Innovation, Diffusion, and Misallocation: Understanding China's Recent Productivity Dynamics*

Emin Gahramanov† Yan Liang‡ Xueli Tang§ Chengsi Zhang¶

April 2026

Abstract

Since the 2008 financial crisis, China has experienced a marked slowdown in productivity growth despite exceptionally strong innovation activity. We study this puzzle using an endogenous growth model calibrated to the Chinese economy over 2007–2017. The analysis identified two central forces. First, innovation has become increasingly costly as China approaches the technological frontier. Second, intensified involutionary competition diverts innovation effort towards activities with limited social gains. Other factors, including weaker knowledge diffusion and greater resource misallocation, play a minor role. While industry policies support innovation, they are insufficient to offset these broader structural forces. Sustaining productivity growth therefore requires improving efficiency and direction of innovation activity, particularly by curbing involutionary competition.

JEL Codes: E22, L13, O31, O33.
Keywords: Productivity slowdown, China, involutionary competition, innovation costs, knowledge diffusion, misallocation.

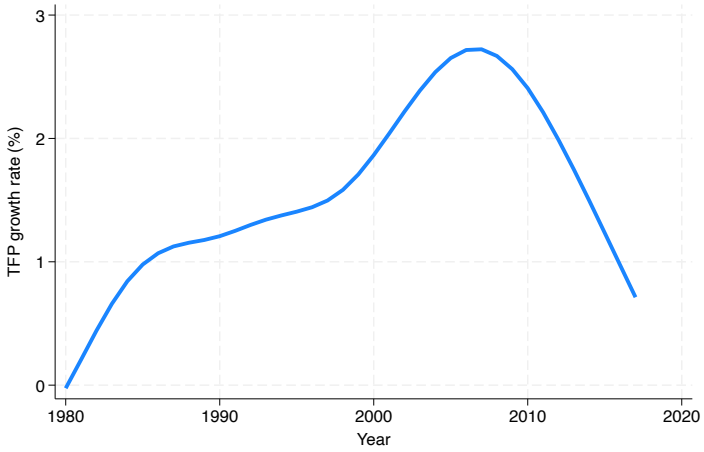
*We are grateful to Xiaobo Zhang for helpful discussions and insightful comments. We also thank conference participants at the 35th Annual Conference of the Chinese Economics Society Australia (CESA) for their constructive feedback. All remaining errors are our own.

†American University of Sharjah; egahramanov@aus.edu.
‡Deakin University; yan.liang@deakin.edu.au.
§Deakin University; xueli.tang@deakin.edu.au.
¶Renmin University of China; zhangcs@ruc.edu.cn

1 Introduction

The Chinese economy has experienced remarkable growth since the start of the reform and opening-up policies in 1978, with particularly strong expansion following the country's accession to the WTO in 2001. More recently, however, economic growth has slowed, with the slowdown in total factor productivity (TFP) growth being especially pronounced (Figure 1). The decline in productivity growth has become a central concern for understanding China's long-run economic prospects.

Figure 1: Trend in TFP Growth



Notes. Long-run trend in China's total factor productivity (TFP) growth, computed from Penn World Table (PWT) 9.1 using the Hodrick–Prescott (HP) filter. We use this version to maintain internal consistency across the empirical and calibration components of the paper.

What makes China's experience particularly striking is that the slowdown in productivity has occurred alongside extraordinarily strong research and development (R&D) activity. R&D expenditure as a share of GDP has risen steadily since the late 1990s, and the measure of research output, such as patent applications per million people, have increased rapidly during the same period (Figure 2). These patterns suggest that the slowdown in TFP growth cannot be attributed to a lack of innovation effort. Instead, they point to a disconnect between the scale of innovative activity and its contribution to productivity growth.

This raises a natural question: What factors could have driven China's productivity slowdown despite sustained increases in both innovation inputs and outputs? More broadly, what do these patterns imply about the effectiveness of existing innovation and industrial policies?

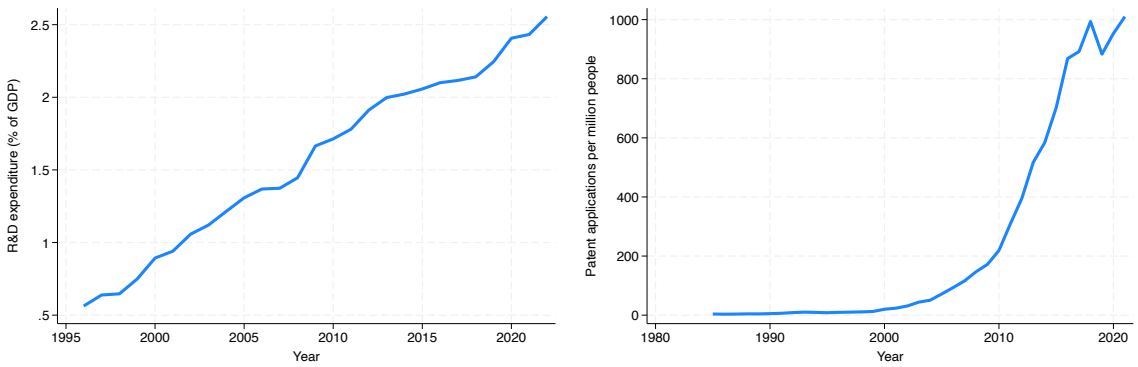
To address these questions, we build on the endogenous growth framework with step-by-step innovation of Akcigit and Ates (2023), which was originally developed to study the decline of business dynamism in the US since the 1980s. In this framework, firms invest in innovation to improve their technological position, and aggregate productivity growth arises from the accumulation of these incremental advances. We extend this framework along several extensions to better capture China's growth experience.

First, we incorporate resource misallocation following Hsieh and Klenow (2009). Such misallocation diverts resources away from technologically more advanced firms. By weakening the profitability of technological leaders, misallocation dampens their incentives to innovate.

Figure 2: Research and Development Activity

(a) R&D Expenditure

(b) Annual Patent Applications



Notes. Research and development (R&D) activity in China. The left panel shows R&D expenditure as a share of Gross Domestic Product (GDP); the right panel shows annual patent applications per million people. Data are from Our World in Data.

Because potential entrants and follower firms are forward-looking, lower expected returns to leadership further discourage innovation more broadly. As a result, resource misallocation reduces both the level and growth of TFP.

Second, and most importantly, we introduce involuntary competition ("*neijuan*") into our framework. Involution refers to a form of excessive competition over a fixed set of opportunities, where individual activities yield private returns but generate little or no social gains. In the context of innovation, this manifests as R&D efforts aimed at acquiring or consolidating market positions, capturing short-term rents, rather than producing genuine technological progress. We model this as a diversion of innovative effort towards socially unproductive activities, which reduces the effectiveness of aggregate R&D in generating productivity growth.

Third, we allow incumbent firms to engage in defensive R&D to protect their market positions and limit knowledge spillovers to competitors. This introduces a strategic dimension to innovation, whereby firms may allocate resources towards restricting diffusion rather than advancing the technological frontier. Finally, we incorporate physical capital into this framework to reflect China's high investment share in GDP.

We discipline the model using Chinese data and use it to quantify the relative importance of different mechanisms in accounting for the observed slowdown in TFP growth. The analysis highlights several key findings.

First and most importantly, involutionary competition emerges as the primary driver of China's productivity slowdown. A substantial shared of R&D resources is diverted toward activities that generate little productivity improvement, which helps explain the coexistence of strong research activity and week TFP growth seen in Figures 1–2. This mechanism captures a form of inefficiency in the direction of innovation, rather than its scale. It accounts for more than half of the TFP slowdown between 2007 and 2017.

Moreover, rising innovation costs also play a significant role. As China approaches the world technology frontier, further innovations become intrinsically more difficult and costly. This reduces both incumbent innovation and the incentives for new firms to enter the market, thereby lowering the pace of aggregate productivity growth. Rising innovation costs explain

roughly one fifth of the productivity slowdown.

Furthermore, resource misallocation also contributes to China's TFP slowdown. Although its qualitative impact is much smaller than those of involutionary competition and rising innovation costs, it nonetheless represents a meaningful drag on aggregate productivity growth.

By contrast, the slowdown in knowledge diffusion plays a relatively limited role in China's experience, contributing only about one percent to China's TFP slowdown. This stands in sharp contrast to the US experience, where a slowdown in knowledge diffusion has been identified as a major driver of declining business dynamism.

These findings have implications for policy. Measures such as reductions in corporate income tax and increases in R&D subsidies are directionally appropriate, as they support innovation incentives and research activity, but their quantitative impact appears insufficient to offset the deeper structural forces. In particular, these policies are less effective in addressing the rise of involutionary competition and the increasing difficulty of innovation near the technology frontier. Thus, a key priority for policymakers should be the design of policies aimed at improving the efficiency and direction of innovative activity, rather than those that focus solely on increasing its scale.

This paper contributes to several strands of literature. First, it relates to the extensive literature on endogenous growth, particularly those featuring step-by-step innovations (e.g., Acemoglu and Akcigit, 2012; Liu et al., 2022; Akcigit and Ates, 2023). Most closely related is Akcigit and Ates (2023), who investigated the decline in business dynamism in the US. We extend their framework by incorporating involutionary competition, resource misallocation, and strategic defensive R&D, thereby providing a richer model to better capture China's growth experience. Our findings highlight that the drivers of growth slowdown differ across countries: while slower knowledge diffusion is central in the US, China's slowdown is primarily driven by the misdirection of innovative effort and rising innovation costs.

Second, our work relates to the growing literature on China's growth experience, particularly empirical studies using detailed firm-level data (e.g., Brandt et al., 2012, 2023; Cerdeiro and Ruane, 2024; Barwick et al., Forthcoming). For instance, Brandt et al. (2012) documents that firm entry contributed substantially to China's remarkable growth during 1998–2007. Subsequent research (Brandt et al., 2023) shows that the contribution of new entrants declined markedly between 2008 and 2013, partly explaining the post-global financial crisis slowdown. We contribute to this literature by offering a unified framework that connects firm-level behaviour with aggregate dynamics. In doing so, it complements existing empirical work by providing a structural interpretation of recent productivity trends.

Third, our paper is also related to the literature on resource misallocation and growth (e.g., Acemoglu et al., 2018; König et al., 2022; Aghion et al., 2025). Studies in this area focus on different dimensions of misallocation. For example, Acemoglu et al. (2018) examines the misallocation of resources between R&D activities and overhead production costs, while Aghion et al. (2025) investigates R&D resource misallocation between highly innovative firms and firms that are more efficient in production but less capable of innovation. Our paper contributes to this literature by highlighting how distortions that reduce the profitability of leading firms can have broader equilibrium effects by discouraging innovation incentives across the entire

distribution of firms.

The remainder of the paper is structured as follows. Section 2 presents the stylised facts. Section 3 outlines the theoretical framework. Section 4 reports quantitative results. Section 5 concludes.

2 Stylised Facts

Before developing the theoretical model, we document several stylised facts about business dynamism in the Chinese economy. We use the China Stock Market & Accounting Research Database (CSMAR), which covers all publicly listed firms since 1990.[1] We begin with a brief overview of these facts:

1. Market concentration exhibits an inverted U-shape, rising until around 2007 and then declining afterward, mirroring the long-run TFP growth pattern shown in Figure 1. This contrasts with the US, where market concentration has increased steadily since the 1980s.

2. Aggregate markups have risen over time, though the increase is relatively mild compared with the US over the same period.

3. Profit shares of GDP have increased over time, but again, the rise is milder than in the US.

4. Dispersion of firm growth rates has declined over time, similar to patterns observed in the US.

5. Productivity dispersion has increased, and the productivity gap between frontier and laggard firms has widened, a pattern also observed in the US during the same period.

6. Entry rates—the rate at which new firms get listed—have declined slightly, but the sales-weighted entry rate, reflecting the market share captured by newly listed firms, has dropped sharply over this period.

