Trimming the Tree to Help it Grow: Selective Exit and Productivity Growth in China's State Sector *

Adam A. Oppenheimer† Yutong Zhong‡

June 28, 2025

Abstract

Chinese state-owned enterprises achieved productivity convergence with private firms between 1998 and 2007, with 80% of this improvement driven by selective firm exits. To estimate the effects of exits on productivity, we develop a new decomposition method taking into account that exiting firms do not only have low levels of productivities, but low predicted productivity growth rates. Accounting for the differences in growth rates is key: the contribution of exit is only 12% when we use traditional methods that only take into account level differences. The decomposition analysis reveals two distinct phases in state sector productivity growth during this transformative decade. From 1998 to 2002, "grasp the large, let go of the small" reforms led to systematic elimination of underperforming firms. From 2003 to 2007, institutional reforms enabled restructuring and realigned incentives, driving genuine operational improvements among state firms. These findings suggest China's industrial transformation followed a critical sequence - market selection eliminated underperformers ("trimming the tree") before firms achieved productivity improvements ("growing the tree").

*We are grateful to Anmol Bhandari, Hannes Malmberg, Todd Schoellman, and Kjetil Storesletten for their invaluable guidance. We also thank Jason Hall, Kyle Herkenhoff, Robert Parham, and Lieyuan Yang for insightful discussion.

†University of Minnesota, Department of Economics. Email: oppen040@umn.edu.

‡University of Minnesota, Department of Economics. Email: zhong462@umn.edu.

1 Introduction

This paper resolves a puzzle in China's economic transformation: between 1998 and 2007, state-owned enterprises (SOEs) achieved remarkable productivity convergence with private firms (Hsieh and Song (2015), Chen et al. (2021), and Brandt et al. (2022)), yet their share of revenue declined from around 50% to 30% (Hsieh and Song (2015)).[1] This paradox of improving productivity metrics alongside declining market presence, while retaining preferential access to inputs,[2] challenges conventional understanding of how competition drives industrial evolution. We develop a new decomposition method that reveals selection effects, specifically the systematic exit of underperforming firms, account for approximately 80% of the observed productivity improvements during this transformative decade.

China's SOE reforms in the late 1990s, known as "grasp the large, let go of the small,"[3] represent one of history's largest industrial restructurings. The key component of the policy was to keep only a few large SOEs in strategic sectors and to merge, privatize, and close most other medium-to-small SOE firms. Yet these reforms did not occur in isolation. They were part of a comprehensive shift in government philosophy - from using SOEs as instruments of social policy to expecting them to compete on market terms. This paradigm shift manifested through complementary policies: the Company Law in 1993 created legal separation between state ownership and operational control of SOEs, the Commercial Bank Law in 1995 provided a framework to impose hardening budget constraints on SOEs (Chow (2015)), WTO accession in 2001 exposed domestic firms to international competition (Khandelwal et al. (2013) and Brandt et al. (2017)),[4] private sector liberalization removed entry barriers

[1] There are some exceptions in certain industries like energy, tobacco, and telecommunications where the state maintained a dominant market presence.

[2] It is well documented in the literature that state-owned firms had preferential access to inputs such as labor, capital, land and export licenses. See Bai et al. (2006), Hsieh and Klenow (2009), Song et al. (2011b), Brandt and Zhu (2001), and Khandelwal et al. (2013) for details.

[3] The policy was formally announced in the Fourth Plenum of the Communist Party's Central Committee in 1999. This policy represented a fundamental shift from a system in which SOEs were not permitted to fail to one in which closure was permitted and in certain cases even encouraged.

[4] In fact, the WTO indirectly facilitated domestic SOE reforms by overcoming certain domestic opposition. Acemoglu et al. (2006) show that vested interests can lead to development traps because they make it difficult to shut down inefficient but politically connected firms. The WTO tied the hands of reform-minded policymakers and helped them overcome domestic opposition. Premier Zhu Rongji, who spearheaded China's WTO entry, explicitly appealed to this logic in the quote "Rapid Waters Should Wash Away Dirty Sands."

(Brandt et al. (Forthcoming)), and reductions in migration costs increased worker mobility (Tombe and Zhu (2019)).[5] Together, these created an environment where market forces, rather than political directives, increasingly determined firm survival. Understanding productivity growth during this period thus requires overcoming the challenge of disentangling the effects of selection from genuine operational improvements.

To address this challenge, we develop a novel method that can decompose changes in entire productivity distributions into three components: (i) within-all effects capturing efficiency changes across all firms; (ii) exit effects measuring selection through firm closure;[6] and (iii) entry effects from new firms. While our approach is similar in spirit to standard methods like Baily et al. (1992) and Foster et al. (2001) (BHC-FHK),[7] such decompositions implicitly assume exiting firms would have evolved at a rate proportional to the rate of growth among survivors had they not exited. As a consequence, these methods tend to underestimate the role of selection. Our approach improves upon these methods by considering the additional deterioration in productivity that exiting firms would have experienced had they continued to operate.[8]

Our method is motivated by a statistical framework for Markovian productivity dynam-

Frustrated by bureaucratic opposition to SOE reforms, he viewed WTO accession as an important step toward reducing the role of SOEs. This argument suggests that China's entry into the WTO was as much about internal reforms as it was about standard external reforms, such as lower tariff barriers. See the discussion from Khandelwal in Hsieh and Song (2015) for more details.

[5]Before the economic reform started in 1978, working outside one's housing registration (hukou) location or category was prohibited. This prohibition was relaxed in the 1980s but prior to 2003, workers without local hukou still had to apply for a temporary residence permit. This was difficult, so many did not and risked arrest and deportation by the local authorities. For details see Chan (2010).

[6]Our definitions of "within-all effects" and "exit effects" are in the same spirit as Backus (2020)'s "treatment effects" and "selection effects." However, while Backus (2020) focuses on identifying the causal effect of competition on firm productivity accounting for selection bias, our paper addresses the different question of how to decompose observed TFP growth into meaningful components based on the underlying stochastic processes governing firm dynamics. In addition, his approach assumes that selection follows a threshold rule, while we allow selective exit to manifest as differential exit rates by productivity, which is a closer match to the relationship between productivity and exit observed in the data.

[7]Other popular accounting-style decompositions of aggregate productivity growth include Griliches and Regev (1995), Olley and Pakes (1996) and Melitz and Polanec (2015). What distinguishes these decompositions is primarily their choice of reference points for measuring the contributions of entering and exiting firms.

[8]In this regard, our approach is similar in spirit to model-based decompositions like Gollop et al. (1987), Basu and Fernald (2002), Petrin and Levinsohn (2012), and more recently, Baqaee and Farhi (2020) in the sense that it is amenable to counterfactual analysis and is able to provide economically interpretable components. However, rather than relying on equilibrium concepts, our approach is rooted in the statistical properties of productivity distribution dynamics.

ics. Since many decompositions are consistent with this framework (including our proposed method, as well as popular methods like BHC-FHK), in order to distinguish between productivity decompositions, we propose three properties that we argue any economically meaningful decomposition should satisfy. We then show that our decomposition satisfies these properties while standard decompositions do not.

We provide a short description of the framework. Suppose that firm productivity evolves as a Markov process $P(z'|z)$, and each period, firms first draw a new productivity z', then exit with probability $Q(z')$.[9] After continuing firms draw whether to survive or exit, a mass λ of new firms enter with productivity drawn from distribution $G(z)$.

The three properties we propose for an economically meaningful decomposition build on the primitives of the underlying process, P, Q, and λG. First, the "pure effects" property asserts that a change in one component of the underlying process should be allocated entirely to one effect.[10] Second, the "independence of irrelevant alternatives" property asserts that within-all effect measurements should remain unaffected by shifts in exit probabilities. Third, the "no residuals" property asserts that the within-all, exit, and entry effects should fully explain the change in the distribution of productivities, with no residual terms. The first property prevents changes in one component from showing up in the decomposition as changes in another (what we term "cross-contamination"), obscuring the true sources of productivity growth; the second property maintains exogeneity of within-all growth to exit rates; and the third property precludes residual terms that can be economically meaningful in magnitude but difficult to interpret.

For estimation, we first show how our proposed decomposition is directly computable using the primitives introduced in the statistical model above, then show how to estimate

[9] We follow the timing in Hopenhayn (1992) where firms see next period's shock before exiting. With this timing, exit thresholds are based on future productivity, while in Jovanovic (1982) and Olley and Pakes (1996), current productivity determines exit. However, unlike Hopenhayn (1992), we do not have a general equilibrium model to determine the productivity cutoff level.

[10] In particular, the pure within-all effect means that if productivity transitions change while exit probabilities and the mass and productivity distribution of entrants remain fixed, this change should be allocated entirely to the within-all effect; the pure exit effect means that if exit probabilities change while productivity transitions and the mass and productivity distribution of entrants remain fixed, this change should be allocated entirely to the exit effect; and the pure entry effect means that if the mass and productivity distribution of entrants change while productivity transitions and exit probabilities remain fixed, this change should be allocated entirely to the entry effect.

the primitives in practice. We prove identification of the primitives given observables under a weak set of assumptions, then discuss estimation via a direct solution to the system of equations and via maximum likelihood. Estimation in this context presents computational challenges - one contribution of the paper is the introduction of a novel regularization approach that penalizes the condition number of the estimated transition matrix to ensure stable and economically meaningful estimates.

Using microdata from China's Annual Survey of Industrial Firms, we replicate and extend two seemingly puzzling empirical patterns observed during the reform period from 1998 to 2007 that have been previously noted in the literature. These patterns hold in virtually all industrial sectors and appear contradictory. First, we establish a convergence in productivity between SOEs and private firms and further the literature by showing the convergence holds for the entire distribution, not just the mean, and also quantify the rate of convergence. Second, we show that this period saw a dramatic decline in SOE market presence.

The combination of improving productivity but declining market presence is especially puzzling given that SOEs retained preferential access to bank credit and government contracts, as is well documented in the literature (see Bai et al. (2006), Riedel et al. (2020), Song et al. (2011b), Song et al. (2011a), and Liu et al. (2021)).

To rationalize these two patterns, we propose a resolution: selective firm exit. Under this hypothesis, the apparent productivity improvements observed among surviving SOEs would reflect selection effects rather than genuine competitive gains. We then document systematic patterns in firm exit and survival in the data that suggest selection effects may have been pervasive during the transition period we study, and potentially played a key role in driving the observed convergence patterns.

To explain the puzzling pattern in China's state sector, we apply our decomposition method. Our baseline results show that the exit effect drove around 84.7%, the within-all effect drove around 14.7%, and the entry effect drove the remaining 0.6% of the productivity convergence between state-owned and private firms during the transition. By looking at the contribution of each effect year-by-year, we also uncover that China's industrial transformation occurred in two distinct phases. During the initial reform period (1998-2002), selection effects through firm exit explain effectively all of the productivity growth in the state sector.

At the same time, within-all effects are close to zero, or even negative. Together, these effects mean that without exit, firms would have on average become worse over time; but the exit of approximately two-thirds of all SOEs during this period was so strongly related to performance that this period actually observed positive growth.

While 1998-2002 witnessed dramatic restructuring, with annual exit rates around 45-58% among the least productive SOEs compared to around 25-30% for the most productive,[11] the transformation entered a new phase in 2003 following the establishment of the State-owned Assets Supervision and Administration Commission (SASAC) as the legal owner of the state-owned groups.[12,13] SASAC professionalized state asset management and enabled surviving SOEs to pursue efficiency-enhancing consolidation alongside strategic expansion. After 2005, the split-share structure reform fundamentally realigned incentives by making state shareholders' returns dependent on stock market performance, catalyzing a dramatic acceleration in operational improvements.[14] Consistent with the timing of these reforms, the results of our decomposition show that the within-all effect turned positive in 2003, though exit effects continued to be the dominant driver of productivity growth. By 2005 and continuing through to 2007 (the last year of our data), within-all effects had become the primary driver of productivity growth in the state sector, indicating that state firms had genuinely transformed their operations.

