Reforms," Working Paper 31395, National Bureau of Economic Research.

Chen, W., X. Chen, C. Hsieh, and Z. Song (2019): "A forensic examination of China's national accounts," *Brookings Papers on Economic Activity*, 1.

Christiano, L. (1989): "Understanding Japan's saving rate: the reconstruction hypothesis," *Quarterly Review*, 13, 10–25.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015): "The next generation of the Penn World Table," *American Economic Review*, 105, 3150–82.

Fernald, J. G., I. Malkin, and M. M. Spiegel (2013): "On the reliability of Chinese output figures," *FRBSF Economic Letter*.

Fukao, K. and T. Settsu (2021): "Japan: Modern economic growth in Asia," in *The Cambridge Economic History of the Modern World*, ed. by S. Broadberry and K. Fukao, Cambridge University Press, vol. 2, 100–128.

Gollin, D. (2002): "Getting income shares right," *Journal of Political Economy*, 110, 458–474.

Hall, R. E. and C. I. Jones (1999): "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quarterly Journal of Economics*, 114, 83–116.

Herd, R. (2020): "Estimating capital formation and capital stock by economic sector in China: The implications for productivity growth," Tech. rep., World Bank No 9317.

Jiang, S., J. Miao, and Y. Zhang (2022): "China's housing bubble, infrastructure investment, and economic growth," *International Economic Review*, 63, 1189–1237.

König, M., K. Storesletten, Z. Song, and F. Zilibotti (2022): "From imitation to innovation: Where is all that Chinese R&D going?" *Econometrica*, 90, 1615–1654.

McGrattan, E. R. and J. A. Schmitz (1999): "Maintenance and repair: too big to ignore," *Quarterly Review*, 23, 2–13.

Naughton, B. J. (2018): *The Chinese Economy: Adaptation and Growth*, MIT Press, 2nd ed.

Parente, S. L. and E. C. Prescott (1994): "Barriers to technology adoption and development," *Journal of Political Economy*, 102, 298–321.

Pritchett, L. and L. H. Summers (2013): "Asiaphoria meets regression to the mean," *Proceedings, Federal Reserve Bank of San Francisco*, 1–35.

Song, Z., K. Storesletten, and F. Zilibotti (2011): "Growing like China," *American Economic Review*, 101, 196–233.

Whelan, K. (2002): "A guide to US chain aggregated NIPA data," *Review of Income and Wealth*, 48, 217–233.

Wu, H. X. (2014): "China's growth and productivity performance debt revisited - Accounting for China's source of growth with a new data set," Tech. rep., The Conference Board.

Xiong, W. (2019): "The Mandarin model of growth," Mimeo.

Yao, W. and X. Zhu (2021): "Structural change and aggregate employment fluctuations in China," *International Economic Review*, 62, 65–100.

Young, A. (1992): "A tale of two cities: Factor accumulation and technical change in Hong Kong and Singapore," in *NBER Macroeconomics Annual 1992, Volume 7*, National Bureau of Economic Research, 13–64.

——— (1995): "The tyranny of numbers: Confronting the statistical realities of the East Asian growth experience," *Quarterly Journal of Economics*, 110, 641–680.

Zhu, X. (2012): "Understanding China's growth: Past, present, and future," *Journal of Economic Perspectives*, 26, 103–24.

Appendix

The following appendices provide further details about the data employed in the paper and offer additional information regarding Hong Kong and Singapore and additional experiments.

A TFP Construction

We use two different TFP measurements from PWT 10.0 to construct our TFP:

- $RTFP^{NA}$:
  - TFP at constant national prices (2017=1), constructed using $RGDP^{NA}$.
  - $RGDP^{NA}$ is based on the national account GDP growth rate.
  - PWT reports the growth rate of $RTFP^{NA}$, normalizes the $RTFP^{NA}$ to 1 in 2017 for all economies, and backs out the TFP level for other years for each economy.
- $CTFP$:
  - TFP level at current PPPs (USA=1), constructed using $CGDP^{O}$.
  - PWT computes each economy's $CTFP$ level and uses U.S. $CTFP$ to normalize them.

We construct China's real TFP relative to the U.S., $\frac{TFP_{CN,t}}{TFP_{US,t}}$, as follows:

- Take the value of $CTFP_{CN,2017}$ and assume $TFP_{CN,2017} = CTFP_{CN,2017}$.
- The relative TFP in any year, take 2019, for example, is

$$\begin{aligned}\frac{TFP_{CN,2019}}{TFP_{US,2019}} &= \frac{CTFP_{CN,2017}}{CTFP_{US,2017}}\frac{RTFP^{NA}_{CN,2019}}{RTFP^{NA}_{US,2019}} \\ &= CTFP_{CN,2017}\frac{RTFP^{NA}_{CN,2019}}{RTFP^{NA}_{US,2019}}.\end{aligned}$$

Note that $CTFP_{US,2017} = 1$.