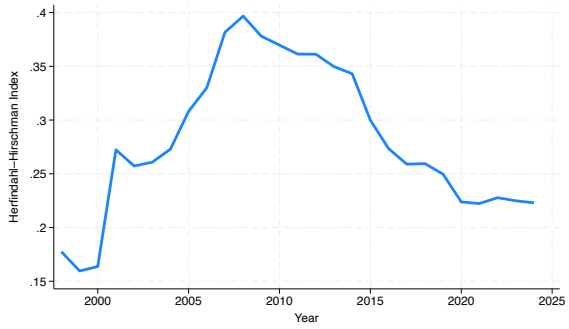
We then examine these patterns in detail below.

First, we find that market concentration has followed an inverted U-shape (Figure 3). It increased from 1998 until the 2007 financial crisis, after which it began to decline and has continued to do so ever since. Market concentration is measured using the Herfindahl–Hirschman Index (HHI). We first calculate the HHI within each 2-digit industry, then average across industries using industrial value added as weights. This inverted U-shape closely mirrors the long-term inverted-U pattern observed in TFP growth.

Second, we document the evolution of aggregate markups and profit shares of output over time. We follow the approach developed by De Loecker and Warzynski (2012) and De Ridder et al. (Forthcoming) to jointly estimate firm-level markups and productivity using GMM and instrumental variable techniques. We then compute the aggregate markup as the sales-weighted average across firms and industries. Figure 4 shows the evolution of the aggregate

[1] We analyse data from 1998 onward to avoid artificial entry effects due to the early collection of the data.

Figure 3: Market Concentration



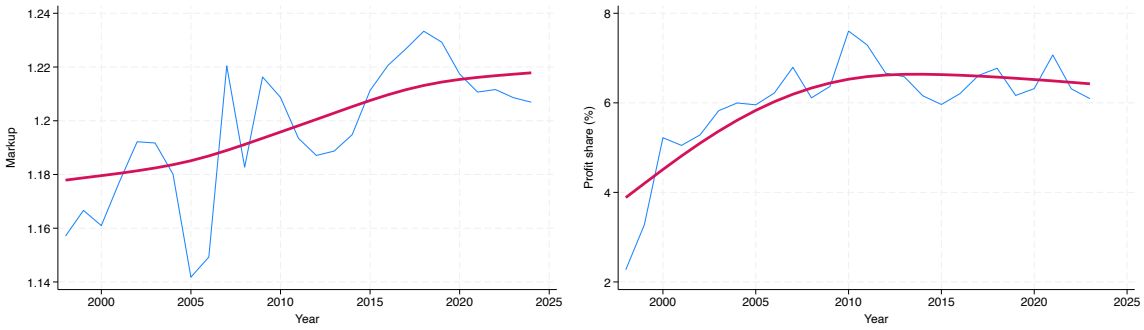
Notes. Market concentration in China, measured by the Herfindahl–Hirschman Index (HHI) using data from the China Stock Market & Accounting Research (CSMAR) Database.

markup along with its long-term trend, which is computed using the Hodrick–Prescott (HP) filter.

Figure 4: Markup and Profit Share

(a) Sales-Weighted Average Markup

(b) Profit Share of Output



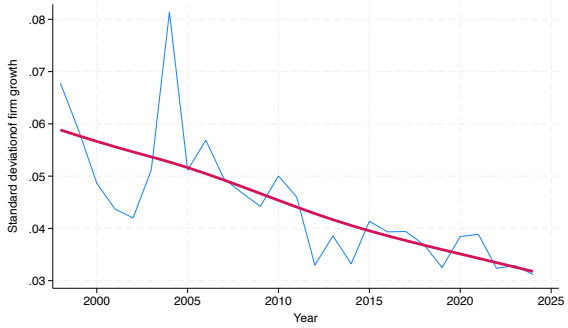
Notes. Markups and profit shares in China. The left panel shows the sales-weighted average markup, with markups estimated using the production approach of De Loecker and Warzynski (2012) and firm-level data from CSMAR. The right panel shows the aggregate profit share of output, based on data from the National Bureau of Statistics of China. The blue line denotes the raw series and the red line the HP-filtered trend.

From the figure, it is evident that the aggregate markup in China has been increasing over time. The trend started at around 17% in 1998 and rose to approximately 22% by the end of the period. This increase, however, is much milder compared to the US. For instance, De Loecker et al. (2020) documents that the aggregate markup in the US rose from about 1.21 in the 1980s to nearly 1.61 more recently (2016). The slow increase in China's aggregate markup partly reflects the decline in market concentration observed in Figure 1. As resources are diverted from productive, high-markup firms, the composition effect tends to slow down aggregate markup growth.

We now examine the evolution of the aggregate profit share in output. We document the profit share using data from the annual statistical yearbook published by the National Bureau of Statistics of China. The figure shows the evolution of profit shares over time along with the long-term trend. Evidently, the profit share has been increasing over time, although the pace of growth has been slow. The profit share of output rose from just under 5% of sales in 1998 to slightly over 6% in 2023. This movement closely mirrors the evolution of aggregate markups.

Next, we examine firm growth dispersion. Using the productivity estimates obtained from the De Loecker and Warzynski estimator, we compute firm growth over time. We then calculate growth dispersion as the standard deviation of firm growth rates within each industry, and aggregate across industries using industry value added as weights. The figure shows the evolution of firm growth dispersion along with its long-term trend. Strikingly, the dispersion of firm growth has been steadily declining over time, from about 0.08 in 1998 to slightly above 0.03 by the end of the sample period.

Figure 5: Dispersion of Firm Growth



Notes. Firm growth dispersion over time, computed by the authors using CSMAR data. The blue line denotes the raw index and the red line the HP-filtered trend.

Typically, high dispersion of firm growth indicates a dynamic economy, where younger and more productive firms grow faster than older, more mature firms, thereby contributing more to overall industrial growth. The slowdown in firm growth dispersion is a symptom of declining business dynamism, and our endogenous growth framework will need to account for this pattern.

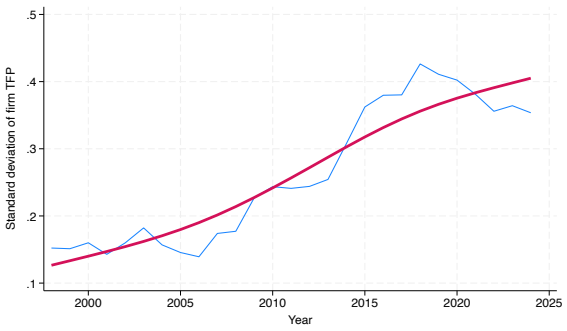
The next diagram (Figure 6a) shows productivity dispersion instead of growth dispersion. Productivity dispersion is computed as the standard deviation of firm-level productivity within each industry, and the resulting measure is then aggregated across industries using value added as weights. Both the original time series and the long-term trend are presented. Evidently, the productivity gap has been steadily increasing over time. This is somewhat expected given the decline in firm growth dispersion. As younger and newer firms no longer grow faster, incumbent firms pull further ahead in terms of productivity, causing overall productivity dispersion to rise.

We can also examine this phenomenon from another perspective. Figure 6b plots the productivity gap between frontier and laggard firms within each industry, aggregated across industries using value added as weights. Frontier firms are defined as the top 10% most productive firms in an industry, while laggard firms are the bottom 10% least productive. The figure clearly shows that the productivity gap has increased over time, rising from about 0.1 in 1998 to nearly 0.18 in 2023.

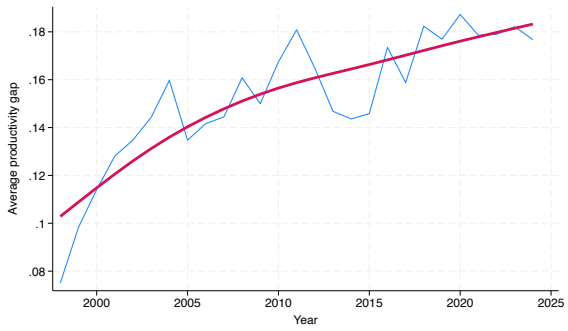
Typically, a larger productivity gap would lead one to expect increased market concentration in larger firms, which should in turn raise aggregate markups and profit shares. However, Figures 1–3 show that market concentration in China has actually declined since 2007, and aggregate markups and profit shares have not increased dramatically over time. This suggests

Figure 6: Productivity Dispersion and Frontier–Laggard Gap

(a) Productivity Dispersion



(b) Frontier–Laggard Productivity Gap



Notes. Productivity dispersion and frontier–laggard gaps in China. The left panel shows productivity dispersion, and the right panel shows the frontier–laggard productivity gap. Both are computed using CSMAR data. The blue line denotes the raw series and the red line the HP-filtered trend.

the presence of market frictions that prevent resources from flowing from less productive firms to more productive ones. Consequently, we introduce resource misallocation as a key feature in our theoretical framework.

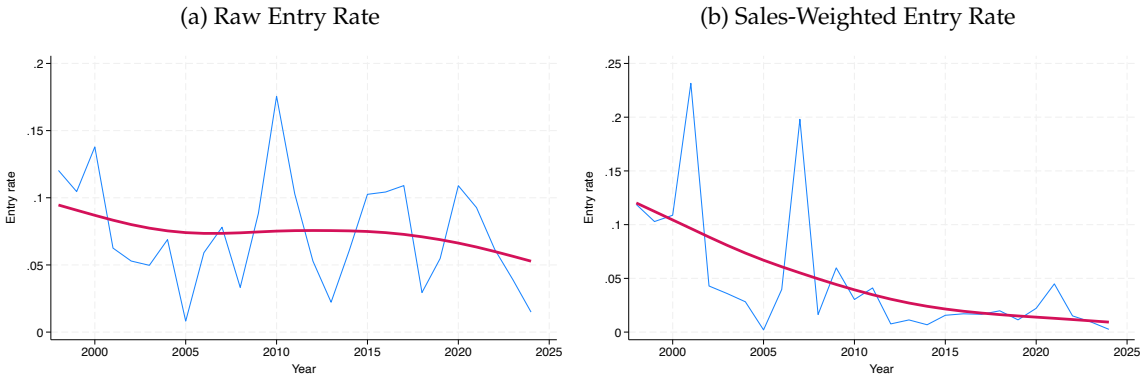
Next, we examine the "entry rates." Note that the CSMAR database covers only large, publicly listed firms. Therefore, the entry rate here refers to the number of newly listed firms as a share of the total number of publicly listed firms in the dataset. It does not represent the genuine entry rate when small- and medium-sized firms are included. Nevertheless, we use this terminology so that the empirical concept can be easily mapped to the theoretical model. These results are intended only as indicative; when calibrating the theoretical model, we will rely on related findings about entry contributions from the literature.

Figure 7 shows the raw entry rates, that is, the number of newly listed public firms over time. The entry rate trends slightly downward, but the decline is very mild. Throughout the sample period, the entry rate remains fairly stable, around 7–8%. However, when we examine the sales-weighted entry rates, a very different pattern emerges. In particular, the sales-weighted entry rate declines sharply over time, from about 12% in 1998 to around 3% in 2023.

The sales-weighted entry rate provides information about the contribution of new entrants. If newly entering firms are, on average, more productive, they capture a larger market share, and the sales-weighted entry rate will be higher. Conversely, a low sales-weighted entry rate indicates that new entrants are less productive and fail to capture a significant market share. In other words, according to the CSMAR database, the quality of new entrants appears to have declined.

Although we rely on large, publicly listed firms to document this fact, the pattern is consistent with findings in the literature. For example, Brandt et al. (2023) show that the contribution of entry to industry growth in China declined from about 33% during 1998–2007 to roughly 15% around 2013. We use these figures to calibrate our model because their dataset is more representative and includes small- and medium-sized firms.