Throughout this reform decade, new entrants contributed to aggregate TFP growth in the state sector by replacing less productive exiting firms, but crucially, these entrants were no more productive than incumbents. This suggests that both new and surviving SOEs operated with similar underlying technologies and capabilities.

The decomposition results remain robust across alternative firm-level productivity mea-

[11] Berkowitz et al. (2017) also document that SOEs that subsequently exited were the least productive SOEs from 1998 to 2002, indicating selection on SOE liquidation by productivity.

[12] Following the sixteenth National Congress's policy guidelines that emphasized both strengthening state-owned enterprises and promoting private sector growth, the Chinese government established SASAC in March 2003 (Yang (2015)). SASAC itself is under the authority of the State Council. The national-level commission directly supervises central SOEs, which are the largest in China's key industries, while indirectly overseeing local SOEs, which are under the direct supervision of local governments (Song (2018)).

[13] The only exceptions are state-owned tobacco companies (controlled by the State Tobacco Monopoly Administration) and state-owned financial institutions (owned by the Ministry of Finance and Central Huijin Investment Ltd.).

[14] The split-share structure reform liberalized the state-owned shares of listed SOEs with legitimate trading rights on the secondary market. See Liao et al. (2014) for details on the reform.

sures, different estimation specifications, and when estimating using an AR(1)-style functional form assumption.

To highlight the importance of accounting for selection effects, we compare our results with those recovered by the widely-used BHC-FHK decomposition. While both methods show similar temporal patterns - declining exit effects and increasing within-firm effects over time - they disagree on the primary drivers of aggregate productivity growth. BHC-FHK attributes productivity growth mainly to within-firm improvement (33.7%) and the entry effect (33.9%), with exit playing a minor role (12.0%). Our selection-corrected decomposition reveals the opposite: the exit effect dominates (84.7%), while the within-all effect contributes modestly (14.7%) and entry is nearly negligible (0.6%). This discrepancy stems from two methodological differences: first, BHC-FHK measures exiting firms' productivity in only their final observed year, missing the full decline trajectory that our approach captures through a selection correction; and second, the methods differ conceptually in measuring "within-firm improvement" - BHC-FHK includes only surviving firms while our approach includes all firms' productivity trajectories, including declining exiters. These differences fundamentally alter our understanding of Chinese SOE reform, shifting the narrative from firm-level improvements and new entry to market selection and the removal of unproductive firms.

Our findings both complement and challenge the influential work of Hsieh and Song (2015), who document similar productivity convergence patterns but reach different conclusions about underlying mechanisms. While they attribute most productivity gains in the state sector to improvements among surviving SOEs and the entry of large SOEs while finding minimal effects from exit, our decomposition reveals that exit drove approximately 80% of productivity growth throughout the 1998-2007 period. This deviation stems from how we model exiting firms' productivity deterioration preceding exit. In order to properly evaluate the effectiveness of the policy reforms, the correct counterfactual is one in which SOEs lacked the option to exit. This requires carefully matching the pre-exit productivity trajectories of exiting firms in order to accurately predict how they would have performed had they been forced to continue operating. However, the static approach in Hsieh and Song (2015) cannot match these trajectories, while our approach is specifically designed to match

this dynamic. On the other hand, the two approaches align in finding that surviving SOEs eventually became major contributors to growth in the state sector, but our decomposition method precisely identifies when this transition occurred, finding that it coincides with the establishment of SASAC and the introduction of the split-share structure reforms.

Our paper has several important implications. First, the dominance of exit effects suggests that China's early reform gains came primarily through eliminating the least productive firms rather than improving existing operations. This finding supports theories emphasizing the importance of market selection in driving aggregate productivity growth. Second, our results provide new perspective on the sequencing of economic reforms. The finding that selection effects preceded within-firm improvements suggests that creating competitive pressure may be a necessary precondition for organizational change within state firms. Third, our findings speak to broader questions about state capitalism as an economic model. The ability of reformed SOEs to generate productivity improvements after the selection process indicates that state ownership need not be incompatible with efficiency gains, provided appropriate institutional reforms create competitive pressure first.

2 Conceptual Framework: Decomposing Changes in the Productivity Distribution

In this section we propose our decomposition of firm growth. We begin by formalizing a Markovian framework for the evolution of the distribution of firm-level productivities. The framework includes three primitives: one governing firm-level productivity growth, one governing selective firm exit, and one governing firm entry. We then propose a decomposition to attribute the effects of changes in primitives to changes in the distribution of productivities. In order to differentiate our decomposition from others, we propose a set of axioms that we argue should be satisfied by any economically meaningful decomposition of productivity growth, and show that our decomposition satisfies these properties while standard decompositions do not. In Appendix E we demonstrate through a simple firm dynamic model how our approach correctly distinguishes between productivity growth from selection effects

versus broad-based within-firm efficiency improvements, while standard decompositions do not. Proofs appear in Appendix C.

2.1 Environment

We characterize the evolution of the distribution of firm-level productivities through three primitives.

The Markov kernel $P(z'|z)$. $P(z'|z)$ captures how a firm's productivity z evolves from one period to the next. This function encapsulates all factors that influence productivity transitions, including management quality, technology adoption, and learning.

The exit probability $Q(z)$. $Q(z)$ governs which firms choose to leave the market. Unlike the deterministic threshold rule in Hopenhayn (1992), where firms exit with certainty if their productivities fall below some fixed threshold $\underline{z}$, our framework allows for probabilistic exit decisions. This generalization better matches empirical patterns where firms with similar productivity levels may make different exit decisions due to unobserved factors like location-specific shocks or implicit government safety nets. Moreover, it accommodates cases where firms may temporarily maintain operations despite low productivity due to adjustment costs or expectations of future recovery.

The mass and distribution of potential entrants $\lambda G(z)$. $G(z)$, the potential entrant distribution, represents the initial productivity draw available to potential entrants after they make their entry decision. This distribution reflects the technological frontier and capabilities accessible to new firms, independent of any selection based on expected profitability or entry costs. λ then represents the mass of entrants relative to incumbents. λ serves two purposes in our framework. First, it ensures the existence of a non-degenerate stationary distribution by replacing exiting firms with new entrants. Second, it captures the overall scale of entry in the economy. To simplify exposition and our empirical analysis, rather than consider $G(z)$ and λ separately, we consider them jointly as $\lambda G(z)$.

These three primitives interact to drive the evolution of the productivity distribution, which we formally define as follows:

Definition 1. *(Distribution of productivities)*

The distribution of productivities $\pi(z)$ satisfies the following:

a. $\pi(z) \geq 0 \ \forall z \in \mathcal{Z}$ *and* $\int_{z \in \mathcal{Z}} \pi(z) = 1$,

b. $\pi(z)$ *updates according to*

$$\pi'(z|\pi, P, Q, \lambda G) = \alpha_{\pi,P,Q,\lambda G} \left(\lambda G(z) + \int_{z' \in \mathcal{Z}} \pi(z') P(z|z')(1 - Q(z)) dz' \right), \quad (1)$$

where $\alpha_{\pi,P,Q,\lambda G}$ *normalizes the probabilities such that (a) holds, and*

c. the stationary distribution solves the fixed point of (b).

The distribution update rule (Equation 1) provides a framework for analyzing productivity dynamics. This formulation allows us to measure distinct contributions to changes in the productivity distribution from within-firm efficiency changes (i.e., changes in P), selection through exit of firms (i.e., changes in Q), and entry of new firms (i.e., changes in λG). Moreover, it enables evaluation of counterfactual policies. When the economy transitions from regime A to regime B, the primitives $\{P, Q, \lambda G\}$ change accordingly, leading to a distributional shift from π_A to π_B. Using superscripts $i \in \{A, B\}$ to denote variables under alternative policies, our goal is to decompose the distribution change $\pi_B - \pi_A$ into economically interpretable components that isolate the contribution of each channel. This decomposition provides valuable insights for understanding the mechanisms through which policy affects aggregate productivity.

2.1.1 Timing and Economic Interpretation

The productivity evolution process described in Equation 1 follows a specific timing that is crucial for correctly understanding firm dynamics. In our framework, the sequence unfolds as follows. In each period, firms produce using their current productivity level z. At the beginning of the next period, firms draw new productivity z' according to the transition kernel $P(z'|z)$. Based on their new productivity z', firms decide whether to continue or exit according to $Q(z')$.[15] At the end of the period, new firms with mass λ enter with productivity

[15] We abstract from dynamic considerations and focus purely on current period profit maximization for exit decisions. However, the model could be easily extended to include dynamics if we have $Q(\mathbb{E}(z''|z'))$. In

drawn from the distribution $G(z')$.

We believe that this timing reflects the natural sequence of business decisions: firms operate with their current capabilities, experience productivity changes (through innovation, management changes, or external shocks), and then make decisions about whether to continue or exit based on their updated productivity level.

2.2 Properties of an Economically Meaningful Decomposition

Before presenting our decomposition, it is important to establish desirable properties that an economically meaningful decomposition should satisfy.

Property 1. *(Pure Within-all Effect) A distributional change from a regime change that affects the transition kernel P but not $\{Q, \lambda G\}$ should be attributed entirely to the within-all effect.*

Property 2. *(Pure Exit Effect) A distributional change from a regime change that affects the exit probability Q but not $\{P, \lambda G\}$ should be attributed entirely to the exit effect.*

Property 3. *(Pure Entry Effect) A distributional change from a regime change that affects the mass and distribution of potential entrants λG but not $\{P, Q\}$ should be attributed entirely to the entry effect.*

Property 4. *(Completeness - No Residual Term) The decomposition should account for the entire change in the productivity distribution without any error or residual terms, such that $\Delta\pi = W + X + E$, where W, X, and E represent the within-all effect, exit effect, and entry effect, respectively.*

Property 5. *(Independence of Irrelevant Alternatives (IIA)) The contribution attributed to changes in one primitive should not be affected by changes in other primitives. Sup-*

this setup, firms make their exit decision based on expected future productivity given their current state. This would make it harder to disentangle the exit effect from changes in the transition kernel, because the exit decision would be based on the entire transition process $P(z'|z)$ rather than just the new productivity level.

pose we have three regimes A, B, and C where regime $A \coloneqq \{P_A, Q_A, \lambda_A G_A\}$*, regime* $B \coloneqq \{P_B, Q_B, \lambda_B G_B\}$*, regime* $C \coloneqq \{P_C, Q_C, \lambda_C G_C\}$*. Then one of the following must hold:*

a. *The within-all effect should be independent of the exit effect: if* $P_B = P_C \neq P_A$*, then the within-all effects in decomposing* $\pi_A - \pi_B$ *or* $\pi_A - \pi_C$ *should be identical.*

b. *The exit effect should be independent of the within-all effect: if* $Q_B = Q_C \neq Q_A$*, then the exit effects in decomposing* $\pi_A - \pi_B$ *or* $\pi_A - \pi_C$ *should be identical.*

The pure effects properties establish that when only one economic mechanism changes while others remain constant, the decomposition should attribute changes solely to that mechanism. This provides a clear test of whether a decomposition correctly identifies underlying economic forces. Suppose only the transition kernel changes (perhaps due to technological innovation). In this case, we argue that a proper decomposition should attribute all distributional changes to within-firm effects. When only exit probabilities change (perhaps due to regulatory shifts), a proper decomposition should attribute all distributional changes to exit effects. When only entry patterns change (perhaps due to changes in entrepreneurship policies), a proper decomposition should attribute all distributional changes to entry effects.

The completeness property requires that a decomposition fully accounts for all observed changes in the productivity distribution without unexplained residuals. It precludes residual terms that can be economically meaningful in magnitude but difficult to interpret.

The independence of irrelevant alternatives (IIA) property requires that the estimated contribution of one mechanism should not be affected by changes in other mechanisms. This independence ensures that decomposition results remain stable and interpretable even when multiple forces are at work simultaneously.

These properties collectively establish a standard against which different decomposition methods can be evaluated. When a decomposition violates these properties, it may provide misleading signals about the economic forces driving productivity changes, potentially leading to misguided policy implications. In the following sections, we will introduce our decomposition approach and demonstrate how it satisfies these desirable properties.