The constructed TFP is plotted in Figure 5.1.

B Different Measures of China's per Capita GDP

As discussed in the main text, there are two GDP series for China in PWT 10.0: one is based on GDP growth adjusted by Wu (2014), denoted by "adjusted" in Figure B.1, and the other is based on the GDP growth rates from China's official national accounts. Figure B.1 plots the differences in levels (left panel) and growth rates (right panel).

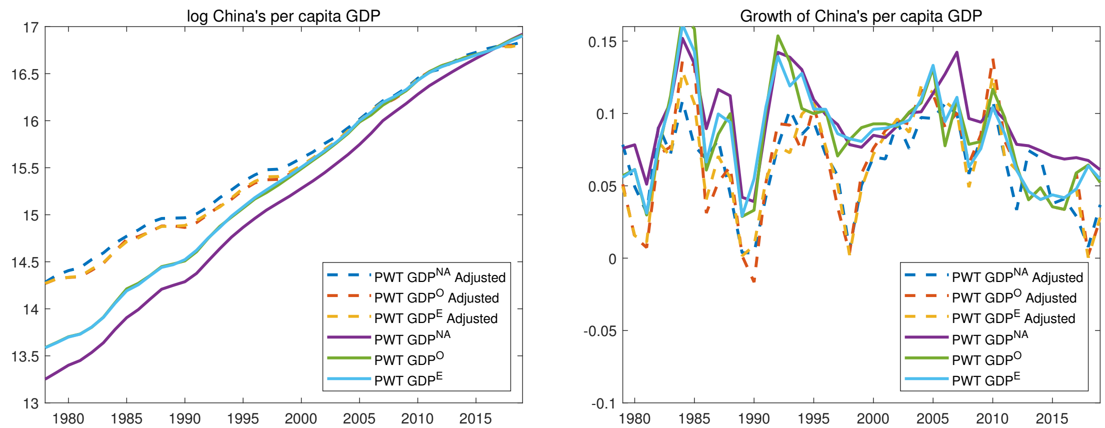

Figure B.1: GDP per capita: log-level and growth rate

We also report the investment rates from China's National Bureau of Statistics and compare them with those from PWT 10.0 in Figure B.2.[19]

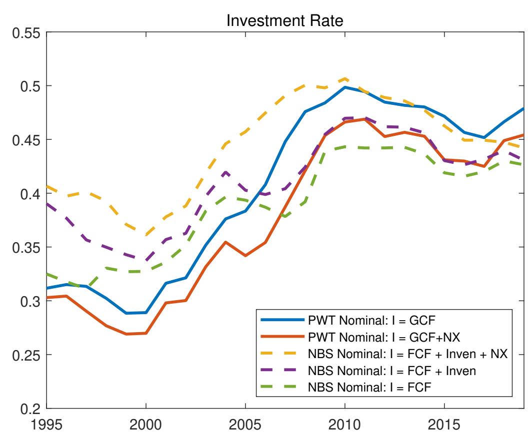

Figure B.2

[19]PWT denotes the investment rate reported by PWT 10.0, where GCF is gross capital formation, and NX is net exports. NBS denotes the data from the National Bureau of Statistics in China, where FCF is the fixed capital formation, and Inven denotes changes in inventories. The investment rates are in nominal ratios, which avoids the need to compute relative prices.

C Facts including Hong Kong and Singapore

We exclude Hong Kong and Singapore from our benchmark facts because their population are relatively small. We add them now. Hong Kong's first year in the data is 1960, with a GDP per capita level of $6004. China's GDP per capita level reached $5953 in 2004. Hence, we plot the figure from 2004 below.

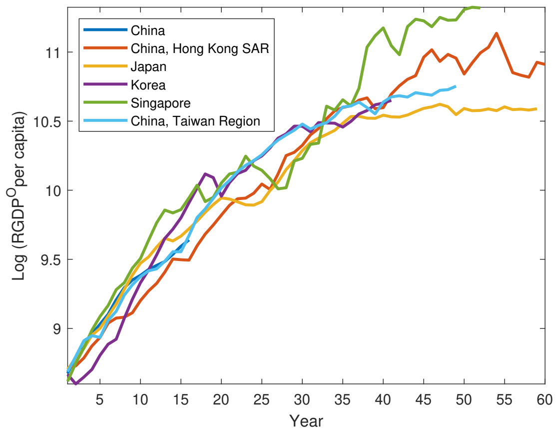

Figure C.1: GDP per capita, normalized at China's GDP per capita in 2004

Figure C.1 shows that the fact on output growth is robust if we include these two cities. Hong Kong and Singapore performed well during the early 2000s, leaving the other East Asian economies behind, but their growth rate eventually slowed down, as in all the other economies.