Figure 7: Entry Rates (Raw and Sales-Weighted)



Notes. Firm entry rates in China. The panels show raw and sales-weighted entry rates (defined as the rate at which firms become publicly listed), computed using CSMAR data. The blue line denotes the raw series and the red line the HP-filtered trend.

3 Theoretical Framework

Our theoretical framework builds on Akcigit and Ates's (2023) endogenous growth model. This framework features step-by-step innovation, whereby market leaders innovate to widen the productivity gap in order to capture more market share and higher profits; follower firms also innovate to catch up with the leader, anticipating that they may eventually leapfrog the leader. Similarly, potential entrants innovate, hoping to break into the market and earn economic rents in the future. The interaction of market leaders, followers, and entrants determines the productivity gaps within each industry and the distribution of industries over productivity gaps in the steady state.

To better reflect China's growth experience, we introduce several features into this step-by-step innovation framework. These include physical capital, defensive R&D, involutionary competition, and resource misallocation à la Hsieh and Klenow (2009). We introduce physical capital because in China, final demand consists heavily of investment and less of domestic consumption, which is relevant for future welfare analysis. We endogenise knowledge diffusion by allowing firms to conduct defensive R&D to prevent knowledge from diffusing to follower firms, while taking the degree of intellectual property rights (IPR) protection as given. We also introduce involutionary competition, as it is a notable feature of the contemporary Chinese economy—a phenomenon that has recently attracted attention in both the media and policy circles. Finally, we incorporate resource misallocation, since it is an intrinsic feature of many developing countries, including China. Hsieh and Klenow (2009) shows that such misallocation can reduce a country's TFP. Here, we are particularly interested in whether such static misallocation has dynamic implications for TFP growth, rather than just affecting its level.

3.1 Preferences, Technologies, and Market Structures

We consider the following continuous-time economy. There is a representative household that derives logarithmic utility from consumption,

$$U_0 = \int_0^\infty e^{-\rho t} \ln C_t dt, \tag{1}$$

where C_t denotes consumption and ρ is the discount rate. The representative household faces the following budget constraint,

$$C_t + \dot{\mathcal{A}}_t = w_t L_t + r_t \mathcal{A}_t + T_t, \tag{2}$$

where $\mathcal{A}_t$ denotes asset holdings, L_t represents labour supply, and T_t denotes the lump-sum transfer from the government. A dot above a variable indicates its derivative with respect to time. We normalise labour supply, $L_t = 1$. The representative household takes as given the wage rate w_t, the interest rate r_t, and the price of the consumption good, which we set as the numeraire. The household's total assets consist of physical capital holdings, K_t, and the aggregate value of firms, V_t.

The representative household's problem implies the standard Euler equation,

$$\frac{\dot{C}_t}{C_t} = r_t - \rho, \tag{3}$$

whereby consumption grows at a rate equal to the interest rate less the discount rate.

There is a single final good, Y_t, which is used for consumption, R&D, and investment in physical capital:

$$Y_t = C_t + R_t + \dot{K}_t + \delta K_t, \tag{4}$$

where R_t denotes R&D expenditure and δ represents the depreciation rate of physical capital.

The final good itself is produced in perfectly competitive markets using the following Cobb-Douglas production technology:

$$\ln Y_t = \int_0^1 \ln y_{jt} dj, \tag{5}$$

where y_{jt} denotes the amount of intermediate good $j \in [0,1]$. The demand for each intermediate good is given by $y_{jt} = Y_t / p_{jt}$, where p_{jt} is its price. Each intermediate good is produced by imperfectly competitive intermediate firms, as described below.

In each product line j, there are two incumbent firms, indexed by $i \in \{1,2\}$, producing imperfectly substitutable varieties of the intermediate good y_{ijt}. The total industry output, y_{jt}, is given by

$$y_{jt} = \left(y_{1jt}^{\frac{\sigma-1}{\sigma}} + y_{2jt}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{6}$$

where $\sigma > 1$ denotes the elasticity of substitution between varieties, with higher values indicating that the varieties are closer substitutes.

The intermediate input is produced using a Cobb-Douglas production technology that

combines capital and labour:

$$y_{ijt} = q_{ijt} k_{ijt}^{\alpha} l_{ijt}^{1-\alpha}, \tag{7}$$

where k_{ijt} denotes the capital input, l_{ijt} denotes labour services, and q_{ijt} is the associated productivity. Within an industry, we use i to denote the leading firm and $-i$ to denote the follower firm, such that $q_{ijt} > q_{-ijt}$. When $q_{ijt} = q_{-ijt}$, we say that the two firms are neck and neck.

The two incumbent firms engage in Bertrand price competition. Each sets a price, p_{ijt}, to capture a larger share of the market for good j and to increase profitability from supplying the intermediate variety for that good.[2] In equilibrium, both firms set prices such that there is a markup over marginal cost. In particular, if $s_{jt} \equiv p_{ijt}/p_{-ijt}$ denotes the relative price in the industry, the markup for the leader is $m_{ijt} = \frac{\sigma + s_{jt}^{1-\sigma}}{\sigma - 1}$, and the markup for the follower firm is $m_{-ijt} = \frac{\sigma s_{jt}^{1-\sigma} + 1}{(\sigma-1)s_{jt}^{1-\sigma}}$. If there were no market distortions, the *laissez-faire* market share for the leader would be $\omega_{ijt}^{LF} = \frac{1}{1+s_{jt}^{\sigma-1}}$, and the *laissez-faire* market share for the follower firm would be $\omega_{-ijt}^{LF} = \frac{s_{jt}^{\sigma-1}}{1+s_{jt}^{\sigma-1}}$. In this oligopoly competition environment, firms' markups, along with the market shares they capture, determine their profits.

To better reflect China's growth experience, we introduce resource misallocation within industries in the spirit of Hsieh and Klenow (2009). Specifically, in the Bertrand equilibrium, leader i's market share in product line j is given by

$$\omega_{ijt} = \eta_{jt} \left(\frac{1}{1 + s_{jt}^{\sigma-1}} \right)^{1-\chi}, \tag{8}$$

where χ is a progressive wedge that captures the degree of market distortions diverting resources away from productive leaders to less productive follower firms. Moreover, η_{jt} is an endogenous industry shifter that determines the average level of distortions in market j. Since there are only two firms competing in this product line, the market share of the follower is $\omega_{-ijt} = 1 - \omega_{ijt}$.

This specification is drawn from the literature on income progressivity, such as Bénabou (2002). A higher χ indicates that more resources are being diverted from leaders to followers within a given industry. For a given χ, industries with larger productivity gaps—and hence larger differences in intermediate variety prices—tend to suffer more from inefficient resource misallocation, reflecting the "progressive" nature of the misallocation. Systematic misallocation from more productive firms to less productive ones is a key feature of developing economies and a central mechanism through which such misallocation lowers TFP (see, e.g., Hsieh and Klenow, 2009). In the absence of resource misallocation, $\chi = 0$ and $\eta_{jt} = 1$.

Given these markups and market shares, in the Bertrand equilibrium, the profits of the two

[2] Formally, the firm's optimisation problem is to choose p_{ijt} to maximise $\left(p_{ijt} - c_t/q_{ijt}\right) y_{ijt}$, subject to $p_{ijt}y_{ijt} + p_{-ijt}y_{-ijt} = Y_t$ and $y_{ijt}/y_{-ijt} = \left(p_{ijt}/p_{-ijt}\right)^{-\sigma}$. Here, $c_t \equiv (r_t/\alpha)^{\alpha} (w_t/(1-\alpha))^{1-\alpha}$ is the common cost of the capital-labour bundle.

incumbent firms $i \in \{1, 2\}$ are given by

$$\Pi_{ijt} = \left(1 - \frac{1}{m_{ijt}}\right) \omega_{ijt} Y_t. \tag{9}$$

In this environment, resource misallocation affects firms' profitability mainly through its effects on their market shares.

Firms' productivity evolves through successive innovations. When the leader innovates between t and $t + \Delta t$, with probability ε, its productivity level increases by a factor of λ, such that the new productivity level is $q_{ijt+\Delta t} = \lambda^{N_{ijt}+1}$, where N_{ijt} denotes the number of past improvements up to time t. With probability $1 - \varepsilon$, however, the leader's innovation yields no substantive improvement, so its productivity level remains at $\lambda^{N_{ijt}}$. A higher value of $1 - \varepsilon$ implies a lower marginal benefit from innovation, but the leader is still willing to devote resources in part to maintain its market position—a phenomenon we call "involutionary competition."[3]

Involutionary competition is excessive competition that results in lower marginal gains, especially from a social perspective. Here, the leader innovates partly to improve its productivity level and partly to maintain its market position, even if some of the innovative efforts produce no meaningful social gains. Such involutionary competition also applies to follower firms and new entrants, as we shall see later. We refer to the parameter $1 - \varepsilon$ as the involutionary competition parameter.

We now describe innovations by the follower firms. Following Acemoglu and Akcigit (2012), we assume that an innovation by a follower may be either incremental or drastic. With probability ϕ, the follower's innovation brings about a drastic improvement in productivity, allowing it to catch up with the leader at once—a process we call "quick catch-up." With probability $(1 - \phi)\,\varepsilon$, the follower's innovation improves its productivity level by one step, bringing it up to $\lambda^{N_{-ijt}+1}$, a process called "slow catch-up." With probability $(1 - \phi)\,(1 - \varepsilon)$, the follower's innovation leads to no essential improvement in productivity. In this case, the follower's innovative effort is partly devoted to involutionary competition.

Apart from the followers' innovative efforts, there is also a process of knowledge diffusion between firms within the industry. Specifically, with a small Poisson arrival rate κ, the follower learns the technology of the leader and closes the technology gap at once. This occurs when the follower effectively learns from the technology frontier, or when the leader's existing patents expire and no newer innovation has yet emerged. Hence, this exogenous probability of quick catch-up essentially captures the knowledge spillovers from the best firms to the rest of the industry. We label κ as the knowledge diffusion parameter.

In each product line, the productivity levels of the leader and the follower are separated by a certain number of gaps, which reflect the differences in the number of past successful innovations by each of these firms, N_{ijt} and N_{-ijt}. Formally, we denote $n_{ijt} = N_{ijt} - N_{-ijt} \geq 0$ as the technological gap between the leader i and the competitor $-i$. The relative productivity

[3]For example, in competitive product markets, many R&D cycles yield minor or purely cosmetic changes (e.g. small feature updates) that do not advance the underlying technology, while only occasionally producing genuine quality improvements. Firms nevertheless incur the cost of all such attempts, as maintaining a steady flow of releases is important for preserving their market position.

of the two incumbent firms is then given by

$$\frac{q_{ijt}}{q_{-ijt}} = \lambda^{N_{ijt}-N_{-ijt}} = \lambda^{n_{ijt}}. \tag{10}$$

The technological gap n_{ijt} will be an important state variable for firms when choosing innovation strategies, as they balance the costs and benefits of such activities.

Next, we describe innovation by entrants. Each period, a potential entrant in each product line invests in R&D to enter the market. If the entrant successfully innovates, it replaces the existing follower in the product line and becomes a new follower itself. The entrant's productivity is determined as follows. With probability ϕ, the entrant achieves a drastic innovation, replacing the follower and becoming neck and neck with the existing leader. With probability $(1-\phi)\,\varepsilon$, the entrant makes a small innovation on the follower's technology, improving it by one step. It then replaces the follower and competes with the existing leader. With probability $(1-\phi)\,(1-\varepsilon)$, the entrant's innovation leads to no essential improvement in the follower's technology. However, it still replaces the existing follower and becomes a new follower. In all cases, part of the entrant's innovation effort is devoted to generating new discoveries and improving technology, while the rest is devoted to involutionary competition. In this scenario, the entrant is able to enter the market even if its efforts produce no substantive productivity improvement from society's perspective.