2.3 A Decomposition of Changes in the Distribution of TFP

Consider an economy that transitions from regime A to regime B. In the following, we interchangeably refer to a change in a primitive as a "channel." To isolate the contribution of each channel, we construct counterfactual distributions that allow only one primitive to change at a time while holding all others constant. We can then decompose the effect of a regime change.

Definition 2. *(Distributions by channel) Given the initial productivity distribution π_0, primitives under Regime $A \coloneqq \{P_A, Q_A, \lambda_A G_A\}$, and primitives under Regime $B \coloneqq \{P_B, Q_B, \lambda_B G_B\}$, we define:*

1. *Distribution with no channels:*
$$\pi'_A(z) = \pi'(z|\pi_0, \textit{Regime } A). \tag{2}$$
2. *Distribution with all channels:*
$$\pi'_B(z) = \pi'(z|\pi_0, \textit{Regime } B). \tag{3}$$
3. *Single channel distributions*
   a. *Distribution with within-all channel:*
   $$\pi'_P(z) = \pi'(z|\pi_0, P_B, Q_A, \lambda_A G_A). \tag{4}$$
   b. *Distribution with exit channel:*
   $$\pi'_Q(z) = \pi'(z|\pi_0, P_A, Q_B, \lambda_A G_A). \tag{5}$$
   c. *Distribution with entry channel:*
   $$\pi'_{\lambda G}(z) = \pi'(z|\pi_0, P_A, Q_A, \lambda_B G_B). \tag{6}$$
4. *Dual channel distributions*
   a. *Distribution with within-all and exit channels:*
   $$\pi'_{P,Q}(z) = \pi'(z|\pi_0, P_B, Q_B, \lambda_A G_A). \tag{7}$$
   b. *Distribution with within-all and entry channels:*
   $$\pi'_{P,\lambda G}(z) = \pi'(z|\pi_0, P_B, Q_A, \lambda_B G_B). \tag{8}$$
   c. *Distribution with exit and entry channels (selection):*
   $$\pi'_{Q,\lambda G}(z) = \pi'(z|\pi_0, P_A, Q_B, \lambda_B G_B). \tag{9}$$

There are two obvious approaches we could take when decomposing the effects of changes in P, Q, and λG on the distribution of productivities.

One approach is to consider direct effects. We could define the direct effect of changing P (within-all effect) as the change from π'_A to π'_P; the direct effect of changing Q (exit effect) as the change from π'_A to π'_Q; and the direct effect of changing λG (entry effect) as the change from π'_A to $\pi'_{\lambda G}$. This gives the direct effect of changing one primitive at a time. However, with this approach, the effect of changing multiple primitives simultaneously will induce 4 cross-terms. Consequently, it fails Property 4. Importantly, the cross-terms can be non-negligible in magnitude and that are difficult to interpret.

We instead choose a sequential decomposition that gives no cross terms. There are two sensible orders. First, we can change P, then Q, then λG; second, we can change Q, then P, then λG. In the first decomposition, the within-all effect is the change from π'_A to π'_P; the exit effect is the change from π'_P to $\pi'_{P,Q}$; and the entry effect is the change from $\pi'_{P,Q}$ to π'_B. The second decomposition is defined similarly but for the alternative sequence of Q, then P, then λG.

Formally, we define the effects of interest as follows:

Definition 3. *(Channel effects) Given the initial productivity distribution π_0, base on Definition 2, we define:*

1. *Within-all effect*
$$f_P(z) = \pi'_P(z) - \pi_A(z). \tag{10}$$
2. *Exit effect*
$$f_{P,Q}(z) = \pi'_{P,Q}(z) - \pi'_P(z). \tag{11}$$
3. *Entry effect*
$$f_{P,Q,\lambda G}(z) = \pi_B(z) - \pi'_{P,Q}(z). \tag{12}$$

Definition 3 motivates our decomposition:

Definition 4. *(Decomposition of changes in the distribution of TFP, OZ decomposition)*

Changes in the distribution of TFP can be decomposed as

$$\Delta\pi(z) \equiv \pi_B(z) - \pi_A(z) = f_P(z) + f_{P,Q}(z) + f_{P,Q,\lambda G}(z). \tag{13}$$

We can then assign the share of growth attributable to each component:

a. *Share of growth attributable to the within-all effect:* $\frac{f_P(z)}{\pi_B(z)-\pi_A(z)}$

b. Share of growth attributable to the exit effect: $\frac{f_{P,Q}(z)}{\pi_B(z)-\pi_A(z)}$

c. Share of growth attributable to the entry effect: $\frac{f_{P,Q,\lambda G}(z)}{\pi_B(z)-\pi_A(z)}$

The labels on the decomposition are meaningful. The within-all effect measures the effect of changing P on the distribution of productivities for all firms, including exiting firms. The selection effect measures the effect of the relationship between productivity and exit by comparing the distribution of productivities when firms receive new productivities following P_B, but firms are allowed to exit according to their initial exit probabilities Q_A, and the distribution when firms are allowed to exit according to their new exit probabilities Q_B. Finally, the entry effect measures the effect of entrants on the distribution of productivities, after the distribution for incumbents has fully transitioned to its new state.[16]

Notice that our decomposition framework satisfies all the key properties of an economically meaningful decomposition. The pure effects properties are satisfied by construction: if $P_A = P_B$, then $f_P(z) = 0$; if $Q_A = Q_B$, then $f_{P,Q}(z) = 0$; and if $\lambda_A G_A = \lambda_B G_B$, then $f_{P,Q,\lambda G} = 0$. This ensures that when a component of the underlying process remains unchanged, its contribution to distributional changes is measured as zero.

The decomposition achieves completeness because the sum of all components exactly equals the total change in the distribution without any residual terms.

Our framework also satisfies the independence of irrelevant alternatives property with respect to the within-all effect being independent of changes in exit probabilities. This is because our within-all effect $f_P(z) = \pi'_P(z) - \pi_A(z)$ compares the distribution under the new transition kernel P_B with the initial distribution while keeping exit probabilities fixed at Q_A. Therefore, the measurement of within-firm productivity improvements remains the same regardless of how exit patterns might change from regime A to regime B.

Proposition 1. *Our decomposition method satisfies properties 1-5 in Section 5.2 and thus is an economically meaningful decomposition.*

[16] In Appendix D.1, we present an alternative decomposition that alters the sequence of applying transformations. Specifically, we first define the exit effect by altering Q only, then define the within-all effect by altering P using Q as the reference point, and lastly define the entry effect. In this ordering, the exit effect now represents the pure impact of selection criteria changes, and is no longer subject to any productivity transition changes. This alternative ordering provides complementary insights while satisfying properties 1-5 in Section 5.2.

2.4 Existing Decompositions

Motivated by the same questions that concern this paper, Baily et al. (1992) (BHC), Griliches and Regev (1995) (GR), Foster et al. (2001) (FHK), and more recently, Melitz and Polanec (2015) (DOPD) have developed methods to decompose aggregate productivity changes into distinct components arising from intensive and extensive margins. This section examines these existing approaches and demonstrates why they fail to satisfy the properties required for an economically meaningful decomposition.

The aforementioned decomposition methods share a common conceptual framework: they partition changes in market-share-weighted average firm-level TFP into contributions from continuing, entering, and exiting firms. These methods differ primarily in their choice of reference points when measuring these contributions.[17]

Despite these refinements in reference points, all existing methods suffer from critical limitations that prevent them from satisfying the desirable properties laid out in Section 2.3. For expositional purposes, we focus on BHC-FHK and show how it violate different properties.

Using our notation, the change in market-share-weighted average firm-level TFP can be defined as

$$\Delta TFP \equiv \int_{z' \in \mathcal{Z}} z' s'(z') \pi'(z') dz' - \int_{z \in \mathcal{Z}} z s(z) \pi(z) dz$$

where $s(\cdot)$ and $s'(\cdot)$ are market shares in the first and second periods, respectively. Each period, market shares must integrate to one (i.e., $\int_{z \in \mathcal{Z}} s(z)\pi(z)dz = 1$ and $\int_{z' \in \mathcal{Z}} s'(z')\pi'(z')dz' = 1$). Moreover, the distribution of firm productivities $\pi'(z)$ evolves according to Equation 1.

Defining $\bar{Z} \equiv \int_{z \in \mathcal{Z}} z s(z) \pi(z) dz$ as the reference point, then the terms in BHC-FHK can

[17] The seminal BHC decomposition does not establish reference points for entry and exit terms. As long as the data panel is unbalanced (with firms entering or exiting), there will be non-zero entry and exit components in the decomposition. Later methods attempt to improve on BHC by introducing various reference points. GR uses the mean of aggregate TFP across periods as the reference point. FHK uses the first-period aggregate TFP as the reference point. DOPD argues that these earlier reference points produce biased results and instead proposes using first-period survivor TFP as the reference point for exit and second-period survivor TFP as the reference point for entry (see Melitz and Polanec (2015) for details).

be written using the underlying primitives as:

$$\text{Within} \equiv \int_{z\in\mathcal{Z}}\int_{z'\in\mathcal{Z}} \Big(z'-z\Big)s(z)\omega(z,z')dz'dz \tag{14}$$

$$\text{Between} \equiv \int_{z\in\mathcal{Z}}\int_{z'\in\mathcal{Z}} \Big(z-\bar{Z}\Big)\Big(s'(z')\alpha_1 - s(z)\Big)\omega(z,z')dz'dz \tag{15}$$

$$\text{Cross} \equiv \int_{z\in\mathcal{Z}}\int_{z'\in\mathcal{Z}} \Big(z'-z\Big)\Big(s'(z')\alpha_1 - s(z)\Big)\omega(z,z')dz'dz \tag{16}$$

$$\text{Entry} \equiv \int_{z'\in\mathcal{Z}} (z'-\bar{Z})\omega_e'(z')dz' \tag{17}$$

$$\text{Exit} \equiv -\int_{z\in\mathcal{Z}} \big(z-\bar{Z}\big)\,\omega_x(z)dz \tag{18}$$

where $\omega(z,z') = \pi(z)P(z'|z)(1-Q(z'))$. Moreover, $\omega_e'(z') = \alpha_1 s'(z')\lambda G(z')$, and $\omega_x(z) = s(z)\pi(z)\int_{z'\in\mathcal{Z}} P(z'|z)Q(z')dz'$. The sum of five terms above is equal to the change in market-share-weighted average firm-level TFP.

Looking at different BHC-FHK decomposition terms defined using our notation, we can immediately see that this decomposition fails the pure effects properties. To clarify, the pure within-all effect property requires that if only the transition kernel $P(z'|z)$ changes, the within-all term should explain all of the change in aggregate TFP. However, we can see from equation Equation 18 that $P(z'|z)$ also enters the exit effect. As long as the exit probability $Q(z)$ is not constant, the exit effect will be contaminated in this decomposition. In addition, the pure exit effect property requires that if only the exit probability changes, the exit term should explain all of the change in aggregate TFP. However, as we can see from Equation 14, Equation 15 and Equation 16 that $Q(z)$ also enters the "within", "between" and "cross" terms. Lastly, the pure entry effect property requires that if only the mass and productivity distribution of entrants changes, the entry term should explain all of the change in aggregate TFP. However, changes in firm entry will drive changes in the market share under the new regime $s'(z)$. Since $s'(z)$ enters both the "between" and "cross" terms, BHC-FHK does not satisfy the pure entry effect property.

BHC-FHK violates the "no residual" property by having the "cross" term (see Equation 16), which lacks economic interpretation. This term mechanically captures interactions between changes in productivity and changes in market share among incumbents but pro-

vides no meaningful insight. While such second-order effects might be negligible in continuous time, they become magnified with annual frequency data when productive firms gain market share from less productive competitors. The cross term thus undermines the decomposition's purpose of providing economically meaningful attribution of the source of productivity growth.

The contamination across decomposition terms also causes BHC-FHK to violate the Independence of Irrelevant Alternatives (IIA) property. Since the pure effects properties fail, the contribution attributed to one primitive becomes dependent on changes in other primitives. This interdependence fundamentally compromises the decomposition's ability to provide stable, economically interpretable attributions of the source of productivity growth.