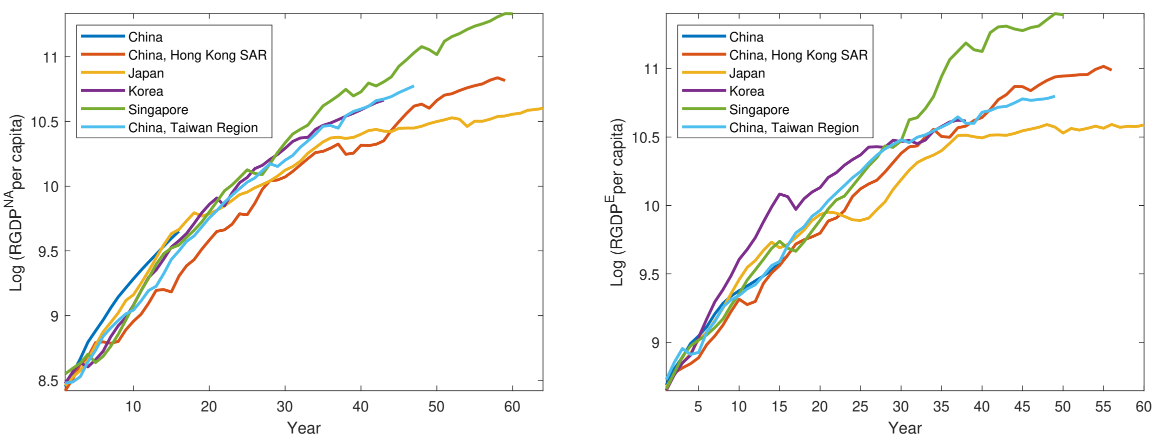

Figure C.2: GDP per capita, alternative definitions

Similarly, alternative measures of GDP (Figure C.2) do not change the picture much.

Figure C.3 plots the investment rates. Hong Kong's and Singapore's investment rates differ from those of other East Asian economies. Look at Singapore's case, where investment reached levels as high as 70% in terms of gross capital formation over GDP, but with negative investment discounting net exports.

However, this is not a surprise. Singapore was the recipient of large foreign direct investment (FDI) flows. Given the size of its economy earlier in the sample, a few large FDI projects can induce wild fluctuations in measured investment, as shown in Figure C.3. Recall, for instance, that Mobil started building a large oil refinery in Jurong (Pioneer Road) in 1966, while Esso started another one at Pulau Ayer Chawan (now part of Jurong Island) in 1970, making Singapore a refinery hub for the region. These observations reinforce our choice of leaving Hong Kong and Singapore out of the main text's discussion.

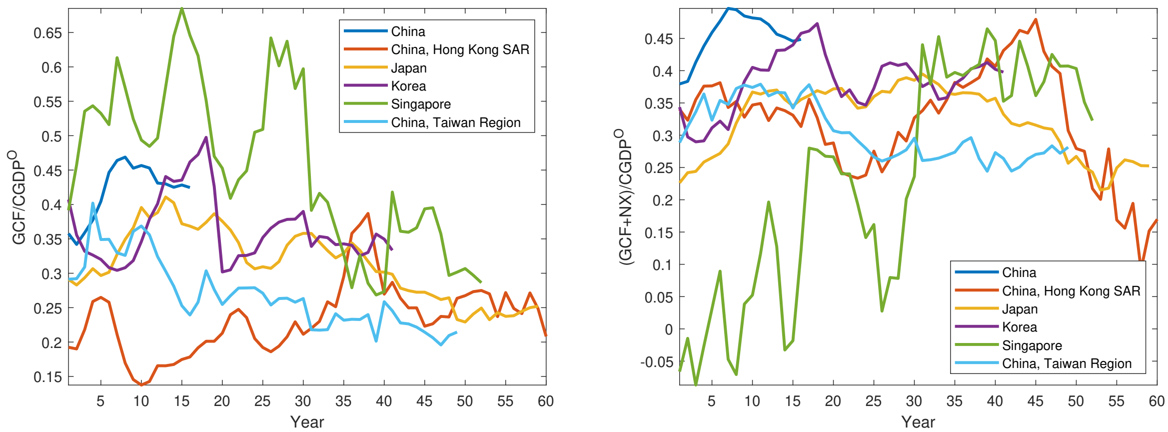

Figure C.3: Investment rates

D Growth Accounting

An interesting complement to our quantitative experiments is undertaking a standard growth accounting exercise for China. This exercise shows that, indeed, TFP growth is a key driver of China's GDP growth and, therefore, that as TFP growth is bound to slow down, so is China's GDP growth.