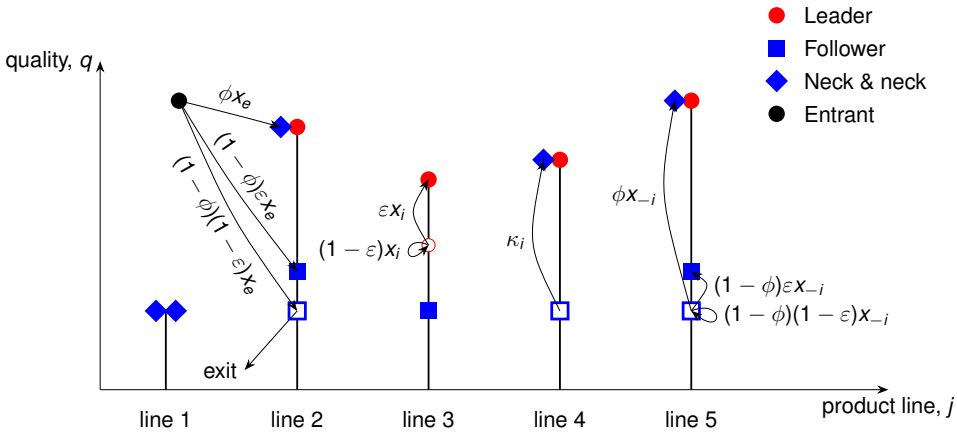
Firms invest in R&D to gain or maintain market leadership. To conduct R&D, we assume incumbent firms use final goods as input. To achieve an innovation with the Poisson arrival rate x_{ijt}, the incumbent firm incurs a convex cost of $G\left(x_{ijt}\right) = \zeta x_{ijt}^{\gamma}/\gamma$, where ζ represents the average (inverse) innovation efficiency, and γ is the (inverse) elasticity of R&D with respect to research expenditures. We further assume that the innovation technology is the same for both market leaders and their followers. Moreover, entrants are assumed to have a similar innovation technology. In particular, to generate an innovation with rate x_{ejt}, the entrant incurs a convex research cost of $G_e\left(x_{ejt}\right) = \zeta_e x_{ejt}^{\gamma}/\gamma$. Here, we allow the entrant's innovation technology to differ from that of incumbent firms through the average efficiency ζ_e, while the elasticity of R&D with respect to R&D expenditures, γ, is the same for incumbents and entrants. We separate the two innovation technologies because entrants' innovation may partly reflect entry barriers and government regulations and may therefore evolve differently from that of incumbent firms.

In addition to investing in R&D to maintain market leadership, incumbent leaders can engage in defensive R&D to limit knowledge spillover to competitors. Specifically, by incurring a convex cost $G\left(\tilde{x}_{ijt}\right) = \zeta \tilde{x}_{ijt}^{\gamma}/\gamma$, a leader i can reduce the knowledge diffusion rate from κ to $\kappa_{ijt} = H\left(\tilde{x}_{ijt}\right)$. We assume $H\left(\tilde{x}_{ijt}\right) = \kappa \exp\left(-\beta \tilde{x}_{ijt}\right)$, where β represents the (inverse) effectiveness of defensive R&D. It follows that $H\left(0\right) = \kappa$, so in the absence of defensive R&D, knowledge diffuses at the baseline rate κ. As the intensity of defensive R&D $\tilde{x}_{ijt}$ increases, the diffusion rate κ_{ijt} declines. However, defensive R&D exhibits diminishing returns. Together with the increasing convex cost of such effort, this determines the equilibrium rate of knowledge diffusion.

Figure 8 summarises how leadership positions evolve as a result of incumbent and entrant

innovations. The vertical axis shows the quality or productivity level of a given product line, while the horizontal axis depicts several representative product lines. Line 1 shows the case in which two incumbent firms are neck and neck. Line 2 begins with a situation where the leader is technologically ahead of the follower and describes how this changes with the entry of a new firm. Line 3 illustrates a case in which the leader successfully innovates and widens the technological gap in the industry. Line 4 depicts the situation where the follower catches up with the leader at once due to knowledge diffusion within the industry. Finally, line 5 illustrates the effects of a successful follower innovation on market leadership.

Figure 8: Entry, Exit, and Leadership



Notes. Evolution of product lines. The diagram depicts how innovation by entrants (line 2), incumbent leaders (line 3), and incumbent followers (line 5), together with neck-and-neck competition (line 1) and within-industry knowledge diffusion (line 4), shape entry, exit, and the evolution of technological leadership. Open squares or circles mark the previous positions of firms that innovate or exit.

3.2 Firm Values and Innovations

The step-by-step innovation model has a recursive structure, and at the core of this structure are firms' values and their innovation decisions. We now describe the stock market value of incumbent and entrant firms, as well as their innovation decisions.

To this end, we note that along the BGP, stock market values grow at the same constant rate g as final output. We therefore focus on market values normalised by GDP to make them stationary along the BGP.

We begin with the market value of incumbent leaders. The payoff-relevant state variable is the technological gap n between the leader and the follower. Hence, all leaders with the same technological gap behave in the same way. In particular, leader i's value in product line j with technological gap n is given by

$$\begin{aligned}\rho v_{nt} - \dot{v}_{nt} = \max_{x_{nt},\tilde{x}_{nt}} \Big\{ & (1-\tau_t)\,\pi_{nt} - (1-s_t)\,G\left(x_{nt}\right) + \varepsilon x_{nt}\left[v_{n+1t} - v_{nt}\right] \\ & + \phi\left(x_{-nt} + x_{ent}\right)\left[v_{0t} - v_{nt}\right] + (1-\phi)\,\varepsilon\left(x_{-nt} + x_{ent}\right)\left[v_{n-1t} - v_{nt}\right] \\ & - (1-s_t)\,G\left(\tilde{x}_{nt}\right) + H\left(\tilde{x}_{nt}\right)\left[v_{0t} - v_{nt}\right]\Big\}.\end{aligned}$$

The left-hand side represents the present value of the leading firm, accounting for any capital gains or losses.

The first component on the right-hand side is the firm's after-tax profit from producing its variety as a market leader, where $\pi_{nt} \equiv \Pi_{nt}/Y_t$ denotes normalised profit, and τ_t is the corporate income tax, assumed to be common across product lines.

The second component is the optional value of innovation as a market leader. Specifically, by investing an amount $G(x_{nt})$ in R&D, subject to an R&D subsidy s_t, the leader successfully innovates with Poisson arrival rate x_{nt}. However, not all innovation leads to productivity improvements. Only with probability ε does the leader improve its productivity level by one step, thereby increasing the market value from v_{nt} to v_{n+1t}. With probability $1-\varepsilon$, the innovation does not lead to any substantive productivity improvement, but the leader retains its leadership position. In this case, there is no change in the value function.

The next component describes what happens when competitors and entrants innovate successfully. When a follower firm or a new entrant innovates successfully (with Poisson arrival rates x_{-nt} and x_{ent}, respectively), it catches up with the leader at once with probability ϕ, reducing the market value of the leader from v_{nt} to v_{0t}. With probability $(1-\phi)\varepsilon$, the follower or entrant catches up with the leader by one technological step, decreasing the leader's value from v_{nt} to v_{n-1t}. With probability $(1-\phi)(1-\varepsilon)$, the follower or entrant does not make any essential productivity improvements, leaving the leader's market value unaffected.

The final component is the leader's optional value of investing in defensive R&D. Specifically, by investing an amount $G(\tilde{x}_{nt})$ in R&D, subject to R&D subsidies, the leader reduces the knowledge diffusion rate from κ to $\kappa_{nt} = H(\tilde{x}_{nt})$. When knowledge diffusion occurs, it decreases the market value of the leader from v_n to v_0.

Similarly, the value of an n-step follower is given by

$$\begin{aligned}\rho v_{-nt} - \dot{v}_{-nt} = \max_{x_{-nt}} \Big\{ & (1-\tau_t)\,\pi_{-nt} \\ & - (1-s_t)\,G(x_{-nt}) + \phi x_{-nt}\left[v_{0t} - v_{-nt}\right] + (1-\phi)\,\varepsilon x_{-nt}\left[v_{-n+1t} - v_{-nt}\right] \\ & + \varepsilon x_{nt}\left[v_{-n-1t} - v_{-nt}\right] + x_{ent}\left[0 - v_{-nt}\right] + \kappa_{nt}\left[v_{0t} - v_{-nt}\right] \Big\}.\end{aligned}$$

Under oligopoly competition, the follower also produces and generates positive profit, represented by the first component on the right-hand side of the value function.

The follower firm also has an optional value of innovation. By investing an amount $G(x_{-nt})$ in R&D, subject to R&D subsidies s_t, the follower successfully innovates with a Poisson arrival rate x_{-nt}. With probability ϕ, the follower catches up with the leader immediately, increasing its market value from v_{-nt} to v_{0t}. With probability $(1-\phi)\varepsilon$, the follower improves its technology by one step, raising its market value from v_{-nt} to v_{-n+1t}. With the remaining probability, the innovation does not generate any meaningful productivity improvement, leaving the follower's market value unchanged.

The next two components describe the effects of successful innovation by the leader or an entrant. When the leader innovates successfully, with probability ε, it widens the technological lead by one step, reducing the market value of the follower from v_{-nt} to v_{-n-1t}. In contrast,

when a new entrant innovates successfully, it always replaces the follower firm, reducing the old follower's market value from v_{-nt} to 0.

The last component captures the effect of knowledge diffusion. With a small Poisson arrival rate κ_{nt}, knowledge diffusion occurs within the industry, allowing the follower to learn from the technological frontier and catch up with the leader at once, increasing its market value from v_{-nt} to v_{0t}.

Next, we describe the value of a neck-and-neck incumbent firm. This case requires special treatment because there is no distinction between quick and slow technological catch-up.

$$\rho v_{0t} - \dot{v}_{0t} = \max_{x_{0t}} \left\{ (1-\tau_t)\,\pi_{0t} - (1-s_t)\,G\left(x_{0t}\right) + \varepsilon x_{0t}\left[v_{1t} - v_{0t}\right] + \varepsilon x_{-0t}\left[v_{-1t} - v_{0t}\right] + \frac{1}{2}x_{e0t}\left[0 - v_{0t}\right] \right\}. \tag{11}$$

The first component on the right-hand side is the profit of a neck-and-neck incumbent firm, subject to the corporate income tax τ_t.

The second component is the optional value of innovation. By spending an R&D expenditure of $G\left(x_{0t}\right)$, subject to an R&D subsidy s_t, the firm successfully innovates with Poisson arrival rate x_{0t}. With probability ε, the innovation raises the technology level by one step, increasing the firm's market value from v_{0t} to v_{1t}. However, with probability $1-\varepsilon$, the innovation does not lead to any substantive productivity improvement, and the market value of the firm remains unchanged.

If the competitor innovates successfully (with Poisson arrival rate x_{-0t}), with probability ε it becomes a one-step-ahead leader, and the firm becomes a one-step-behind follower, reducing its market value from v_{0t} to v_{-1t}. With probability $1-\varepsilon$, the competitor's innovation does not improve productivity, leaving the firm's market value unchanged. In equilibrium, both neck-and-neck firms innovate at the same rate, so x_{0t} is equal to x_{-0t} for the competitor.

Finally, if a potential entrant innovates successfully, it replaces one of the two incumbent firms. In this case, the market value of the incumbent firm decreases from v_{0t} to 0. Note that when an entrant innovates, it always replaces the incumbent firm regardless of whether the innovation leads to genuine productivity improvements. Hence, ε does not appear in the neck-and-neck incumbent firm's value function.