These violations arise from inherent limitations in the methodological approach: existing decompositions rely primarily on accounting identities derived from observed data, without explicitly modeling the underlying productivity transition process. Without accounting for unobserved counterfactual scenarios, these methods necessarily incorporate implicit assumptions to complete their analysis. One crucial implicit assumption they make is that the productivities of exiting and surviving firms grow at the same rate. Furthermore, the algebraic grouping of terms - first constructed from observed accounting identities and subsequently assigned economic interpretations - results in components whose economic significance may be difficult to interpret clearly in relation to the fundamental drivers of productivity dynamics.

A key limitation is their treatment of exit effects. An economically meaningful exit effect should fully capture exiting firms' contribution to aggregate TFP changes, including the critical role of freeing up resources for surviving and newly-entered firms. However, exit effects measured using compositional shifts in the observed firm distribution fail to capture these resource reallocation dynamics, leading to results that are not only non-interpretable but potentially severely biased.

The consequences of violating these properties, and the resulting departure from an economically meaningful decomposition of changes in aggregate TFP, can be arbitrarily large. As we demonstrate in Appendix E through a simple firm dynamic model with endogenous exit, the components defined and computed from traditional decompositions are not merely

theoretical concerns but can lead to substantially misleading conclusions about the drivers of productivity growth.

While we have demonstrated the ability of our proposed decomposition to identify within-all, exit, and entry effects through given primitives in both regimes, implementing this decomposition with actual data presents challenges. In practice, we only observe the productivity dynamics of surviving firms, but the productivity dynamics of exiting firms - crucial for measuring exit effects - remain unobserved. The following section addresses this empirical challenge by developing a systematic procedure to recover the complete firm dynamics from observable data.

3 Estimation Strategy

Estimating the productivity transition process (P), exit probabilities (Q) and distribution and mass of entrants (λG) from observed data presents two key challenges.

First, we want our approach to be able to accommodate a broad class of primitives. We therefore pursue a discretized, non-parametric estimation strategy that estimates all three primitives under a weak set of assumptions. As robustness, we also consider an AR(1)-style process, which enables us to estimate a continuous process at the cost of flexibility.

Second, the non-parametric approach presents computational challenges. The system of equations relating observed transitions ($\tilde{P}$) and observed exits rates ($\tilde{Q}$) to the underlying, unobserved parameters (P and Q) is ill-conditioned in the sense that small perturbations in the data can lead to large changes in the estimated parameters. We introduce a novel regularization approach that penalizes the condition number of the estimated transition matrix P, ensuring stable and economically meaningful estimates.

These considerations motivate the empirical implementation strategy we introduce in this section. We begin by reiterating our statistical framework, and derive some additional key expressions that are crucial for estimating the primitives from data. For tractability, we consider a discrete model that preserves the key features of our framework. We introduce a non-parametric estimator of the underlying parameters, establish identification, and derive asymptotic properties of the estimator. We finally show through numerical examples how

our penalized maximum likelihood approach overcomes the ill-conditioned problem.

3.1 Identification

We now reiterate the statistical model from our conceptual framework described in the previous section and derive key equations that will be used when estimating the model. Suppose that firm productivity evolves as a Markov process, and each period, firms first draw a new productivity, then exit with a probability that is a function of their new productivity draw. We write that productivity evolves according to the Markov kernel $P(z'|z)$ and firms exit with probability $Q(z')$.

We consider a discrete-state version of the model, and prove identification of all parameters and derive asymptotic properties. While the process is discrete, the process can always be made more granular to more closely approximate a continuous process.

This model is tainted by selection, so observed transitions and exits are not necessarily aligned with the underlying process. More formally, data is tainted by selection in the following sense:[18]

$$\tilde{Q}(z) = \sum_{z'} P(z'|z)Q(z'), \quad \text{and} \tag{19}$$

$$\tilde{P}(z'|z) = \frac{P(z'|z)(1-Q(z'))}{\sum_{z''} P(z''|z)(1-Q(z''))}. \tag{20}$$

$\tilde{Q}$ gives the observed exit rates conditional on the last observed productivity of exiting firms. To arrive at this expression, note that from their last observed productivity, exiting firms first transition to their new productivities in the next period, then exit based on rates that are a function of their new productivity draw. $\tilde{P}$ gives the observed transition rates conditional on firm survival.[19] The numerator in Equation 20 captures the share of firms

[18]The notion of a selection effect (i.e., $P \neq \tilde{P}$ and $Q \neq \tilde{Q}$) depends on the timing in the model. The model we propose includes a selection effect. If we instead consider a model where each firm exits conditional on its first period productivity, then draws its second period productivity, there is no selection effect. These two models are observationally equivalent, so we follow the timing of Hopenhayn (1992), which will give a selection effect and is more plausible in data with long intervals between observations, such as in our context.

[19]If the exit rate is constant (i.e., $Q(z) = c, \forall z \in \mathcal{Z}$), then the system of equations becomes trivial and $P(z'|z) = \tilde{P}(z'|z)$. Conversely, if the observed exit rate $\tilde{Q}(z)$ is not constant, then the true exit rate $Q(z)$ is also not constant. Moreover, $Q(z)$ must be more dispersed than $\tilde{Q}(z)$ since $\tilde{Q}(z)$ is a weighted average of

that both transition from z to z' and survive. The denominator normalizes these shares to ensure the transition rates sum to one.

The observed distribution of productivities for entrants is $\tilde{G}(z)$. Even though the entry distribution might also be tainted by selective entry, from an empirical perspective we do not observe potential producers who do not enter the market, making it more difficult to measure this dimension of the selection effect without additional structure. We choose not to separately model the selection process for entrants; we instead assume that $\tilde{\lambda}\tilde{G}(z)$ serves as an estimator for the true productivity distribution of entrants $\lambda G(z)$.

Let $\tilde{P}$ represent the matrix form of $\tilde{P}(z'|z)$ and $\tilde{Q}$ represent the vector form of $\tilde{Q}(z)$. We introduce the following assumption:

Assumption 1. *$diag(1-\tilde{Q})\tilde{P}$ is full rank, where diag denotes the diagonalization operator.*

Assumption 1 implies that $diag(1-\tilde{Q})$ and $\tilde{P}$ are both individually full rank. The full rank condition on $diag(1-\tilde{Q})$ means that there cannot be a productivity bin in which exit is guaranteed tomorrow. The condition on $\tilde{P}$ imposes linear independence of observed transitions across productivity bins. This prevents, for instance, two productivity bins from having identical observed transition rates.

This leads us to our first proposition:

Proposition 2. *Suppose that Assumption 1 holds. Then we have*

a. $Q = 1 - 1 \oslash \left(diag(1-\tilde{Q})\tilde{P}\right)^{-1} 1$, *where* 1 *denotes the vector of ones and* $\oslash$ *denotes element-wise division; and*

b. $P = diag(1-\tilde{Q})\tilde{P}diag(1-Q)^{-1}$.

We have the following corollaries:

Corollary 1. *(Consistency) Denoting the matrix of empirical transitions as $\hat{\tilde{P}}$ and vector of empirical exit rates as $\hat{\tilde{Q}}$, then:*

a. $\hat{Q} = 1 - 1 \oslash \left(diag(1-\hat{\tilde{Q}})\hat{\tilde{P}}\right)^{-1} 1$ *is a consistent estimator of Q, and*

$Q(z)$.

b. $\hat{P} = diag(1 - \hat{\tilde{Q}})\hat{\tilde{P}}diag(1 - \hat{Q})^{-1}$ *is a consistent estimator of* P*.*

Corollary 2. *(Asymptotic normality) There is a variance-covariance matrix* Ω *such that*

$$\sqrt{n}\left(\begin{bmatrix}\hat{Q}\\ vec(\hat{P})\end{bmatrix} - \begin{bmatrix}Q\\ vec(P)\end{bmatrix}\right) \xrightarrow{d} N(0, \Omega),$$

where vec denotes the vectorization operator.

3.2 Implementation

We estimate the model via a modified maximum likelihood approach. This allows us to achieve two objectives. First, we can impose constraints on $\hat{P}$ and $\hat{Q}$, guaranteeing values fall in the interval $(0, 1)$. Second, we can add a penalty on the condition number of P, addressing the ill-conditioning problem which will be discussed in detail later.

Corollary 1 to Proposition 1 gives consistent estimators for Q and P. However, they have poor small sample performance. Two primary issues arise. First, **the constraint problem**: the estimators $\hat{Q}$ and $\hat{P}$ that solve the system of equations are unconstrained. Second, **the ill-conditioned problem**: the solution is unstable in the sense that small perturbations in either $\tilde{P}$ or $\tilde{Q}$ can cause large perturbations in $\hat{Q}$ and $\hat{P}$.

To bypass **the constraint problem** we estimate using maximum likelihood and add constraints that $\hat{Q}$ gives valid probabilities and $\hat{P}$ is a valid transition matrix. Then, we can use equations 19 and 20 to back out the associated $\tilde{P}$ and $\tilde{Q}$. We can then compute the log-likelihood of the observed data:

$$\begin{aligned}\mathcal{L}(X|P, Q, \tilde{Q}) = \sum_{i=1}^{n} 1\{i \text{ is exiter}\} \log\Big(\tilde{Q}(X_1^i)\Big) \\ + 1\{i \text{ is incumbent}\}(\log\big(P(X_2^i|X_1^i)\big) + \log\big(1 - Q(X_2^i)\big))\end{aligned} \tag{21}$$

The main benefit of this approach is that when going from $\tilde{P}$ and $\tilde{Q}$, it is possible to back out estimates for $\hat{P}$ and $\hat{Q}$ that are not valid probabilities. However, equations 19 and 20 guarantee that any valid $\hat{P}$ and $\hat{Q}$ will imply valid $\tilde{P}$ and $\tilde{Q}$.

To bypass **the ill-conditioned problem**, we penalize the condition number of $\hat{P}$. This is motivated by the matrix form of equation 19, $PQ = \tilde{Q}$. This is a linear equation, with Q as the unknown. If we write this as $Q = P^{-1}\tilde{Q}$, we can see that the condition number of P^{-1} (which is also the condition number of P) bounds the effect of deviations in $\tilde{Q}$ on Q. Penalizing the condition number of $\hat{P}$ thus ensures we estimate parameters that satisfy the requirement that small deviations in observed data lead to small deviations in the underlying parameters, and avoid parameters that have the property that small deviations in observed data lead to large deviations in the underlying parameters. Beyond intuition, in Appendix F.1, we show that our penalized maximum likelihood estimator noticeably improves estimation stability and accuracy.

Previous research has estimated transition matrices using regularization methods (see Kreinin and Sidelnikova (2001)). Regularization is also common in other contexts. As an example, Tikhonov Regularization is a well known regularization technique for estimating Fredholm integral equation of the first kind (FIE1) (see Groetsch (1984), and Appendix F.2 for an application using our data). However, to our knowledge, this is the first paper to estimate a matrix using maximum likelihood with a penalty on the condition number. We believe this is an important contribution because of the gains in small sample precision we observe in numerous experiments.

As a robustness check, we also estimate a continuous, parametric AR(1)-style model. See Appendix G.1 for details.

In the previous section, we introduced a conceptual framework for studying changes in the distribution of firm-level productivities over time and decomposing these changes into within-all, exit, and entry channels. In this section, we demonstrated how to recover the necessary primitives to implement this decomposition in practice from observed data. Building on this, the following sections illustrate the empirical relevance of our framework and estimation method by applying them to China's state-owned enterprise (SOE) reform - a large-scale economic transformation in the late 1990s and early 2000s that provides an ideal setting to test our approach. This reform directly impacted all three channels in our framework: it aimed to improve the operational efficiency of existing firms (within-all channel), imposed harder budget constraints on underperforming enterprises (exit channel),

and created opportunities for new entrants (entry channel).

We will describe the institutional background of the reform, outline our data sources, present empirical findings on productivity catch-up of SOEs with private firms and several patterns that point to selection effects as the primary driver of productivity improvements in the state sector, and finally apply our decomposition methods to quantify the relative contribution of each channel to productivity growth in the state sector during this pivotal transition period.