Let us assume a production function of the form $Y_t = K_t^{\theta}(A_t L_t)^{1-\theta}$, where $TFP = A_t^{1-\theta}$. Following Hall and Jones (1999), we can decompose GDP per capita into:

$$\frac{Y_t}{pop_t} = A_t \left(\frac{K_t}{Y_t}\right)^{\frac{\theta}{1-\theta}} \frac{L_t}{pop_t}$$

where pop_t is total population. Hence, the growth of GDP per capita can be expressed as:

$$\begin{aligned} \Delta\left(\frac{Y_t}{pop_t}\right) &= \Delta A_t + \frac{\theta}{1-\theta}\Delta\left(\frac{K_t}{Y_t}\right) + \Delta\left(\frac{L_t}{pop_t}\right) \\ &= \frac{1}{1-\theta}\Delta TFP_t + \frac{\theta}{1-\theta}\Delta\left(\frac{K_t}{Y_t}\right) + \Delta\left(\frac{L_t}{pop_t}\right) \end{aligned}$$

where Δ denotes a growth rate.

Table D.1 shows the growth accounting exercises with the national account approach TFP, capital, and GDP ($RGDP^{NA}$). Since the contribution of total factor productivity growth is weighted by $1/(1-\theta)$, the growth contribution of TFP is $2 \times 3.11/8.30 = 75\%$ of the growth in GDP per capita.

Table D.1: Growth accounting in the data, national accounting approach

	Data			
	GDP per capita ($RGDP^{NA}$)	TFP	Capital/output ratio	Emp/pop ratio
Aver. annual growth rate	8.30	3.11	2.09	-0.05
Cont. to per capita GDP growth	100.00	74.83	25.21	-0.56

Unfortunately, data limitations preclude us from undertaking this growth accounting using production-based GDP, the concept that we use in our quantitative model, in a consistent way (see our discussion in Section 2 of the main paper).

E Discussion of Labor Income Shares

E.1 Labor Shares from the PWT 10.0

Computing the labor share from income data is challenging because we need to calculate the labor compensation of self-employed workers, which includes both a labor and a capital income component. Gollin (2002) discusses possible treatments of self-employed income. PWT 10.0 estimates the labor share in the following five ways, among which Adjustment 1-3 is from Gollin (2002):[20]

- **Naive share**: Labor income only includes labor compensation of employees, which does not include self-employed income.

- **Adjustment 1**: Treat all self-employed income as labor compensation.

- **Adjustment 2**: Self-employed workers use labor and capital in the same proportion as the rest of the economy. Although adjustment 2 is considered the best measure, only Japan in our sample reports mixed income.

- **Adjustment 3**: The self-employed earn the same average wage as employees, but this greatly overestimates labor share in East Asian economies, where self-employed workers are mostly low-skilled. PWT 10.0 considers adjustment 3 as the best measure for Taiwan.

[20] See the details in Feenstra et al., 2015, Appendix C.

- **Adjustment 4**: Assume that all self-employed work takes place in the agricultural sector. Thus, the entire value added in agriculture is added to labor compensation. PWT 10.0 considers Adjustment 4 the best measure for Korea, Hong Kong, and Singapore.

Table E.2 shows the labor shares PWT 10.0 constructs for each economy. PWT 10.0 chooses what it considers the best measure of labor income in each economy from the five adjustments, and uses it to construct the TFP.

Table E.2: Labor shares reported by PWT 10.0

Economy	Sample Period	Different Adjustment Methods					Best Labor Share	
		Naive (1)	Adj 1 (2)	Adj 2 (3)	Adj 3 (4)	Adj 4 (5)	Method	Value
U.S.	1995-2019	0.56	0.64	0.61	0.6	0.57	Adj 2	0.61
China	1995-2019	0.58			0.61	0.7	Naive	0.58
Japan	1950-2019	0.51	0.63	0.6	0.68	0.54	Adj 2	0.6
Korea	1972-2019	0.46				0.55	Adj 4	0.55
Taiwan	1962-2019	0.53			0.73		Adj 3	0.73
Hong Kong	1960-2019	0.48			0.54	0.48	Adj 4	0.48
Singapore	1960-2019	0.42			0.52	0.43	Adj 4	0.43

Figure E.1 plots the different labor shares from PWT 10.0. In a blog entry, Restrepo and Reinbold also report Adjustment 4 for Taiwan, which is very similar to the naive share.[21] However, we did not find such information in PWT 10.0 to replicate their computations in a model-consistent way.

E.2 Discussion of Adjustment of Labor Income Shares

China: China's National Bureau of Statistics provides annual data on the labor share for each province but not for the aggregate economy. Bai et al. (2006) construct the aggregate labor share as the average of the provincial labor shares weighted by each province's GDP share. This gives an average labor income share of 0.5.