We now describe the entrants' problem. In this model, entry is targeted at particular product lines. The problem of an entrant targeting a product line with an n-step gap is as follows:

$$\max_{x_{ent}} \left\{ -G_e\left(x_{ent}\right) + \phi x_{ent} v_{0t} + (1-\phi)\,\varepsilon x_{ent} v_{-n+1t} + (1-\phi)\,(1-\varepsilon)\,x_{ent} v_{-nt} - 0 \right\}. \tag{12}$$

By investing an amount $G_e\left(x_{ent}\right)$ in R&D, the entrant achieves an innovation with a Poisson arrival rate x_{ent}. With probability ϕ, the innovation is drastic, allowing the entrant to catch up with the leader at once and obtain a market value of v_{0t}. With probability $(1-\phi)\,\varepsilon$, the entrant improves the follower's technology by one step, replaces the existing follower, and obtains a market value of v_{-n+1t}. With probability $(1-\phi)\,(1-\varepsilon)$, the entrant's innovation yields no essential productivity gains over the follower. However, it still replaces the follower and obtains a market value of v_{-nt}. This represents a situation in which the entrant engages in

involutionary competition to break into the market, without generating any substantive social gains. The last term captures the entrant's outside option, normalised to zero.

The problem of an entrant targeting a product line in a neck-and-neck state is defined similarly.

$$\max_{x_{e0t}} \left\{ -G_e\left(x_{e0t}\right) + \varepsilon x_{e0t} v_{1t} + \left(1-\varepsilon\right) x_{e0t} v_{0t} - 0 \right\}. \tag{13}$$

By devoting resources to R&D, the entrant successfully innovates with Poisson arrival rate x_{e01}. It replaces one of the incumbents with probability ε, obtaining a market value of v_{1t}. With probability $1-\varepsilon$, the innovation does not lead to substantive productivity improvements, but the entrant still replaces one of the incumbents and competes with the other in a neck-and-neck state. In this case, it obtains a market value of v_{0t}. The final term reflects the entrant's outside option (set to zero).

We have now completed the description of the value functions for different types of firms. Solving these value functions yields the innovation rates for these firms, which in turn determine the aggregate growth rate of the economy.

3.3 Distribution of Industries

Before describing the aggregate economy, we first describe the evolution of the distribution of industries. To this end, we denote μ_{nt} as the measure of product lines in which the technological gap between the leader and the follower is n. Naturally, the measures of different product lines sum to one, $\sum_{n=0}^{\infty} \mu_{nt} = 1$.

The law of motion of μ_{nt} for $n > 1$ is as follows:

$$\dot{\mu}_{nt} = -\left(\varepsilon x_{nt} + \left(\phi + \left(1-\phi\right)\varepsilon\right)\left(x_{-nt} + x_{ent}\right) + \kappa_{nt}\right)\mu_{nt} + \varepsilon x_{n-1t}\mu_{n-1t} + \left(1-\phi\right)\varepsilon\left(x_{-n-1t} + x_{en+1t}\right)\mu_{n+1t}. \tag{14}$$

The first term on the right-hand side represents the departures from state n when there is a substantive innovation by incumbents or entrants in those industries, or when knowledge diffusion occurs, allowing follower firms to catch up quickly with market leaders. The second term represents additions to state n when the incumbent leader in industries with state $n-1$ successfully innovates and widens the technological gap by one step. The third term captures inflows into state n from industries with state $n+1$ when a successful and substantive innovation by follower firms or entrants reduces the technological gap by one step.

Similarly, the transition of μ_{1t} follows the differential equation:

$$\dot{\mu}_{1t} = -\left(\varepsilon x_{1t} + \left(\phi + \left(1-\phi\right)\varepsilon\right)\left(x_{-1t} + x_{e1t}\right) + \kappa_{1t}\right)\mu_{1t} + \left(2\varepsilon x_{0t} + \varepsilon x_{e0t}\right)\mu_{0t} + \left(1-\phi\right)\varepsilon\left(x_{-2t} + x_{e2t}\right)\mu_{2t}. \tag{15}$$

The interpretation of this equation is similar to the one above, except for the second term. Here, additions to state 1 occur when either one of the two incumbent firms or a new entrant successfully innovates, improving productivity by one step, and the industry transitions from state 0 to state 1.

Finally, the transition of μ_{0t} is governed by the following equation:

$$\dot{\mu}_{0t} = -\left(2\varepsilon x_{0t} + \varepsilon x_{e0t}\right)\mu_{0t} + \sum_{n=1}^{\infty} \kappa_{nt}\mu_{nt} + (1-\phi)\,\varepsilon\left(x_{-1t} + x_{e1t}\right)\mu_{1t} + \sum_{n=1}^{\infty} \phi\left(x_{-nt} + x_{ent}\right)\mu_{nt}. \tag{16}$$

Briefly, the first term represents departures from state 0, which occur when incumbent firms or entrants successfully make substantive innovations, widening the technological gap by one step. The second term represents additions to state 0 when knowledge diffusion occurs in any unlevelled industries, bringing those industries to the neck-and-neck state. The third term captures inflows into state 0 when there is a successful and substantive innovation by follower firms or entrants in industries with a single technological gap. The last term represents transitions into the neck-and-neck state when incumbents or entrants in unlevelled industries successfully make a drastic innovation and immediately catch up with their leaders.

3.4 Aggregate Economy and Balance Growth Path

The endogenous growth framework is highly tractable and admits closed-form solutions for various aggregate variables. We are now ready to describe these variables. First, the productivity frontier of the economy is given by

$$Q_t \equiv \exp\left(\int_0^1 \ln q_{jt} dj\right), \tag{17}$$

where $q_{jt} \equiv \left(\sum_{i=1}^{2} q_{ijt}^{\sigma-1}\right)^{\frac{1}{\sigma-1}}$ denotes the productivity of industry j. In equilibrium, labour market clearing requires that $1 = L_t = \int_0^1 \sum_{i=1}^{2} l_{ijt} dj$, and capital market clearing requires that $K_t = \int_0^1 \sum_{i=1}^{2} k_{ijt} dj$.

Using these expressions, the final output of the economy can be expressed as

$$Y_t = A_t K_t^{\alpha} L_t^{1-\alpha}, \tag{18}$$

where $A_t \equiv \mathcal{D}_t Q_t$ represents the TFP of the economy. The TFP consists of two components: the productivity frontier Q_t and a distortion index $\mathcal{D}_t$, which arises because different firms charge different markups.

The distortion index can be expressed as

$$\mathcal{D}_t \equiv \exp\left(-\int_0^1 \ln\left[m_{ijt}^{1-\sigma}\frac{\lambda^{n_{ijt}(\sigma-1)}}{\lambda^{n_{ijt}(\sigma-1)}+1} + m_{-ijt}^{1-\sigma}\frac{1}{\lambda^{n_{ijt}(\sigma-1)}+1}\right]^{\frac{1}{1-\sigma}} dj\right) M_t, \tag{19}$$

where $M_t \equiv \left(\int_0^1 \sum_{i=1}^{2} \frac{\omega_{ijt}}{m_{ijt}} dj\right)^{-1}$ denotes aggregate markups, and the remaining term summarises how variations in markups interfere with resource allocation, thereby reducing TFP (see also Peters, 2020). [4]

[4] The terms in brackets are expanded to include both firms in the production line j, the leader i and the follower $-i$, and thus written without a summation over $i = 1, 2$.

Furthermore, in equilibrium, aggregate R&D expenditure is given by

$$R_t = \int_0^1 \left[G\left(x_{ijt}\right) + G\left(\tilde{x}_{ijt}\right) + G\left(x_{-ijt}\right) + G_e\left(x_{ejt}\right) \right] dj. \tag{20}$$

These R&D investments support the innovations of incumbent leaders, followers, and new entrants. These innovation efforts, in turn, determine the evolution of the productivity frontier of the economy. In particular, the productivity frontier evolves according to

$$g_t = \left(\mu_{0t}\left(2x_{0t} + x_{e0t}\right) + \sum_{n=1}^{\infty} \mu_{nt} x_{nt} \right) \varepsilon \ln \lambda, \tag{21}$$

which determines the aggregate TFP growth rate along the BGP.

Along the BGP, per capita consumption and output grow at the same constant rate g. Per capita consumption is given by $C_t = k^{\alpha}\left(1 - \theta_R - \theta_K\right) Y_t$, where $\theta_R \equiv R_t / Y_t$ denotes the share of R&D expenditure in output,

$$\theta_K = \alpha M^{-1} \frac{\delta + g/\left(1-\alpha\right)}{\rho + \delta + g/\left(1-\alpha\right)} \tag{22}$$

is the share of physical capital investment in output and $k = \left(\frac{\alpha M^{-1}}{\rho + \delta + g/(1-\alpha)} \right)^{\frac{1}{1-\alpha}}$ is the per capita capital, which remains constant along the BGP.

Using these expressions, aggregate welfare in the BGP can be expressed as:

$$\mathcal{W}_0 = \int_0^{\infty} e^{-\rho t} \ln C_t dt = \ln\left[A_0 k^{\alpha}\left(1 - \theta_R - \theta_K\right)\right] \frac{1}{\rho} + \frac{g}{\rho^2\left(1-\alpha\right)}, \tag{23}$$

where $A_0 \equiv 1$ is the initial level of TFP of the economy.

4 Calibration

We now take this model to the data to explore which potential drivers may lie behind the slowdown. Specifically, we first calibrate the model to match the Chinese economy in the initial BGP, which we set to the period between 2003 and 2007. We then keep most of the model's structural parameters fixed while varying a subset corresponding to potential channels through which TFP growth may have been affected. By comparing the initial and terminal BGPs, we can identify the proximate causes of the slowdown, thereby laying the groundwork for further investigation into this puzzling yet important phenomenon.

4.1 Model Counterparts of Empirical Moments

Before starting the calibration process, we introduce the model's counterparts for the empirical moments of interest. Our framework is highly tractable, allowing us to derive closed-form solutions for the aggregate and distributional moments used in calibration. This tractability also allows us to use methods such as GMM instead of relying on computationally intensive simulations.

First, an important aggregate moment is the TFP growth rate, which has already been given by equation (21). The TFP growth rate helps pin down the innovation step size, among other parameters. Next, the aggregate R&D share and aggregate investment share in output are given by equations (20) and (22), respectively. The aggregate profit share of GDP is defined by

$$\text{profit share}_t = 1 - \frac{1}{M_t}, \tag{24}$$

where the aggregate markup M_t is given below equation (19). These aggregate shares help pin down the capital's output elasticity in the production function, the elasticity of substitution between varieties within each industry, and the innovation costs for both incumbents and new entrants.

The next important aggregate moment is market concentration within industries. Typically, market concentration within an industry increases with the productivity gap between frontier and laggard firms. However, in economies with distortions, market concentration also reflects the degree of resource misallocation. Therefore, this measure helps pin down the parameter governing resource misallocation in our framework. The average sales concentration, measured by the Herfindahl–Hirschman index, is given by

$$\text{hhi}_t = \sum_{n=0}^{\infty} \mu_{nt} \left[\omega_{nt}^2 + (1 - \omega_{nt})^2 \right]. \tag{25}$$

The productivity gap between frontier and laggard firms is defined as the difference in average productivity between leaders and followers in the model. The economy-wide productivity gap is then calculated as the average of industry-level productivity gaps and is given by

$$\text{productivity gap}_t = \sum_{n=0}^{\infty} \mu_{nt} \ln \lambda^n. \tag{26}$$

This aggregate moment serves as a measure of productivity dispersion.