4 Institutional Background

The reform of state-owned enterprises (SOEs) in China began in 1978 and has since played a central role in the country's economic transformation. From 1978 to 1993, reforms focused on incremental changes within the planned economy framework, including increasing enterprise autonomy, decentralizing control to local governments, and instituting managerial reforms like the "managerial responsibility system" that links managers' and workers' income to the financial outcomes of the enterprises where they work. A dual-track pricing system was implemented, gradually exposing SOEs to market forces while maintaining social and economic stability. However, these reforms embodied a "reform without losers" approach - to avoid social instability, the government typically instructed state-owned banks to bail out struggling SOEs rather than allow them to fail or lay off workers. This "soft budget constraint" weakened the incentives for SOEs to maximize profits, and consequently induced sub-optimal operational efficiency. The absence of exit mechanisms also prevented market selection from serving as a force for improving aggregate productivity in the state sector (Lau et al. (2000), Brandt and Zhu (2001)). Despite these limitations, the reforms did have some positive effects on productivity (Groves et al. (1994), Jefferson et al. (1996), Li (1997)).

However, by the early 1990s, SOEs had become unprofitable and were draining local government budgets (Berkowitz et al. (2017)). By 1998, the state sector was in crisis, with 40.6% of SOEs incurring losses totaling 1.4% of GDP (Lardy (2014)). It had become clear that the strategy of "reform without losers" could no longer be sustained. As a result, the government launched a more radical reform strategy known as "grasp the large, let go of the

small" (zhuā dà fàng xiǎo) in the 1999 Fourth Plenum of the Communist Party's Central Committee.

The reform strategy followed a dual approach. For large, centrally controlled SOEs, the government pursued consolidation and corporatization, creating larger, more competitive enterprise groups. This "grasp the large" approach involved not only restructuring existing firms into limited liability or joint-stock companies but also establishing new large state-owned enterprises in strategic sectors.[20] Meanwhile, small and medium-sized SOEs, mostly loss-making, were either privatized or closed under expanded local government authority.[21] This dual approach resulted in fewer but larger central SOEs and a significant reduction in small state enterprises.

The State-owned Assets Supervision and Administration Commission (SASAC), established in 2003, marked a pivotal reform in China's state enterprise sector. As the formal owner of all (with some exceptions) state-owned enterprise groups, SASAC implemented comprehensive corporate governance reforms including board restructuring, executive performance systems, and asset protection policies. Central SASAC firms were particularly well-positioned during the following years of robust economic growth, given their concentration in booming sectors: natural resources, energy, power, and telecommunications. This strategic positioning yielded exceptional results - SASAC firms achieved approximately 40% annual profit growth, earning 628 billion RMB in 2005, equivalent to 3.4% of China's total GDP (Naughton (2006)).

5 Data

Our analysis uses the Annual Survey of Industrial Firms (ASIF) from China for the years 1998-2007. The ASIF is an annual survey of manufacturing enterprises conducted by the National Bureau of Statistics of China, and is one of the most important sources of informa-

[20] For example, in the steel industry, five large industrial groups were created in the late 1990s and early 2000s. In the automobile industry, state-owned automobile companies were consolidated into six state-owned conglomerates. See Hsieh and Song (2015) for details.

[21] For example, in May 2002, the State Tobacco Monopoly Administration ordered all state-owned firms producing less than 100,000 cigarette cases per year to be closed down, whereas state-owned firms with an annual production below 300,000 cases were "encouraged" to merge with larger firms. See Rubens (2023) for details.

tion on the performance of the manufacturing sector in China.[22] The data is ideally suited for illustrating the advantages of our statistical framework for several reasons. First, the annual frequency of observations matches the typical setting where productivity decompositions are conducted but firm trajectories between time periods are unobserved. Second, the data captures rich variation in entry and exit patterns, with firm coverage expanding from around 100,000 to over 300,000 firms over the sample period. This extensive turnover provides essential variation for identifying our selection parameters. Third, comprehensive coverage of state-owned enterprises allows construction of representative counterfactual distributions. The additional comprehensive coverage of large private firms (those exceeding 5 million RMB in revenue) allows thorough comparison of state and private firm performance.

A key feature of the data that motivates our methodological approach is the prevalence of exit. Among firms present in 1998, approximately 80% exit by 2007, with exit rates varying systematically with productivity levels. Moreover, many firms show declining productivity in the years before exit, highlighting the importance of properly accounting for selection when decomposing aggregate productivity changes.

The unit of observation in the data is a registered firm in a given year. Within state-owned industrial groups, each firm owned by the group is a separate observation in our data. Our baseline variables of interest include the firm's industry at the four-digit level,[23] location, age, ownership status, total assets, total liabilities, industrial production, value added, revenue, number of employees, wage payments, capital stock, intermediate goods purchase, export revenues, and subsidies.[24] After cleaning the data to remove observations with missing or implausible values for key production variables, our final sample contains over 1.7 million firm-year observations.

Our data treatment follows Hsieh and Klenow (2009) (henceforth HK). For the baseline analysis, we define capital stock as the book value of fixed capital net of depreciation. As an alternative, we employ the perpetual inventory method with an annual depreciation rate

[22] A panel is constructed using a method that almost directly follows the approach outlined in Brandt et al. (2012).

[23] To preserve greater sectoral granularity in our analysis, we follow Liu (2019) and map the original 4-digit industry codes to an expanded set of 80 2-digit codes, rather than using the standard classification of 30 2-digit sectors in manufacturing.

[24] Data on value added and industrial production in 2004 are imputed since they are missing in the original data.

of 9% and province-sector-specific investment rates (see the online appendix from Brandt et al. (2012)).

For labor input, we adjust the reported wage payments to account for non-wage benefits, scaling up wages to match the 50% aggregate labor share from Chinese input-output tables. As a robustness check, we also measure labor input using employment.

Our analysis requires precise identification of state-owned enterprises (SOEs). We adopt the definition from Hsieh and Song (2015), which offers a robust framework for classifying state ownership (see Appendix B.1 for a detailed discussion of the complexities in defining state ownership in China).

For our baseline analysis, we measure firm-level productivity following HK. However, recognizing that their assumptions of constant returns to scale and homogeneous factor shares within sectors may be restrictive, we also employ an alternative approach. Following Brandt et al. (2017), we estimate production functions that allow for varying returns to scale and heterogeneous factor shares, deriving firm-level revenue TFP (TFPR) as a residual (see Appendix B.2). This alternative specification serves as a robustness check for our main results.

6 Empirical Evidence: the Puzzle of SOE Performance

This section presents and resolves a conceptual puzzle that arises from two empirical patterns observed during China's state-owned enterprise (SOE) reforms from 1998 to 2007. These patterns hold in virtually all industrial sectors during the period we study, and appear contradictory. First, we establish a convergence in productivity distributions between SOEs and private firms. Second, we show that this apparent productivity improvement coincided with a dramatic decline in SOE market presence.

The combination of improving productivity but declining market presence presents a conceptual challenge. In a neoclassical framework, if SOEs were genuinely achieving productivity convergence with private firms while retaining their preferential access to bank credit

and government contracts - as is well documented in the literature (see Bai et al. (2006), Riedel et al. (2020), Song et al. (2011b), Song et al. (2011a), and Liu et al. (2021)) - we would expect them to at least maintain their relative market position against comparable private incumbents, even in an increasingly competitive environment.

To rationalize these two seemingly contradictory patterns, we propose a resolution: selective firm exit. Under this hypothesis, the apparent productivity improvements among surviving SOEs would reflect selection effects rather than genuine competitive gains - the weaker SOEs exited while even the "stronger" incumbents steadily lost ground to private firms in the marketplace. Going to the data, we document systematic patterns in firm exit and survival that suggest selection effects may have been pervasive during the transition period we study, and potentially played a key role in driving the convergence patterns observed in the data.

With these empirical facts established, in the next section we will use our decomposition method to quantify the contribution of within-all, exit, and entry effects in driving TFP growth among SOEs during the transition period, and will show that our decomposition is able to match the empirical trends we observe in the data.

We are not the first to document the catch-up phenomenon of state-owned firms with private firms or the decline in market presence of state firms.[25] However, we expand upon previous research by showing that the convergence occurs not just for the mean, but for the entire distribution of firm productivities across ownership types between 1998 and 2007, and we are the first to link these two seemingly contradictory patterns together and provide a comprehensive explanation with the objective of resolving this apparent puzzle.

Note that throughout this section, when we refer to "privately-owned firms" or "private firms," we are using these terms interchangeably with "non-state owned firms". This category includes all enterprises that are not state-owned, such as privately-owned domestic firms, foreign-owned enterprises, and collectively-owned enterprises. While our descriptive analysis

[25] Hsieh and Song (2015), Chen et al. (2021) and Brandt et al. (2022) have similar findings that state-owned and privatized firms started from significantly lower TFP levels but then caught up with always-private firms by 2007. Brandt et al. (2022) find that the TFP gap between state-owned and private manufacturing firms narrowed and even became slightly positive after 2007. However, Chen et al. (2021) and Brandt et al. (2022) interpret the convergence in TFP across ownership types with caution. See Appendix C for detailed discussion.

considers all firm ownership categories, our model results focus on comparing firms that maintained consistent ownership throughout the period - either "forever state-owned" or "forever private-owned" - to isolate ownership effects from transition effects.[26]

6.1 TFP Convergence Across Ownership Types during 1998 and 2007

Fact 1: Between 1998 and 2007, there is a convergence in TFP distributions across all ownership types in nearly all manufacturing sectors.

Figure 1 illustrates Fact 1, focusing on two representative sectors: grain milling and garment manufacturing. The figure traces the evolution of firm productivity distributions in these two industries, comparing three distinct ownership categories: forever SOEs, forever private-owned (forever POEs), and privatized firms (enterprises that transitioned from state to private ownership). These two sectors serve as exemplars of broader patterns observed across the manufacturing sector.[27]

In 1998, both grain milling and garment manufacturing exhibit similar patterns: the productivity distribution of SOEs lies to the left of POEs, indicating lower average productivity among state-owned firms. The TFP distributions for SOEs also display greater dispersion, suggesting more varied performance within the state sector. Privatized firms in 1998 show distributions more closely aligned with POEs, though still slightly lagging in average productivity.

By 2007, both industries had undergone dramatic changes. First, a substantial shift from state to private ownership is observed in both sectors. In the grain milling industry, the number of SOEs decreases from 606 to 218, while the number of POEs increases from

[26]It is important to acknowledge the role of privatization in China's economic growth during this time. While privatization of state-owned enterprises has been shown to be a significant driver of the aggregate productivity growth and economic transformation seen between 1998 and 2007 (see Chen et al. (2021) and Wang et al. (2024)), our focus in this study is on comparing the performance trajectories of firms that did not undergo ownership changes. By concentrating on these stable ownership categories, we can better understand how firms with persistent state ownership evolved in terms of productivity relative to the private sector benchmark, providing insights into the effectiveness of reforms within the state-owned sector itself.

[27]Due to space constraints, we present only these two examples in the main text. However, the patterns demonstrated in these cases are consistently observed across the majority of manufacturing sectors analyzed in our study.