Before 2004, the National Bureau of Statistics counted all self-employment income as labor income, corresponding to Adjustment 1 by PWT 10.0's definition. Thus, the estimates of 0.5 from Bai et al. (2006) overstate the true labor share and understate the true capital share. In 2004, the National Bureau of Statistics, for the first time, explicitly excluded the imputed capital income of self-employed workers from the estimates of labor income. Unfortunately, the National Bureau of Statistics does not report the magnitude of this adjustment, and we cannot adjust our labor share estimates for the years before 2004.

[21] https://www.stlouisfed.org/on-the-economy/2018/january/measuring-labor-share-asian-tigers

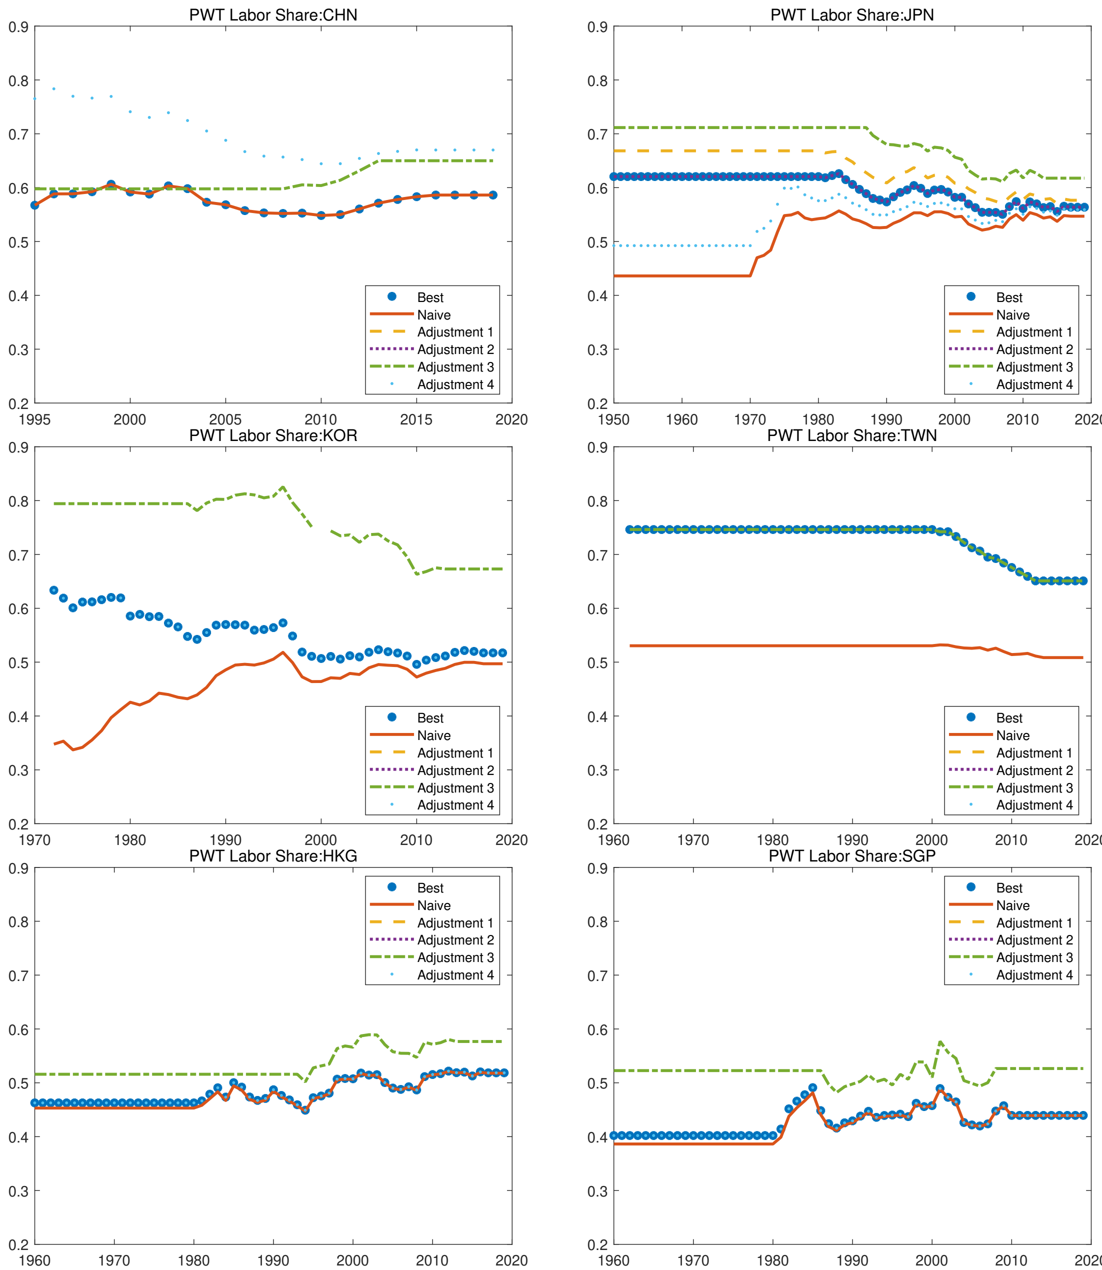

Figure E.1: Labor shares reported by PWT 10.0

PWT 10.0 uses the naive labor share (0.58) for China as the best labor share, which means that labor income only includes the labor compensation of employees and does not include self-employed income at all. However, as explained above, this number is not reported in China's official data; it is unclear whether this is a label mistake or some adjustment has been made.

Taiwan: PWT 10.0 uses Adjustment 3 for Taiwan as the best labor share, which gives a labor share of 0.73. The "same-wage" assumption may not be too far off the mark in advanced economies, where the share of employees in the total number of persons engaged (employees + self-employed) is 85-95%. However, this share is below 50% in many emerging economies and as low as 4%. In those economies, using the information on the wages of employees will overstate the labor income of the self-employed. Hence, the labor share of Taiwan is likely to be too high.

Singapore: PWT 10.0 follows Adjustment 4 by adding the entire value added in agriculture to labor compensation, assuming that all of the self-employed work in Singapore takes place in agriculture. This gives a labor share of 0.43, which is still too low as the labor income of the self-employed outside agriculture is largely ignored. Moreover, the effect of such an adjustment is minimal. The adjusted labor share is very close to the labor share before adjustment, which is 0.42. In the exercises in Subsection F.4, we pick a more conventional value for labor income share of 0.6, nearly the same as the one in the U.S.