We next introduce a second aggregate moment that captures the same underlying dispersion using firm-level productivity. Since the theoretical framework includes a continuum of industries, firm productivity dispersion can be computed using the law of total expectations. Specifically, the dispersion of firm productivity is the sum of the within-industry variation in firm productivity and the dispersion of average firm productivity across industries. Mathematically, it can be expressed as

$$\text{Var}_t \left[\ln q_t \right] = \sum_{n=0}^{\infty} \mu_{nt} \text{Var}_t \left(\ln q_{nt} \right) + \sum_{n=0}^{\infty} \mu_{nt} \left(\text{E}_t \left[\ln q_{nt} \right] - \ln \bar{q}_t \right)^2 = \sum_{n=0}^{\infty} \mu_{nt} \left\{ n^2 + (n - \bar{n}_t)^2 \right\} \frac{\ln \lambda^2}{4}, \tag{27}$$

where average firm productivity is $\ln \bar{q}_t \equiv \sum_{n=0}^{\infty} \mu_{nt} \text{E}_t \left[\ln q_{nt} \right] = \bar{n}_t \ln \lambda$ and the average gap is $\bar{n}_t \equiv \sum_{n=0}^{\infty} \mu_{nt} n$.

Using a similar approach, we can calculate another important aggregate moment: the dispersion of firm growth. Again applying the law of total expectations, economy-wide firm growth dispersion is the sum of growth dispersion within industries and firm types (leaders, followers, or entrants) and the variation in average firm growth across industries and types.

The theoretical expression for firm growth dispersion is given by

$$\mathrm{Var}_t\left[g_t\right] = \frac{1}{3}\sum_{n=0}^{\infty}\sum_{k\in\{l,f,e\}}\mu_{nt}\,\mathrm{Var}_t\left(g_{nt}^k\right) + \frac{1}{3}\sum_{n=0}^{\infty}\sum_{k\in\{l,f,e\}}\mu_{nt}\left(\mathrm{E}_t\left[g_{nt}^k\right]-\bar{g}_t\right)^2, \tag{28}$$

where k indexes the type of firm, namely, a leader (l), a follower (f), or an entrant (e), and $\bar{g}_t$ denotes average firm growth across types and industries: $\bar{g}_t = \frac{1}{3}\sum_{n=0}^{\infty}\sum_{k\in\{l,f,e\}}\mu_{nt}\,\mathrm{E}_t\left[g_{nt}^k\right]$.

Given the Poisson arrival rates of innovation, the expected logarithmic growth of a leader in an industry with gap n is $\mathrm{E}_t\left[g_{nt}^l\right] = x_{nt}\varepsilon\ln\lambda$, and the variance of logarithmic growth for a leader is $\mathrm{Var}_t\left[g_{nt}^l\right] = x_{nt}\varepsilon\ln\lambda^2$. The corresponding expressions for followers and entrants can be derived similarly.[5]

Next, the average markup in the economy is defined as

$$\text{markup}_t = \sum_{n=0}^{\infty}\mu_{nt}m_{nt}. \tag{29}$$

Although the average markup is easy to obtain both theoretically and empirically, recent studies have highlighted limitations of empirically estimating markups based solely on revenue data without information on prices (see, for example, Bond et al., 2021; De Ridder et al., Forthcoming). In particular, empirically estimated markups convey useful information about trends and variation, but their levels may be biased. Given this, we use empirical markups to assess changes in firm profitability over time, but we refrain from using their levels to directly calibrate the model.

Next, the aggregate entry rate is determined by the distribution of industries and the innovation intensity of entrants targeting those industries. Specifically, it is given by

$$\text{entry rate}_t = \frac{1}{2}\sum_{n=0}^{\infty}\mu_{nt}x_{ent}, \tag{30}$$

where the division by two reflects the fact that two firms operate in each product line at all times. A related aggregate moment is the sales share accrued by new entrants, defined as

$$\begin{aligned}\text{entrant sales share}_t = \frac{1}{2}\sum_{n=1}^{\infty}\mu_{nt}x_{ent}\left(\phi\omega_{0t} + (1-\phi)\,\varepsilon\omega_{-n+1t} + (1-\phi)(1-\varepsilon)\,\omega_{-nt}\right)\\ + \frac{1}{2}\mu_{0t}x_{e0t}\left(\varepsilon\omega_{1t} + (1-\varepsilon)\,\omega_{0t}\right).\end{aligned}$$

Although we can compute the entry rate and entrants' sales shares directly from the CSMAR data, we refrain from using these empirical moments directly because the dataset only includes large, publicly listed firms and may not reflect genuine entry by small- and medium-sized firms. Instead, we rely on the entry contributions reported in the literature, such as Brandt et

[5] For a Poisson process with rate x and a random jump size, the mean and variance of the jumps are both equal to x, adjusted for the first and second moments of jump size, respectively. Using this property, the expected log growth of a follower in an industry with productivity gap n is $\mathrm{E}_t\left[g_{nt}^f\right] = \left\{x_{-nt}\left((1-\phi)\,\varepsilon+\phi\varepsilon n\right)+\kappa_{nt}n\right\}\ln\lambda$, and the variance of log growth for a follower is $\mathrm{Var}_t\left[g_{nt}^f\right] = \left\{x_{-nt}\left((1-\phi)\,\varepsilon+\phi\varepsilon n^2\right)+\kappa_{nt}n^2\right\}\ln\lambda^2$. Similarly, the expected log growth of an entrant in an industry with gap n is $\mathrm{E}_t\left[g_{nt}^e\right] = x_{ent}\left((1-\phi)\,\varepsilon+\phi\varepsilon n\right)\ln\lambda$, and the variance is $\mathrm{Var}_t\left[g_{nt}^e\right] = x_{ent}\left((1-\phi)\,\varepsilon+\phi\varepsilon n^2\right)\ln\lambda^2$.

al. (2023).

In what follows, we formalise this concept as an aggregate moment capturing the contribution of entry to TFP growth. In the model, TFP growth is driven by innovations from incumbent leaders, followers, and new entrants. Together, these innovations push the economy's technological frontier forward, leading to sustained economic growth. The entry contribution is defined as the share of TFP growth coming from entrant innovations and is given by

$$\text{entry contribution}_t = \frac{1}{2g_t} \sum_{n=0}^{\infty} \mu_{nt} \, \mathrm{E}_t \left[g_{nt}^e \right], \tag{31}$$

where the division by two reflects that two firms push the productivity frontier outward in any given product line. This measure is particularly useful for identifying differences between the innovation costs of entrants relative to those of incumbents.

Finally, an important set of moments is the distribution of industries across technology gaps. The theoretical counterpart is the measure $\{\mu_{nt}\}_{n=0}^{\infty}$, which can be used to construct the cumulative density function (CDF). We target several key points of the CDF, including the fraction of industries with productivity gaps less than or equal to 0.05, 0.1, 0.2, and 0.5. This completes the discussion of the model's counterparts to empirical moments.

4.2 Data and Calibration

We first calibrate the initial BGP of the model to reflect the state of the Chinese economy before the 2008 global financial crisis. There are a total of 15 structural parameters to be calibrated, $\{\rho, \gamma, \delta, \tau, s, \lambda, \alpha, \sigma, \zeta, \zeta_e, \phi, \kappa, \chi, \beta, \varepsilon\}$. We calibrate five of them externally, using values sourced from the literature, and calibrate the remaining parameters internally using GMM to match empirical moments in the data.

We begin with the externally calibrated parameters. First, we set the discount rate, ρ, to 0.05. Next, we set the curvature parameter of the R&D production function for both incumbents and entrants to $\gamma = 1/0.35$, consistent with values used in previous literature (e.g., Kortum, 1993; Acemoglu and Akcigit, 2012; Akcigit and Ates, 2023). We set capital's depreciation rate to 5%, a value commonly used in macroeconomic studies. We set the corporate income tax rate to 33%, reflecting the statutory corporate tax rate in China before the 2008 financial crisis. Finally, we set the R&D subsidy to 0, since widely available R&D subsidies were virtually non-existent until shortly before the 2008 crisis.

We then calibrate the remaining parameters to match a set of aggregate and distributional moments that capture the key features of the model. Although the GMM estimator jointly estimates all parameters simultaneously, we have selected empirical moments that provide specific information on particular parameters. For example, the TFP growth rate informs the step size of innovation in the model. The aggregate investment share helps discipline the capital elasticity in the production function. The aggregate profit share provides information on the elasticity of substitution between varieties, which affects the market power of incumbent firms and, consequently, their profitability.

Moreover, the aggregate R&D share, together with the entry contribution, helps pin down the average innovation costs for both incumbents and entrants, with the entry contribution

particularly useful for identifying the relative costs of entrant versus incumbent innovations. Market concentration, measured by the Herfindahl–Hirschman index, provides insight into the degree of resource misallocation in the economy, helping to calibrate parameter χ. Finally, the dispersion of firm growth, together with the distribution of industries across productivity gaps, is useful for calibrating the probability of drastic innovation (ϕ), the average rate of knowledge diffusion within industries (κ), the effectiveness of defensive R&D in limiting knowledge diffusion (β), and the degree of involutionary competition in the economy ($1-\varepsilon$).

4.3 Parameter Values and Model Fit

Table 1 summarises the calibrated parameters for the initial BGP, representing the state of the Chinese economy between 2003 and 2007, while Table 2 reports the model's fit to the data. Overall, the model successfully matches the empirical moments. Apart from slightly overstating the aggregate profit share and aggregate R&D share, all other aggregate and distributional moments are matched closely.

Table 1: List of Parameter Values

Parameter	Description	Value
Panel A. Externally calibrated		
ρ	Discount rate	0.05
γ	R&D cost curvature	1/0.35
δ	Depreciation rate	0.05
τ	Corporate taxes	0.33
s	R&D subsidies	0
Panel B. Internally calibrated		
λ	Innovation step size	1.0719
α	Capital's output elasticity	0.5414
σ	Elasticity of substitution	87.0009
ζ	R&D cost scale, incumbents	1.6614
ζ_e	R&D cost scale, entrants	0.4130
ϕ	Probability of drastic innovation	0.0175
κ	Inverse IPR protection	0.1049
χ	Misallocation	0.3150
β	Defensive R&D effectiveness	0.5290
$1-\varepsilon$	Involution	0.0240

Notes: Parameters in Panel A are calibrated externally using values from the literature (see Section 4.2). Parameters in Panel B are calibrated internally by matching model moments to the empirical moments in Table 2.

Turning to the internally calibrated parameters, the innovation step size is estimated to be 1.0719, the capital elasticity in the production function is estimated to be 0.5414, the elasticity of substitution between varieties within industries is estimated to be 87.0009, and the effectiveness of defensive R&D in reducing knowledge diffusion is estimated to be 0.5290. The probability of making drastic innovations is very small, approximately 1.75 percent, meaning that, for the majority of the time, laggard firms had to catch up with leaders slowly and gradually. These five parameters will be held fixed when we calibrate the terminal BGP to investigate the proximate causes of China's TFP growth slowdown.

A few remarks on the remaining parameters are in order. First, the entrants' innovation cost is estimated to be about one-fourth of the incumbents' innovation cost, indicating that entry was relatively easy in China before 2007.

The degree of knowledge diffusion in the economy is estimated to be 0.1049. One way to interpret this parameter is to think of knowledge diffusion as the inverse of intellectual property rights protection. In this case, the estimate suggests that, on average, a patent would expire roughly ten years after being granted in China before 2007. The degree of misallocation is estimated to be 0.3150, suggesting that even before 2007 there was substantial resource misallocation in the Chinese economy, echoing earlier findings in the literature (e.g., Hsieh and Klenow, 2009).

Finally, the intensity of involutionary competition, or hyper-competition without social gains, is estimated to be about 2.40 percent. This implies that, before 2007, involution was not a major problem. This likely reflects the fact that China was still behind the world technology frontier at that time and that ample innovation opportunities allowed firms to direct their innovation efforts more effectively and productively.