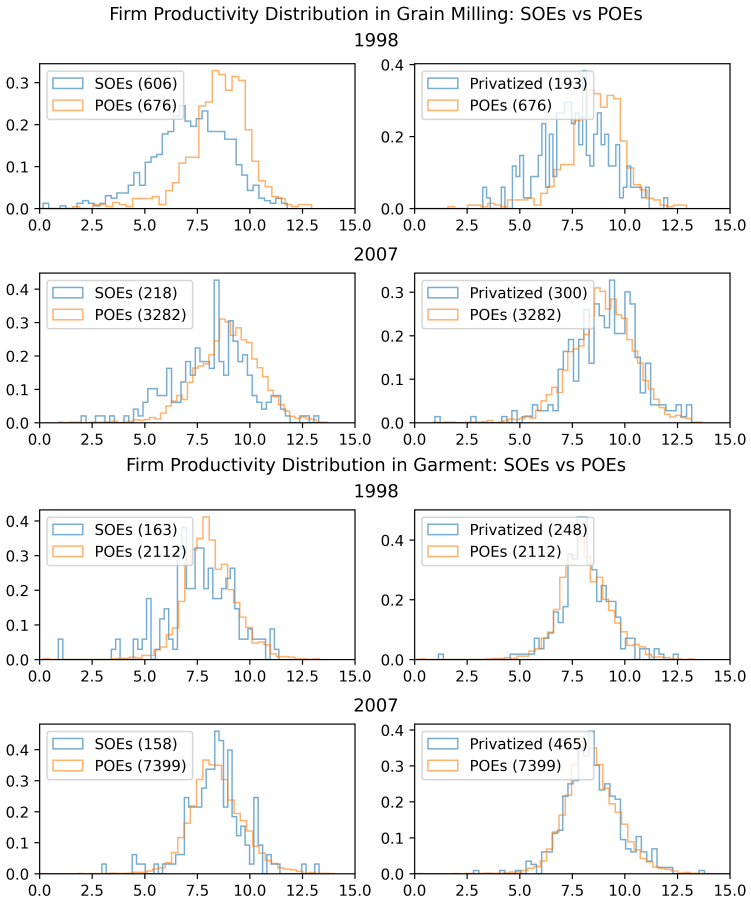
676 to 3,282. The garment industry sees a similar trend, with the number of SOEs slightly decreasing from 163 to 158, while POEs surge from 2,112 to 7,399. This shift underscores the significant transformation in ownership structure across these industries.

Second, the performance of surviving SOEs improved markedly in both sectors. Their productivity distributions shift rightward and become less dispersed, suggesting that less efficient SOEs were either closed or privatized, leaving a core of more productive state-owned firms.

Privatization also plays a key role in this transformation. The number of privatized firms increases in both industries, from 193 to 300 in grain milling and from 248 to 465 in the garment sector. Notably, the productivity distributions of these privatized firms in 2007 closely resemble those of POEs.

Finally, overall industry productivity growth is evident in both sectors, with the productivity distributions of all firm types moving rightward, signifying increased efficiency across the board. This shift appears more pronounced in the grain milling industry, suggesting the different institutional details and market structure can drive variation in rates of convergence across different sectors.

Figure 1: Within-Industry Convergence in TFP distributions



Notes: We define SOEs/POEs as firms that maintain state/private ownership throughout the entire sample period. Privatized firms are those that transition from state to private ownership at some point during the sample period. Numbers in parentheses represent the number of firms under different ownership types in each sector.

Having established the uniform trend of within-sector productivity distribution convergence across all ownership types, we now provide insights into the pace and extent of SOE productivity improvements relative to the private sector benchmark.

Fact 2: Between 1998 and 2007, sectors that began with larger TFP gaps between private and state firms saw faster rates of convergence between the two groups.

Figure 2 illustrates the rate of TFP convergence across ownership types for all manufacturing sectors (excluding tobacco, as it is a state-controlled sector). The top plot shows the relationship between the initial productivity gap in 1998 and subsequent productivity growth differentials between forever SOEs and POEs, while the bottom plot presents the

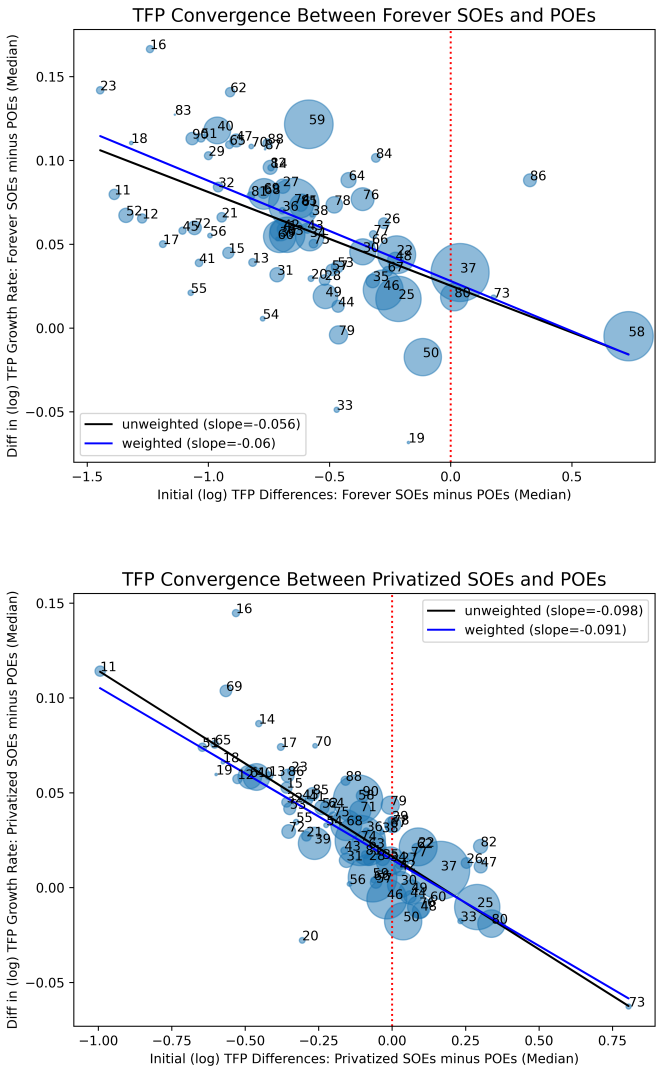
same relationship for privatized firms and POEs.

In both plots, the horizontal axis represents the initial TFP differences in 1998, measured as the log difference between the median TFP of the respective ownership type (SOEs or privatized firms) and POEs in each sector. Negative values indicate that SOEs or privatized firms started with lower productivity compared to POEs in 1998, which is the case for most sectors as evidenced by the majority of data points lying to the left of the vertical dotted line at $x = 0$. The vertical axis shows the difference in annual TFP growth rates between the ownership types, calculated as the growth rate of the median productivity of SOEs or privatized firms minus that of POEs in the sector. Positive values indicate faster TFP growth for SOEs or privatized firms relative to POEs in subsequent years.

Each numbered bubble represents a different manufacturing sector, with the bubble size indicating the sector's market share. The wide spread of these bubbles suggests considerable heterogeneity in both initial productivity differences and subsequent growth patterns across sectors. While at the beginning of the reform period, in most sectors SOEs and privatized firms are less productive than their private counterparts, there is a mix of positive and negative values on the y-axis, indicating varied performance in productivity growth relative to POEs.

The downward-sloping trend lines in both plots imply a convergence effect. In the top plot, the slope of -0.056 for forever SOEs indicates that for every additional log point gap in median TFP between state-owned and private firms within a sector, SOEs within that sector tended to see 5.6% faster annual TFP growth relative to POEs during the transition period. This pattern is consistent with a catching-up phenomenon. The bottom plot reveals a similar trend for privatized firms, but with a steeper slope of -0.098, suggesting an even faster (9.8%) convergence rate. Synthesizing the faster catch-up rate of privatized firms with POEs with the finding in Figure 1, we conclude that privatization appears to be an effective mechanism for enhancing productivity in formerly state-owned enterprises. Without considering selection into privatization, the privatization process seems to unlock efficiency gains that allow these firms to bridge the productivity gap with POEs more quickly than their counterparts that remain under state ownership.

Figure 2: TFP Convergence among All Ownership Types



Notes: The size of each marker represents the market share of the sector in 1998. Median gives the median productivity of firms in the sector. We exclude tobacco manufacturing (24) since it is a state-controlled sector. The unweighted line gives an unweighted OLS fit through the data. The weighted line weights by the market share of the sector in 1998.

6.2 Declining Market Share of State-Owned Enterprises

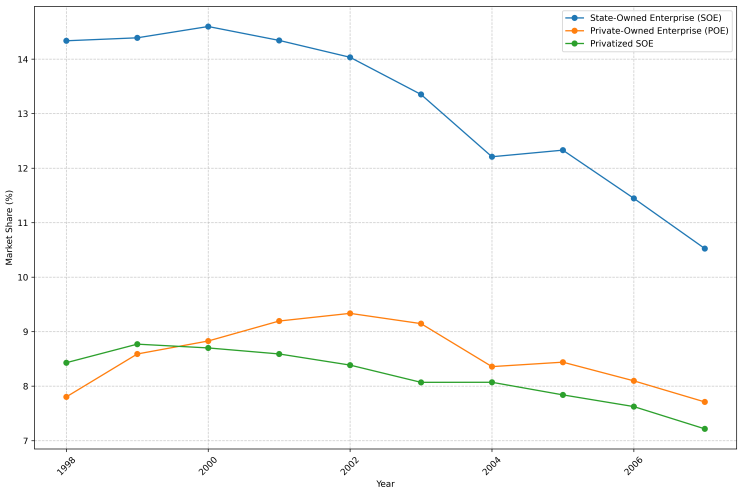
Fact 3: Between 1998 and 2007, across nearly all sectors, state-owned enterprises saw declining market share relative to their private counterparts.

Despite the productivity increase among SOEs documented in the previous subsection, during the transition period we study, the state sector simultaneously experienced a dramatic

decline in market share across nearly all major manufacturing sectors. In grain milling, for instance, SOE market share fell from nearly 60% in 1998 to below 15% by 2007. Similar patterns emerge in other sectors: feed processing saw a decline from 35% to about 10%, while sugar processing witnessed an even more dramatic drop from over 80% to approximately 35% (see Figure 1 in Appendix A.1). For all sectors, our data show SOEs' share of industrial output fell from 54.0% in 1998 to 36.7% by 2007, with their share of value added declining from 57.6% to 37.0%, similar to the findings in Hsieh and Song (2015).

We also compare the market shares of incumbent state and private firms, those present since 1998 (Figure 3). We notice that while incumbent SOEs saw their market share decline by one-third, incumbent private firms maintained or even grew their market share over the same period.[28] Note that since the set of firms is held fixed for this figure, these numbers can be interpreted as proportional to firm-specific market shares - meaning that this trend is not driven by entry and exit. This contrast in maintaining market shares between state-owned and private incumbents, which holds across industries (see Figure 2 in Appendix A.1), reveals that surviving SOEs struggled to compete even in sectors with limited private entry.

Figure 3: Market Share of Incumbent Firms: 1998-2007



[28]In fact, the result is even stronger, since the data include only sufficiently large POEs. So if some POEs fell below the threshold for inclusion in the data without fully exiting the market, we incorrectly classify them as exiters, and thus understate the market share of incumbent POEs over time.