F Transitional Dynamics in China, Taiwan, Hong Kong, and Singapore

F.1 China: Alternative capital income share

In this section, we set China's capital income share to 0.42, the value from PWT 10.0. Figure F.1 shows the new transitional dynamics. Now the model delivers income growth that is a bit lower than in the data (top left panel) but matches the observed TFP much better (bottom left panel). In comparison, the model does a much worse job matching the observed investment rates. The explanation is transparent: with $\theta = 0.42$, the production function is more concave in capital. Thus, the social planner does not want to invest that much in capital, and the resulting growth is lower.

F.2 Taiwan

In this section, we show the results for Taiwan when we use the capital income share directly from PWT 10.0. Figure F.2 shows that the model hugely underestimates observed growth, as the marginal productivity of capital in the model falls very rapidly.

Figure F.2 justifies our choice of disregarding the value of θ reported in PWT 10.0 for Taiwan: the behavior that it would imply for growth in our model (or, for that matter, any other commonly used growth model) is sharply counterfactual.

Figure F.1: China: Transitional dynamics with $\theta = 0.42$

F.3 Hong Kong

Our next set of results deals with Hong Kong, where we use the capital share, $\theta = 0.51$, directly from PWT 10.0. In the top left panel of Figure F.3, the model forecasts well Hong Kong's income per capita, although there is a growing gap at the end of the sample, where it seems that the economy of the city is falling behind what one would expect. The model also does a good job matching TFP relative to the U.S.

The main divergence between the data and the model is in investment, which seems much lower in the data than in the model. However, since Hong Kong is heavily dependent on trade and services, our model might need to be modified to capture all the peculiarities of Hong Kong's data on investment. For example, investments in factories close to but not in Hong Kong (i.e., the Chinese hinterland of Hong Kong) might not be properly recorded in Hong Kong's statistics.