Table 2: Data and Model Moments

Target	Data	Model
Panel A. Aggregate moments		
TFP growth rate	0.0331	0.0331
Aggregate investment share	0.3382	0.3382
Aggregate profit share	0.0601	0.1197
Aggregate R&D share	0.0124	0.0381
Entry contribution	0.3300	0.3300
Market concentration (HHI)	0.3817	0.3817
Dispersion of firm growth	0.0554	0.0541
Panel B. Distribution of industries		
Fraction of industries with productivity gap		
≤ 0.05	0.1667	0.1793
≤ 0.1	0.4444	0.4444
≤ 0.2	0.7778	0.6475
≤ 0.5	0.8889	0.9081

Notes: The model is calibrated to match the Chinese economy in 2003–2007. TFP growth rate is obtained from the Penn World Table version 9.1. Aggregate investment share in GDP and Aggregate R&D expenditure share in GDP are obtained from the World Bank. Aggregate profit share in GDP is taken from De Loecker and Eeckhout (2021). Entry contribution is based on Brandt et al. (2023). Market concentration, the dispersion of firm growth, and the distribution of industries are calculated using data from CSMAR.

In summary, the calibration results provide valuable insights into the state of the Chinese economy between 2003 and 2007, lending confidence in using this framework to investigate the causes of the subsequent slowdown in productivity growth.

4.4 Investigation of Productivity Slowdown

We now move on to calibrate the model to reflect the state of the Chinese economy during the period 2013–2017. By comparing the calibration results between the initial and terminal BGPs,

we can identify the proximate causes of China's TFP growth slowdown. To this end, we keep a subset of the internally calibrated parameters fixed and recalibrate the remaining parameters to match the 2013–2017 data targets.

We select the parameters for recalibration based on several potential channels that might have affected TFP growth. First, changes in industrial policies, such as the corporate income tax rate and R&D subsidies, may influence firms' innovation incentives and, consequently, aggregate TFP growth. Second, changes in innovation costs for incumbents and entrants may directly affect firms' innovation decisions, which could help explain the productivity slowdown. This channel is particularly relevant as China moves closer to the world technology frontier. Recent studies suggest that ideas are becoming harder to find in advanced economies like the US.

Third, we consider the possibility that knowledge diffusion within industries has slowed in China. This could be a contributing factor, especially in light of findings that the decline in business dynamism in the US since the 1980s was mainly driven by weaker knowledge diffusion over time (see, e.g., Akcigit and Ates, 2023).

Finally, we consider two channels particularly relevant to developing economies like China. The first is an increased degree of resource misallocation, which diverts resources from more productive firms to less productive ones. In the context of an endogenous growth model, greater misallocation not only lowers the level of aggregate productivity but also affects TFP growth by altering firms' innovation incentives. The second is the intensification of involutionary competition in China over time. Such competition can be wasteful of innovation resources, generating no social gain, and could potentially explain the puzzling combination of rising R&D expenditures and slowing TFP growth.

To capture these channels, we focus on a set of internally calibrated parameters $\{\zeta, \zeta_e, \kappa, \chi, \varepsilon\}$, as well as two externally calibrated policy parameters, τ and s. We set the corporate income tax rate to 25 percent. The Chinese government significantly revised its corporate income tax system, effective January 1, 2008, unifying it and reducing the standard rate to 25 percent. We thus set the corporate income tax to this new statutory rate.

We set the R&D subsidy rate to 23 percent. Although China's formal R&D tax incentives date back to the 1990s, these mainly took the form of preferential tax policies and direct funding that were not widely accessible to firms. Starting in 2006, however, the government relaxed the eligibility criteria for a significant R&D super deduction, allowing firms to deduct 150 percent of qualified R&D expenses, making these incentives more broadly available. By the end of the period, the effective R&D subsidy rate is estimated at approximately 23 percent, which we use to calibrate the terminal BGP of our model. We then calibrate the remaining parameters to match selected aggregate and distributional moments for the Chinese economy during 2013–2017.

Table 3 shows the new calibration results for the terminal BGP, and Table 4 reports the fit of the model to the data. Overall, the model fit is very good. It generates a significant slowdown in productivity growth, consistent with the observed data. It also closely matches changes in the entry contribution, market concentration, and firm growth dispersion. The model-generated distribution of industries aligns fairly closely with the distribution observed

in the data.

Table 3: Calibrated Parameter Values: Terminal BGP

Parameter	Description	Value
Panel A. Externally calibrated		
ρ	Discount rate	0.05
γ	R&D cost curvature	1/.35
δ	Depreciation rate	0.05
τ	Corporate taxes	0.2500 ($\downarrow$)
s	R&D subsidies	0.2300 ($\uparrow$)
Panel B. Internally calibrated		
λ	Innovation step size	1.0719
α	Capital's output elasticity	0.5414
σ	Elasticity of substitution	87.0009
ζ	R&D cost scale, incumbents	4.7337 ($\uparrow$)
ζ_e	R&D cost scale, entrants	2.8255 ($\uparrow$)
ϕ	Probability of drastic innovation	0.0175
κ	Inverse IPR protection	0.0259 ($\downarrow$)
χ	Misallocation	0.5496 ($\uparrow$)
β	Defensive R&D effectiveness	0.5290
$1-\varepsilon$	Involution	0.6684 ($\uparrow$)

Notes: Parentheses after values in the third column indicate that the parameter is recalibrated; otherwise, it is held fixed at the baseline value reported in Table 1. Up (down) arrows indicate that the estimated value increases (decreases) relative to the baseline in Table 1. Parameters in Panel A are calibrated externally, whereas those in Panel B are calibrated internally using the generalised method of moments (GMM).

Turning to the estimates of the selected parameters, several patterns emerge.

First, innovation costs for both incumbents and entrants have increased substantially toward the end of the sample period. The incumbent innovation costs have nearly tripled (4.7337 in 2017 vs. 1.6614 in 2007), and the innovation costs for entrants have also risen several times (2.8255 vs. 0.4130). This indicates that as China moved closer to the world technology frontier, innovation has indeed become more difficult and costly, consistent with findings in the literature that ideas are generally getting harder to find. More importantly, the relative innovation costs of entrants compared to incumbents have increased from about one-fourth (0.4130/1.6614) in the earlier period to nearly one-half (2.8255/4.7337) in the later period. This suggests that not only have innovation costs increased overall, but innovation has also become less accessible to new entrants, potentially explaining the decline in the entry contribution to the economy over time.

Second, the knowledge diffusion parameter has declined significantly (0.0259 in 2017 vs. 0.1049 in 2007), despite differences in development stages between China and the US. This indicates that a slowdown in knowledge diffusion is also occurring in China, suggesting that knowledge diffusion slowdowns are not specific to the US but may be a more global phenomenon driven by technological and structural changes. A decline in knowledge diffusion could contribute to the TFP slowdown, as knowledge spillovers within industries become less effective.

Third, the degree of misallocation is estimated to have increased substantially toward the

end of the sample (0.5496 in 2017 vs. 0.3150 in 2007). This implies that after the global financial crisis, the functioning of the market deteriorated over time. Greater misallocation could negatively impact TFP growth by altering firms' innovation incentives.

Table 4: Model Fit: Terminal BGP

Target	Data	Model
Panel A. Aggregate moments		
TFP growth rate	0.0083	0.0083
Entry contribution	0.1500	0.1500
Market concentration (HHI)	0.2589	0.2589
Dispersion of firm growth	0.0375	0.0342
Panel B. Distribution of industries		
Fraction of industries with productivity gap		
≤ 0.05	0.0980	0.1707
≤ 0.1	0.3922	0.3922
≤ 0.2	0.6863	0.5444
≤ 0.5	0.9412	0.8034

Notes: The model is recalibrated to match the Chinese economy in 2013–2017. TFP growth rate is obtained from the Penn World Table version 9.1. Entry contribution is based on Brandt et al. (2023). Market concentration, the dispersion of firm growth, and the distribution of industries are calculated using data from CSMAR.

Finally, we estimate that involutionary competition has intensified considerably toward the end of the sample period (0.6684 in 2017 vs. 0.0240 in 2007). The probability of firms engaging in such competition was very low before the financial crisis but increased substantially afterward. One potential reason is that, following the global financial crisis, the decline in foreign demand and weaker domestic demand forced firms to compete fiercely over smaller markets. This competition extends beyond product markets and may influence other aspects of corporate strategy, including product innovation. Firms may allocate substantial resources to innovation activities primarily to consolidate market positions or expand market share, without generating meaningful improvements in product quality or operational efficiency. Such dynamics could further contribute to the slowdown in TFP growth. Overall, our analysis suggests that the channels we have identified are highly relevant for explaining the TFP growth slowdown in China.

Before we decompose the relative importance of each channel in the TFP growth slowdown experienced in China, we note that our model replicates key empirical dynamics that were not directly targeted in the calibration. Table 5 shows how well the model tracks these untargeted empirical patterns. Specifically, it presents a selection of aggregate moments not targeted in the calibration, showing both their changes over time in the data and the model's predictions. The model correctly predicts the direction of change for all these moments except the entry rate. For example, the profit share increased by 0.35 percentage points in the data, and the model predicts an increase of 2.62 percentage points. Similarly, the R&D share in GDP increased by 0.80 percentage points in the data, while the model predicts 1.93 percentage points. The model slightly overstates the increase in aggregate markups but captures the directional change correctly.

Table 5: Level Changes (Δ) in Untargeted Moments

Moment	Data (%)	Model (%)
Δ profit share	+0.35	2.62
Δ R&D share	+0.80	1.93
Δ markup	+2.56	6.77
Δ entry rate	+2.08	−11.04
Δ entrant sales share	−4.53	−11.63
Δ productivity gap	+1.23	6.23
Δ productivity dispersion	+16.17	6.14

Notes: The model is calibrated to match the Chinese economy in 2003–2007. TFP growth rate is obtained from the Penn World Table version 9.1. Aggregate investment share in GDP and Aggregate R&D expenditure share in GDP are obtained from the World Bank. Aggregate investment share is the gross fixed capital formation by the private sector as a share of GDP. Aggregate profit share in GDP is obtained from De Loecker and Eeckhout (2021). Entry rate and the distribution of industries are calculated using the matched patent data.

The empirical increase in the entry rate is somewhat surprising, given the long-run downward trend documented earlier (see Figure 7). Nevertheless, the five-year average entry rate shows a slight increase between 2003–2007 and 2013–2017. Our model captures the long-term decline in the entry rate, as well as the decline in entry sales shares, reasonably well. While the model slightly overstates changes in the productivity gap and understates changes in productivity dispersion, it predicts both in the correct direction. Overall, we conclude that the model provides a reasonable fit to the data.

4.5 Contributions of Individual Channels

We now turn to decomposing the relative importance of each channel in explaining the productivity growth slowdown between 2007 and 2017. To this end, we perform counterfactual experiments in which we shut down one channel at a time. Shutting down a channel means keeping the parameter governing that channel fixed at its initial BGP level, while allowing all other parameters to adjust to their terminal BGP levels. We then examine how much TFP growth slowdown the model generates without that channel. By comparing the hypothetical slowdown to the actual slowdown when all channels are active, we can assess the relative magnitude of each channel's contribution to the productivity slowdown.

Formally, let the aggregate TFP growth rate at time 2007 be denoted by g_{2007}, its value at time 2017 when all channels are active by g_{2017}, and the aggregate productivity growth when a specific channel i is absent by g_{2017}^{-i}. The contribution of channel i to the total TFP growth slowdown can then be expressed as follows:

$$\text{contribution}_i = \frac{g_{2017} - g_{2017}^{-i}}{g_{2017} - g_{2007}}. \tag{32}$$

Table 6 presents the decomposition results, revealing several interesting findings.

First, unlike in the US, slower knowledge diffusion (κ) is not a major driver of China's TFP growth slowdown. The contribution of this channel is positive but relatively small, a point we will return to later.

Second, a key factor behind the slowdown is that ideas are becoming harder to find. Both innovation and entry costs (ζ and ζ_e) have risen substantially toward the end of the sample period. This likely reflects the fact that innovation has become more costly as China approaches the world technology frontier. The country can no longer rely solely on imitating technology from advanced economies and adapting it locally; instead, it increasingly needs to push the frontier itself, which is more difficult and expensive. Together, the increase in innovation and entry costs can account for about 21 percent of the total productivity growth slowdown.