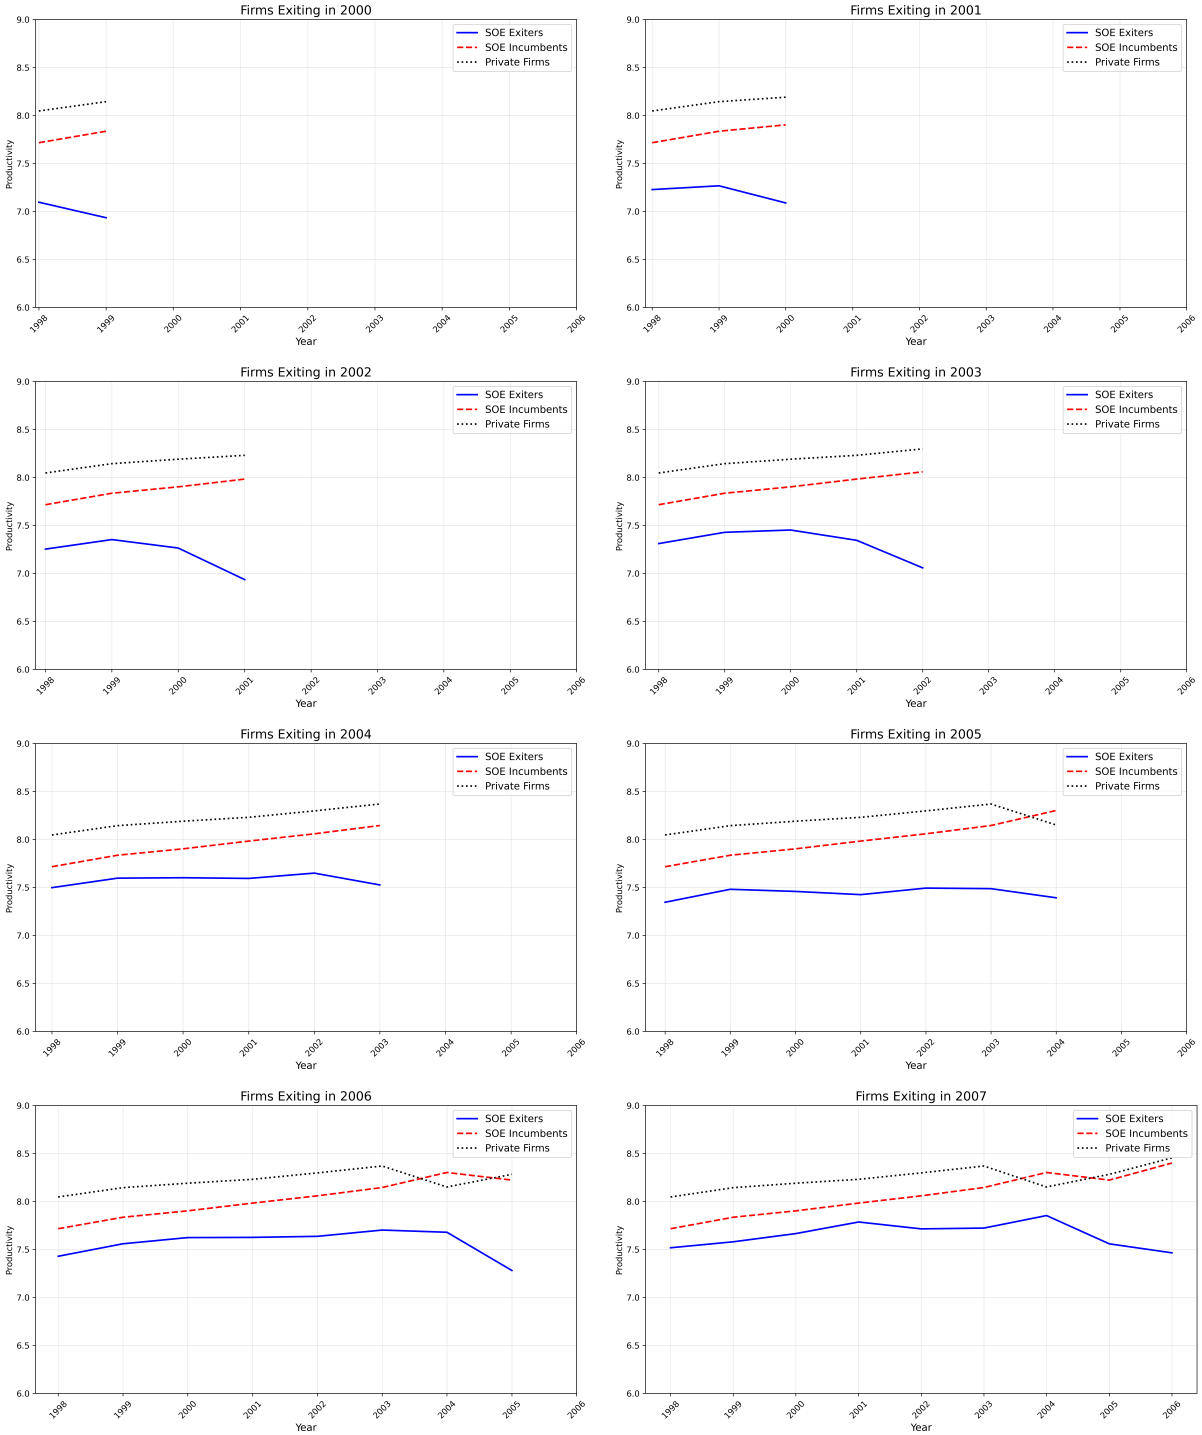
6.3 Evidence of Selective Exit among SOEs

We propose that selective firm exit can rationalize the two seemingly contradictory patterns documented in the previous two subsections.

The selection effect, or survivorship bias, stems from the fact that firms make exit decisions based on factors that are unobserved in the data, potentially because of gaps between when data is collected and when exit decisions are made. This bias is particularly crucial in the context of SOE reforms during this period. As part of these reforms, the government divested many small and medium-sized SOEs that were financially unsound or underperforming. The reduction in government subsidies raised the threshold for market survival, while the increase in firm-level autonomy gave unproductive firms the freedom to make the decision to shut down. Together, these two factors generated substantial exit among less productive SOEs in the late 1990s and early 2000s.

Figure 4 shows evidence of selection patterns by comparing the productivity trajectories of exiting SOEs to those of surviving SOEs and private firms. We highlight three observations. First, a persistent productivity gap exists between exiting firms (blue lines) and survivors (red lines), indicating that exit-bound firms underperformed before their departure. Second, exiting firms on average see a productivity decline one to two years before exit, reflecting real operational changes - our data show these firms actively winding down by selling capital and reducing employment and wage payments. This is further corroborated in the regressions shown in Table 1 in Appendix A.2, which show that exiting state-owned firms have productivity growth rates that are 15.62% lower than surviving state-owned firms. These results are robust to controlling for industry fixed effects. Both of these pieces of evidence indicate that exit is not a sudden event but rather the culmination of a process where firms experience significant productivity declines in their final few years of operation. Third, the selection process becomes more stringent over time - firms exiting in later years (e.g., 2005) had higher initial productivity and showed less dramatic declines prior to exit than firms that exited in earlier years.

Figure 4: Productivity Trajectories of Exiting State-Owned Firms: 1998 - 2007



Note: We track firms with more than one observation to capture pre-exit productivity trends. Year of exit is defined as the first year a firm is absent from the dataset.

7 Application: Decomposing Growth in the Chinese State Sector

In this section, we apply our decomposition framework to Chinese state-owned enterprises during 1998-2007 to quantify the role of firm exit versus efficiency improvements in driving productivity growth. We first present our benchmark estimation results and decompose aggregate productivity changes to identify the contribution of firm exit, within-firm improvements, and firm entry. We then connect these quantitative findings to the institutional context of Chinese SOE reforms to understand the economic mechanisms behind our results. We conduct multiple robustness checks across alternative specifications that validate the stability of our main findings. Finally, we compare our decomposition results with results from the widely-used Baily et al. (1992) and Foster et al. (2001) (BHC-FHK) method to demonstrate how accounting for selection effects fundamentally changes our understanding of the sources of SOE productivity growth during this critical reform period.

7.1 Benchmark Estimation Specification

For our benchmark analysis, we use the productivity measure defined by Hsieh and Klenow (2009) using an elasticity of substitution (σ) of 3 across firms within each industry. We normalize each firm's productivity relative to the median productivity in its industry and include all sample firms after removing outliers (trimming the bottom and top 1% in each year). Our analysis divides the data into nine one-year intervals: 1998-1999, 1999-2000, ..., 2006-2007. We discretize productivity into 10 bins, with bin selection balancing between uniformly spacing bins and equalizing the number of observations per bin.[29]

[29]Given the choice of N_q bins and bin edges $z_2, \ldots, z_{N_q}$ (z_1 and z_{N_q+1} are fixed as the minimum and maximum observed log TFP), defining the bin group indicator $q(z)$, and taking as given the weight on spacing, $\gamma \in [0,1]$, and observations z_i for $i \in \{1, \ldots, N\}$, the objective function is

$$\max_{\{z_j\}_{j=2}^{N_q}} \gamma \sum_{j=1}^{N_q} \log\left(\frac{z_{j+1}-z_j}{z_{N_q+1}-z_1}\right) + (1-\gamma)\sum_{j=1}^{N_q} \log\left(\frac{1}{N}\sum_{i=1}^{N} 1\{q(z_i)=j\}\right).$$

We pick $\gamma = 0.2$ in our benchmark estimation specification but test out other values to ensure γ does not affect the results.

To estimate our model primitives (transition matrix P, exit probability Q, and entry distribution λG for both old and new regimes), we leverage the observed firm productivity data and status (i.e., whether a firm is an incumbent, an entrant or an exiter). For the old regime, we assume the economy begins from a stationary distribution in 1998 with no exit or entry ($Q = 0$, $\lambda G = 0$), and infer P as the transition matrix that would maintain this distribution as stationary. The no exit assumption reasonably approximates China's pre-reform environment, where SOEs were not allowed to fail even though the 1986 Enterprise Bankruptcy Law created the legal possibility of SOE bankruptcy.[30] In order to ensure stationarity, we also assume no entry. For the new regime, we estimate P and Q using the non-parametric method with penalty $\kappa = 0.25$ calibrated to match the observed transition matrix among incumbents $\tilde{P}$ and the observed exit rate $\tilde{Q}$, and directly use the empirical distribution of entrants to estimate λG, assuming no selection in entry. This estimation approach allows us to recover all six primitive objects required for our decomposition while maintaining consistency with the key institutional features of China's economic transition.

7.2 Estimation Results

7.2.1 Transition Matrix Evolution

The productivity transition matrices reveal systematic increases in upward mobility and decreases in downward mobility over time. To facilitate a numerical measure of mobility, we define upward mobility as the average number of bins moved upward conditional on moving upward, $\sum_{i>j}(i-j)P(j,i)$, and similarly define downward mobility as $\sum_{i<j}(j-i)P(j,i)$. Figure 5 shows upward mobility rising from 6.6 in 1999 to around 7.5-8.0 in later years, while downward mobility declines from 7.5 in 1999 to approximately 5.0 by 2007.

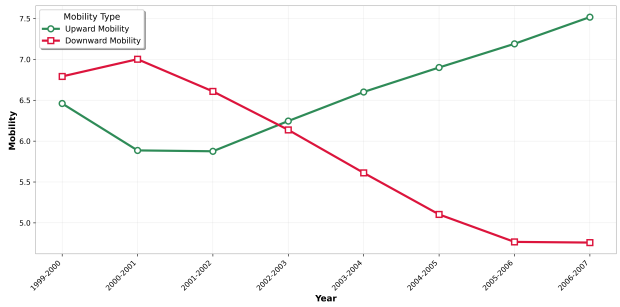
During the period from 1998 to 2000, the estimated transition matrices showed minimal productivity mobility, with firms largely maintaining their relative positions. By 2001-2003, low-productivity firms began showing an increase in the probability of improvement, while mid-productivity firms demonstrated greater upward mobility. The 2006-2007 period con-

[30] Hsieh and Song (2015) documents an annual exit rate of 0.9% between 1991 to 1995 using data from the 1996 China Statistical Yearbook and the microdata of the 1995 Industrial Census.

tinued this pattern, with enhanced dynamics across all productivity levels.

This evolution from static productivity rankings toward broad-based upward mobility provides strong empirical support for our hypothesis that the within-all effect strengthened over time. As firms became increasingly likely to move up rather than down the productivity distribution, within-firm productivity improvements became a more important driver of aggregate productivity growth.[31]

Figure 5: Firms Productivity Mobility Trends: 1998 - 2007



Note: We interpolate the productivity mobility trend in 2004 as missing key variables in 2004 requires imputation, potentially affecting estimated transition kernels for 2003-2004 and 2004-2005.

7.2.2 Exit Probability Patterns

In the early reform period (1998-2001), exit probabilities were relatively high across all productivity levels. Low-productivity SOEs ($\log \text{TFP} < 5$) faced exit probabilities of 45-58%, while middle-productivity firms ($\log \text{TFP} \approx 5 - 8$) experienced rates around 30-40%. Even relatively productive firms maintained substantial exit probabilities of 25-30%.

By the later reform period (2005-2007), exit probabilities had decreased significantly across the board. Low-productivity firms now faced exit rates of only 32-35%, while middle-productivity firms saw rates decline to 15-25%. High-productivity firms experienced the largest proportional decline, with exit rates falling to around 10-15%.

There is a consistent decrease in exit probabilities for firms at all productivity levels between 1998 and 2007. As a consequence of this general reduction in exit pressure, we would expect a diminishing exit effect over time.[32]

[31]Figures 3 and 4 in Appendix A.1 present nine one-year interval estimates of transition matrices and exit probabilities for SOEs.