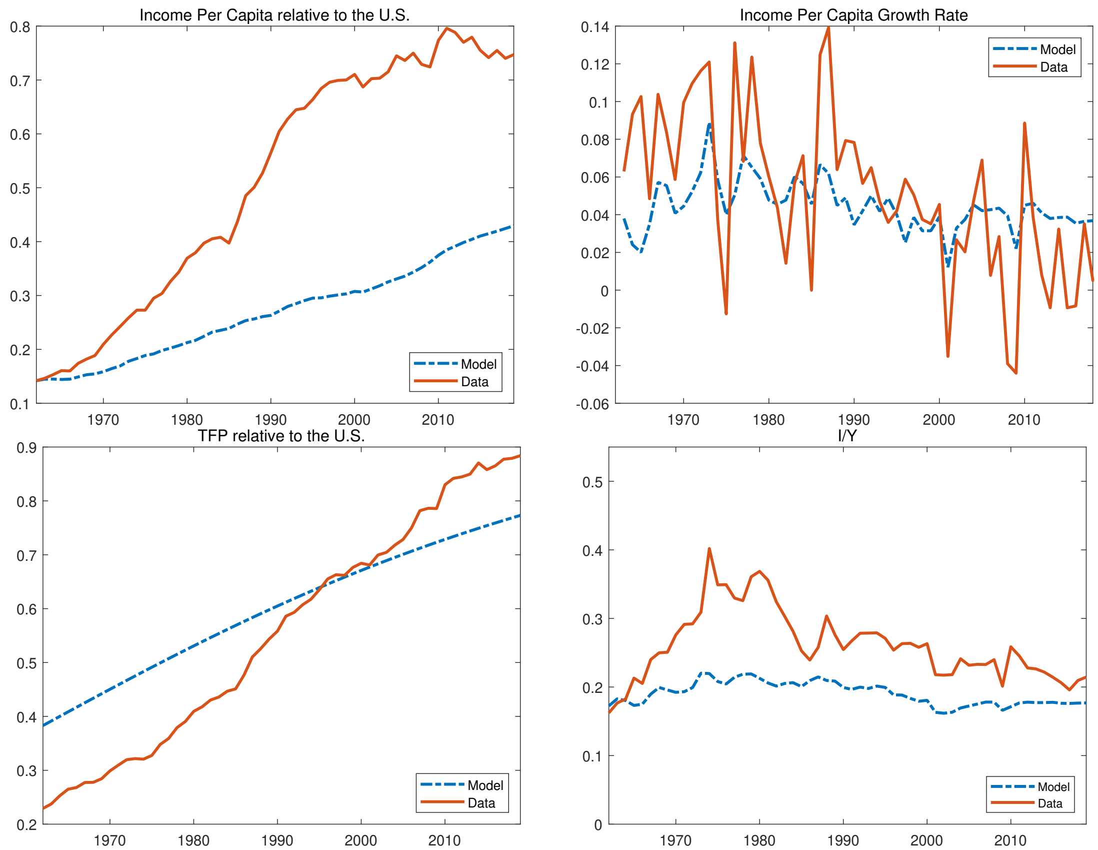

Figure F.2: Taiwan: Transitional dynamics with $\theta = 0.27$

F.4 Singapore

We close our investigation with Singapore, with results in Figure F.4. As in previous cases, the model does a fair job of matching income per capita (top left panel and TFP relative to the U.S. (bottom left panel). In the case of income per capita, Singapore outperforms the model from around 2000 to 2010. However, after 2010, growth relative to the U.S. stopped. Strikingly, Singapore seems to have converged very quickly to 90% of the U.S. TFP (with convergence to this level completed by the mid-1980s). See Young (1995) for an early paper that noticed this pattern.

Also, the model does not do a great job with investment (bottom right panel), at least until the early 2000s. In contrast with the Hong Kong case, now the model underestimates the extent of investment in the data for many years. Young (1992) documented, several decades ago, some of the big differences between the growth process of these two economies despite their apparent similarities.