Table 6: Decomposition results

	Individual Channels						
	ζ	ζ_e	κ	τ	s	χ	$1-\varepsilon$
Contribution (%)	17.78	2.77	1.03	−1.59	−3.31	1.39	56.37

Notes: The table reports the percentage contribution of specific channels to the total model-generated deviation of TFP growth between 2007 and 2017 (as defined in Eq. 32). Negative values indicate that the corresponding channels move TFP growth in the opposite direction to the empirical trend. The specific channels include: (i) increases in innovation costs ζ; (ii) increases in entry costs ζ_e; (iii) decreases in knowledge diffusion κ; (iv) decreases in corporate income taxes τ; (v) increases in R&D subsidies s; (vi) increases in resource misallocation χ; and (vii) increases in involutionary competition $1-\varepsilon$.

Third, and perhaps most importantly, involutionary competition ($1-\varepsilon$) has intensified. This means that firms are devoting more research effort, but achieving fewer productivity gains. This channel alone can explain more than 56 percent of the slowdown. This finding also sheds light on the puzzling pattern highlighted earlier in the paper: despite substantial increases in both R&D expenditures and patent applications, TFP growth has declined since the global financial crisis. The intensified involutionary competition provides a key explanation for this apparent decoupling between R&D effort and productivity growth.

Fourth, rising resource misallocation (χ) also slows TFP growth, contributing about 1.39 percent to the slowdown. At first glance, this channel may seem minor. Upon closer examination, however, interesting lessons emerge. While the magnitude of the misallocation channel appears small, it is roughly comparable to the positive effect of the corporate income tax cut. In other words, if misallocation had not worsened after the crisis, the government could have achieved the same growth rate without resorting to the tax cut.

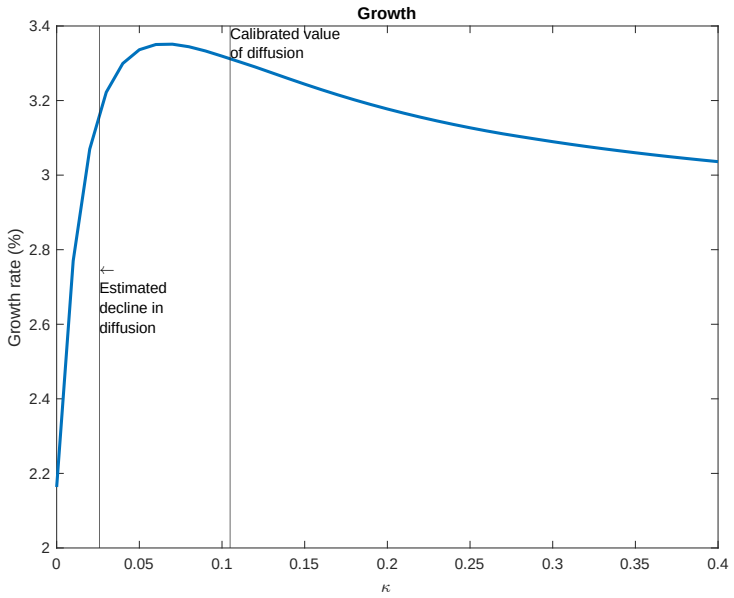
From a theoretical perspective, our framework demonstrates that static resource misallocation, as documented in the literature, can have dynamic implications for TFP growth. When market frictions divert resources away from highly productive firms, it reduces their incentives to innovate. Because laggard firms and potential entrants are forward-looking, such misallocation also diminishes their innovation incentives.

Finally, it should be noted that industrial policy responses, including the corporate tax rate cut (τ) and the increase in R&D subsidies (s), were generally supportive of TFP growth. Nevertheless, policy alone is insufficient to prevent the slowdown. This implies that other policy instruments—particularly those addressing resource misallocation and involutionary competition—may be necessary to support sustained productivity growth in the future.

Before concluding, an important question remains: Why does slower knowledge diffusion

hurt the US but not China? To understand this, it is helpful to note that knowledge diffusion is a double-edged sword. On one hand, it promotes knowledge spillovers across firms, increasing productivity growth. On the other hand, it can reduce technological leaders' incentives to innovate, as they may not be able to fully reap the rewards of their discoveries once these spill over to the wider economy. For these reasons, the relationship between aggregate growth and the degree of knowledge diffusion follows an inverted-U shape, as illustrated in Figure 9. The figure also shows the calibrated knowledge diffusion rate for China before the 2008 financial crisis, along with the estimated decline towards 2017.

Figure 9: Growth Effects of Knowledge Diffusion



Notes. Growth rate as a function of knowledge diffusion based on a calibrated version of the model. Details of the parameter calibration are described in Section 4.2 and summarised in Table 1. Vertical lines show (i) the calibrated value of knowledge diffusion (0.1049) and (ii) the estimated decline in knowledge diffusion between 2007 and 2017 (from 0.1049 to 0.0259).

From the diagram, it is clear that around 2007, China's knowledge diffusion rate was relatively high, likely due to weaker intellectual property rights protection at the time. A subsequent strengthening of IP protection increased firms' innovation incentives and supported growth, at least initially. Even though knowledge diffusion declined towards the end of the sample period, the two effects appear to offset each other, so the overall impact of slower diffusion on China's TFP growth slowdown remains modest.

In contrast, knowledge diffusion in the US was already quite slow before the 1980s due to strong IPR protection. The US economy was thus on the left-hand side of the inverted-U curve, where further declines in diffusion reduce knowledge spillovers and slow growth and business dynamism. In this case, the negative effect on spillovers dominates the incentive effect.

An important implication is that one cannot simply assume a phenomenon observed in one economy will have the same effect in another. Differences in development stages and institutional circumstances can lead to very different outcomes.

5 Conclusion

China achieved extraordinary productivity growth following the start of the economic reforms in 1978. However, since the 2008 global financial crisis, growth has become more muted and puzzling, with TFP growth continuing to decline despite exceptionally strong R&D activity. At the same time, firm-level dynamics point to weakening competitive outcomes, including declining market concentration, limited growth in markups and profits, and a reduced contribution from new entrants. These patterns suggest a disconnect between innovation activity and aggregate productivity performance.

To account for these developments, we build a step-by-step endogenous innovation model by extending the framework of Akcigit and Ates (2023) along several dimensions, including involutionary competition, resource misallocation and strategic defensive R&D. Our quantitative analysis identifies two primary drivers of China's productivity slowdown. First, intensified involutionary competition diverts a substantial amount of innovative resources towards activities that generate limited social returns, thereby reducing the effectiveness of aggregate R&D in driving productivity growth. Second, innovation and entry have become increasingly costly as China approaches the global technology frontier, slowing the pace of productivity growth. Together, these forces account for the bulk of observed slowdown in TFP growth. Resource misallocation and slower knowledge diffusion also contribute, albeit to a much lesser extent. We also find that while industrial policies, such as corporate income tax cuts and R&D subsidies, have supported growth, they were insufficient to offset these broader structural forces.

Our findings highlight involutionary competition as a central challenge for China's future growth. By directing resources away from productivity-enhancing innovation towards zero-sum competition, involution reduces the effectiveness of aggregate R&D investment. From an economic perspective, involution can be understood as arising from excess capacity and the resulting intensified competition over limited demand. When firms compete for limited markets, resources are increasingly allocated to defensive or rent-seeking activities, with little or no aggregate gains. Addressing this challenge is inherently difficult, as it arises from deeper structural imbalances rather than a single market distortion. This suggests that policies aimed at rebalancing structural supply and demand, as well as those that improve the allocation of innovative effort, are likely to be important.

A second key implication is that innovation is becoming costlier for both incumbent firms and potential entrants. While policy factors may play a role, the increase in innovation costs are likely to be driven by technological factors and China's stage of developments. This is consistent with the observation that "ideas are getting harder to find" as economies approach the technological frontier (Bloom et al., 2020). As a result, policies that focus solely on increasing the scale of R&D investment may face diminishing returns unless they are complemented by measures that improve the efficiency of innovation.

More broadly, our analysis point to the importance of allocative efficiency for sustaining productivity growth. Beyond direct measures such as R&D subsidies, policies that allow resources to flow towards more productive and innovative firms remain critical for supporting growth. In this context, policies that reduce involutionary competition can complement traditional innovation policies by redirecting effort towards more socially productive activities and

strengthening the aggregate impact of R&D.

Finally, our analysis opens several avenues for further research. First, the role of involutionary competition in shaping economic outcomes remains insufficiently understood. Further work is needed to develop theoretical frameworks, empirical measures, and provide causal evidence on how such competition affects innovation and productivity. Second, the rise in innovation costs raises important questions about the sources of this trend, the extent to which it reflects global versus country-specific factors, and the policies that can mitigate these costs. Addressing these questions would not only shed further light on China's growth prospect but also offer insights for economies facing similar challenges as they approach the technological frontier.

References

Acemoglu, Daron and Ufuk Akcigit, "Intellectual Property Rights Policy, Competition and Innovation," *Journal of the European Economic Association*, 2012, *10* (1), 1–42.

_ , _ , Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, 2018, *108* (11), 3450–91.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li, "Good Rents versus Bad Rents: R&D Misallocation and Growth," 2025.

Akcigit, Ufuk and Sina T. Ates, "What Happened to US Business Dynamism?," *Journal of Political Economy*, 2023, *131* (8), 2059–2124.

Barwick, Panle Jia, Luming Chen, Shanjun Li, and Xiaobo Zhang, "Entry Deregulation, Market Turnover, and Efficiency: China's Business Registration Reform," *Review of Economics and Statistics*, Forthcoming.

Bénabou, Roland, "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?," *Econometrica*, 2002, *70*, 481–517.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, 2020, *110* (4), 1104–44.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch, "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data," *Journal of Monetary Economics*, 2021, *121*, 1–14.

Brandt, Loren, Johannes Van Biesebroeck, and Yifan Zhang, "Creative accounting or creative destruction? Firm-level productivity growth in Chinese manufacturing," *Journal of Development Economics*, 2012, *97* (2), 3390351.

_ , Johannes Von Biesebroeck, Yifan Zhang, and Luhang Wang, "Where Has all the Dynamism Gone? Productivity Growth in China's Manufacturing Sector, 1998-2013," 2023. BOFIT Discussion Papers 6.

Cerdeiro, Diego A. and Cian Ruane, "China's Declining Business Dynamism," *Canadian Journal of Economics*, 2024, *57* (3), 726–752.

De Loecker, Jan and Frederic Warzynski, "Markups and Firm-Level Export Status," *American Economic Review*, 2012, *102* (6), 2437–71.

_ and Jan Eeckhout, "Global Market Power," 2021. Working Paper.

_ , _ , and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

De Ridder, Maarten, Basile Grassi, and Giovanni Morzenti, "The Hitchhiker's Guide to Markup Estimation: Assessing Estimates from Financial Data," *Econometrica*, Forthcoming.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–82.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

König, Michael, Kjetil Storesletten, Zheng Song, and Fabrizio Zilibotti, "From Imitation to Innovation: Where Is All That Chinese R&D Going?," *Econometrica*, 2022, *90* (4), 1615–1654.

Kortum, Samuel, "Equilibrium R&D and the Patent–R&D Ratio: U.S. Evidence," *American Economic Review*, 1993, *82* (2), 450–57.

Liu, Ernest, Atif Mian, and Amir Sufi, "Low Interest Rates, Market Power, and Productivity Growth," *Econometrica*, 2022, *90* (1), 193–221.

Peters, Michael, "Heterogeneous Markups, Growth and Endogenous Misallocation," *Econometrica*, 2020, *88* (5), 2037–2073.