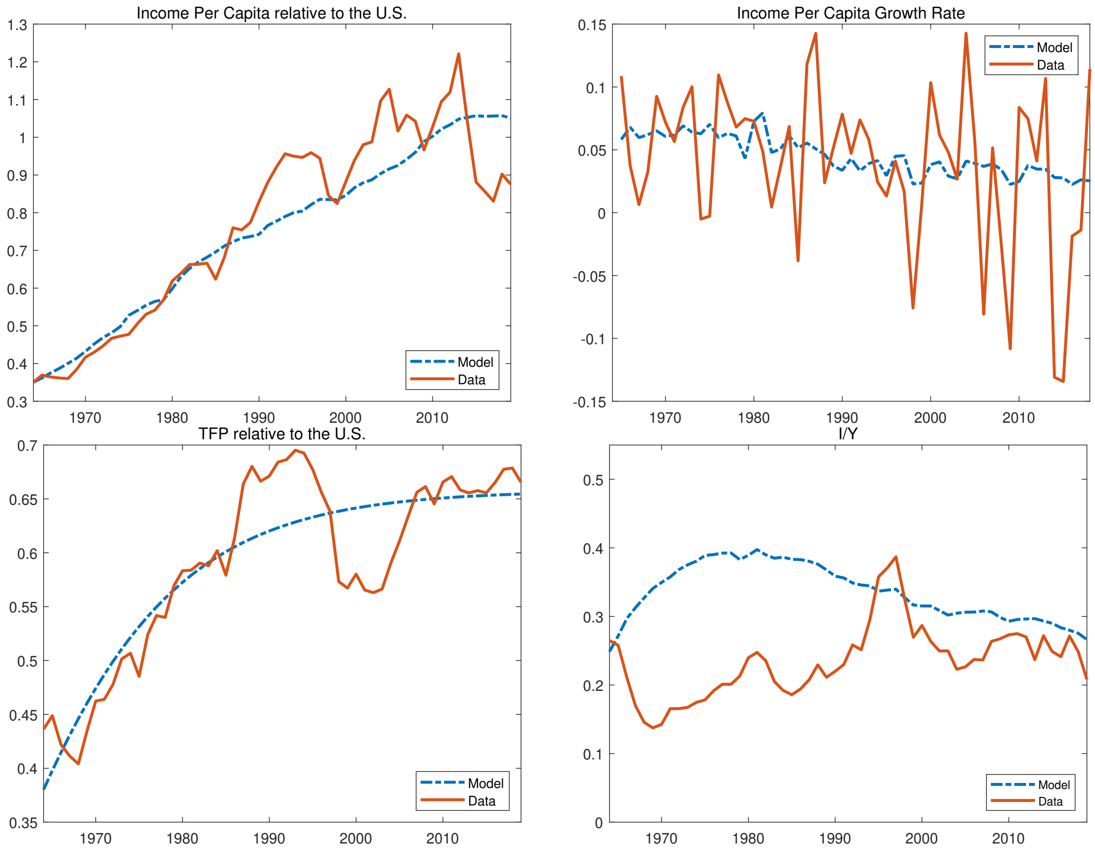

Figure F.3: Hong Kong: Transitional dynamics

F.5 Taking stock

To sum up, our minimalist model has more difficulties in the cases of Hong Kong and Singapore, but even for these two economies, the model can capture the slowdown in growth that we forecast for China.

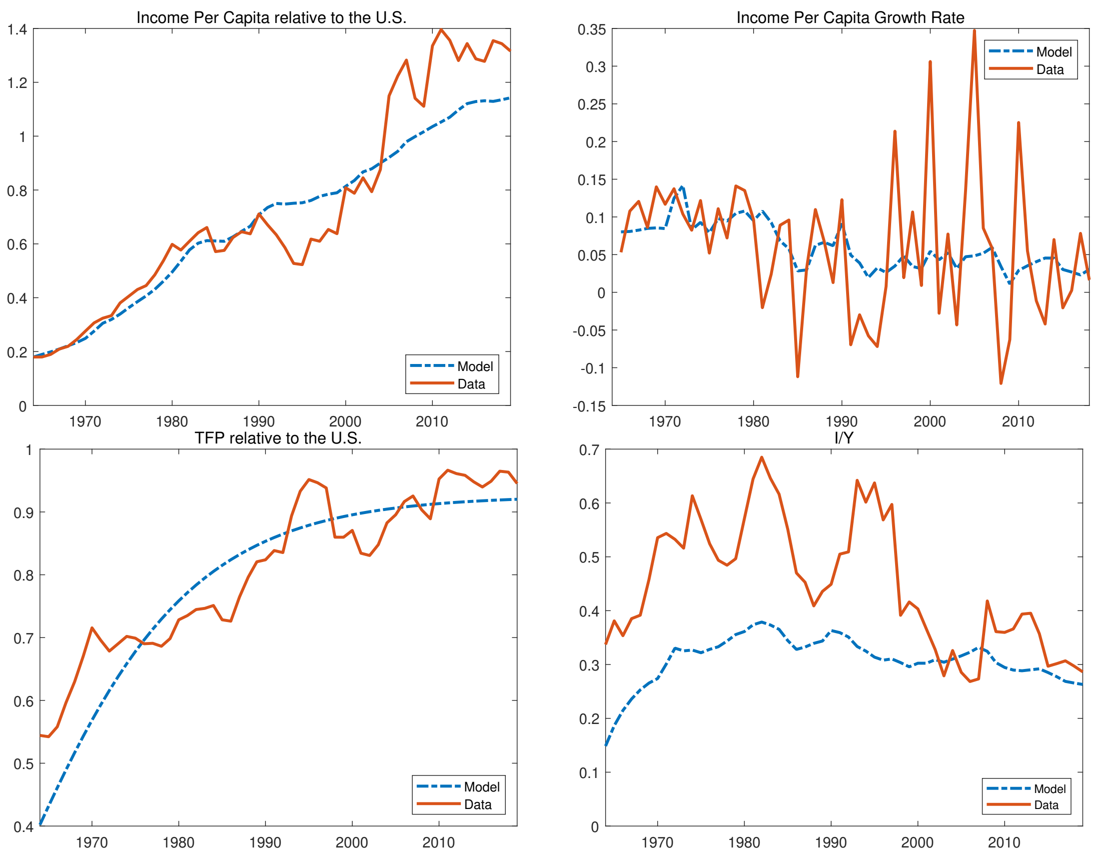

Figure F.4: Singapore: Transitional dynamics $\theta = 0.4$