[32]One persistent anomaly that appears across all years is the unexpected uptick in exit rates at very high

The consistent gap between observed and true exit rates across years (shown in Figure 2) highlights a key methodological insight: firms typically experience declining productivity before exiting, creating measurement challenges for conventional approaches. When correcting for this timing issue by conditioning on exit-time productivity rather than pre-exit productivity, the relationship between productivity and exit becomes substantially stronger than directly observed in the data.

7.3 Decomposition Results

Our benchmark decomposition uses the complete nine-year interval (1998-2007) to provide an initial assessment of productivity changes throughout the reform period. This approach is particularly illuminating as a starting point because pre-reform state-owned firms were not allowed to go bankrupt. By examining this full interval, we can effectively ask: How important were within-all, exit, and entry for overall TFP growth among SOEs during the reform period, compared to the counterfactual of no reform? Once we address the drivers of growth over the full transition, we will study how the importance of these drivers evolved along the transition by analyzing year-by-year changes.

7.3.1 Decomposing the Reform: 1998-2007

Figure 6 illustrates our decomposition for the full transition from 1998 to 2007, separating the total observed difference (black line) into three fundamental components: exit effect (green bars), within-all effect (red bars), and entry effect (blue bars).

The overall productivity distribution shifted rightward significantly, with negative density differences at lower productivity levels (log TFP $\approx$ 2-7) and positive differences at higher levels (log TFP $\approx$ 8-13).

The exit effect dominates this transformation, showing large positive contributions in the entire distribution. This is driven by the strength of the relationship between productivity and exit - the recovered exit probability among the least productive firms is more than 90%,

productivity levels. This pattern may indicate measurement error in extreme productivity estimates, other data issues, or some other factors that fall outside our analysis. Despite this anomaly, the overall pattern of lower exit probabilities for more productive firms holds throughout the sample period.

while the exit probability among the most productive firms is less than 50%. Moreover, the exit probability is monotonically decreasing with firm productivity (see Figure 5 in Appendix A.1). The exit effect played an important role in shifting the entire distribution rightward by removing less productive firms.

The within-all effect shows a mixed pattern. The effect works against the reduction of the share of firms at the lower tail (log TFP $\approx$ 2-4), but contributes positively to the rightward shift in the middle of the distribution (log TFP $\approx$ 4-8) and upper tail (log TFP $\approx$ 9-12). This suggests that firms experienced varying productivity changes, with some low-productivity firms deteriorating further while high-productivity firms improved.

The entry effect is negligible across most of the distribution. It contributes positively at the very high end (log TFP $\approx$ 10-12). This indicates that new entrants were not systematically more productive than incumbent firms, except at the highest productivity levels.

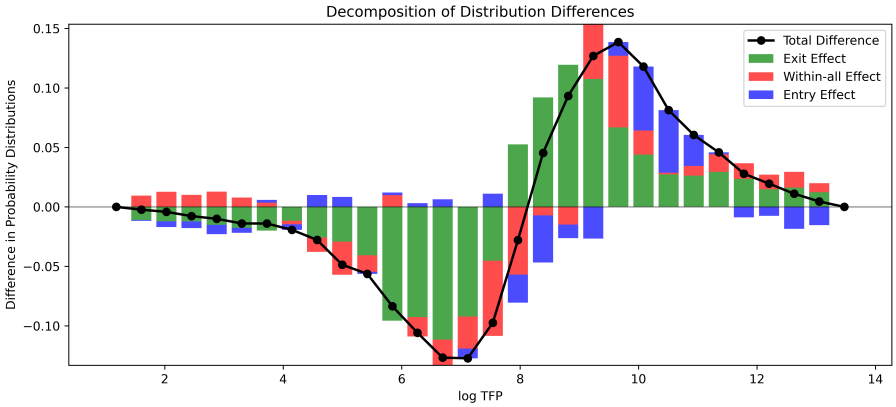
We also quantify the change in the unweighted mean of log TFP. We find that mean log TFP increased by 1.25 log points between 1998 and 2007, with the exit effect and within-all effect accounting for 84.7% and 14.7% of this improvement respectively. The entry effect contributed the remaining 0.06%.

An important methodological consideration is that using the unweighted mean of productivity may introduce bias since larger, more productive firms contribute disproportionately to aggregate economic outcomes. To address this concern, we conduct robustness checks using input shares to compute a weighted mean of log TFP. If we compute the decomposition using the weighted mean of log TFP, we find that the exit effect contributed 85.8%, the within-all effect contributed -17.8%, and the entry effect contributed 32.1% to overall growth. The change across the decompositions is consistent with the observation from Hsieh and Song (2015) that entrants are large in size, which would magnify the positive effect of entrants and explain why the contribution of the entry effect rises from the unweighted decomposition to the weighted decomposition.

One important caveat of these results is that these counterfactuals assume that firms can never exit. So for instance, firms that in reality chose to exit in 1999 would be forced to continue operating indefinitely in our counterfactual scenarios. We setup the counterfactuals in this way to match the pre-reform policy that SOEs were not allowed to go bankrupt.

These findings highlight both the dominant role of selection in driving SOE productivity improvements, while also acknowledging the non-negligible contribution of within-firm factors, particularly among higher-productivity firms. This demonstrates that while the productivity gains in the state sector stemmed primarily from eliminating inefficient firms, there were meaningful improvements among state firms.

Figure 6: Benchmark Decomposition of Changes in TFP Distributions: 1998-2007



7.3.2 Decomposing the Reform: Year-by-Year

Figure 7 presents our year-by-year decomposition of state-owned enterprise productivity drivers for the period 1998-2007. The figure breaks down how much of the growth in unweighted average TFP can be attributed to changes in the transition matrix (P), exit probability (Q), and entrant distribution (λG), compared to the pre-reform regime.

To compute the decomposition, we start from the pre-reform stationary distribution in 1998. We then construct counterfactual scenarios where only one primitive changes while the others remain fixed at pre-reform levels. The within-all effect (blue line) measures productivity gains from allowing only the transition kernel P to evolve post-reform according to our estimated annualized P, while keeping exit probabilities and entry patterns at 1998 levels. The exit effect (orange line) isolates the contribution of changing exit probabilities Q, holding transition dynamics and entry constant at pre-reform levels. The entry effect (green line) captures gains from allowing only the entrant productivity distribution λG to change, with all other parameters fixed.

For each effect, we calculate the difference between the mean productivity under the counterfactual scenario and what the mean would have been if all parameters had remained at 1998 levels. The bottom panel shows year-by-year marginal contributions, while the top panel displays cumulative effects over time.

The decomposition reveals a two-phase pattern of reform. During the early period (1998-2002), productivity gains were dominated by exit effects. The exit channel contributed a substantial 0.172, 0.126, and 0.115 log points in the first three years respectively, before gradually tapering to less than 0.1 log point annually by the end of the reform period. This rapid initial wave reflects the removal of the least productive state-owned enterprises, generating immediate gains through selection.

In contrast, the within-all effect was negative at approximately -0.07 log points at the beginning of the reform period, but underwent a dramatic transformation after 2001. From that point forward, within-all productivity improvements accelerated sharply, eventually becoming the dominant driver of average TFP growth. By 2006-2007, the within-all effect contributed 0.277 log points - nearly eight times the 0.035 log points from exit effects.

The entry effect maintained a modest but consistent role throughout the transition period, contributing between 0.03 and 0.06 log points annually. While never matching the magnitude of the other channels, entry provided steady background productivity enhancement.

These patterns illuminate a fundamental shift in reform strategy. Initially, productivity growth stemmed primarily from "trimming to the tree" - the selective removal of under-performing firms. Over time, however, the approach evolved toward "growing the tree" - enhancing the productivity dynamics of existing state enterprises. This transition is clearly reflected in the cumulative effects: by 2007, cumulative within-all and exit effects stood nearly equal (0.86 vs. 0.81 log points), despite their vastly different temporal trajectories.

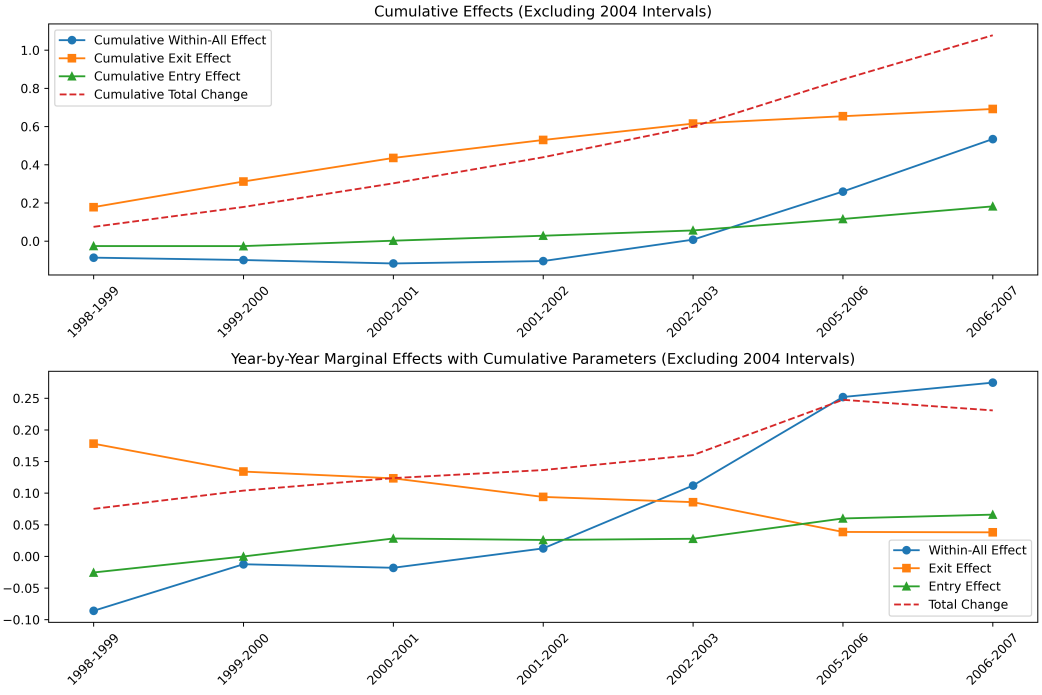
This evolution directly connects to our parameter estimates. The decline over time in exit probabilities across all productivity levels explain the diminishing exit effect, while the transformation from transition matrices with limited mobility for the period 1998-2002 to those enabling significant upward movement during 2005-2007 explains the surge in within-all effects (see Figure 5).

As robustness, we repeat the exercise using the input-share weighted mean of log TFP.

The weighted decomposition results, shown in Figure 6 in Appendix A.1, confirm that our main findings are robust to this alternative specification. The fundamental pattern remains unchanged: exit effects decline in importance over time while within-all effects become the dominant driver of TFP growth after 2005.

Finally, it is important to note that the "Total Change" line tracks the year-to-year change in observed mean log TFP. Since each individual effect (within-all, exit, entry) isolates only one changing parameter while holding others fixed, the gap between their sum and the total change captures interactions among all channels - highlighting that real-world productivity dynamics are shaped by intertwined mechanisms, not isolated forces.

Figure 7: Temporal Pattern of Effects in Driving Average TFP Changes during 1998-2007



Note: Missing key variables in 2004 required imputation, potentially affecting productivity measures for 2004-2005. We exclude the 2004 decomposition results in this figure.

7.4 Robustness Checks

We conduct an extensive set of checks by varying key aspects of the estimation specification. These variations include different penalty values for the estimators, alternative productivity measures beyond the Hsieh and Klenow (2009) specification, industry-specific

estimation to account for sectoral heterogeneity, adjustments to sample trimming thresholds, and tests with different bin configurations (varying both the number of bins and the weighting between equal spacing and equal observation counts). Most importantly, we compare results from our non-parametric estimation approach to those obtained using an AR(1)-style specification. Across all these exercises, our main findings remain essentially unchanged (see Appendix G for detailed estimation and decomposition results).

7.5 Discussion

The timing and nature of our decomposition results - where the within-all effect turned from negative prior to 2003 to positive after 2003 and accelerated sharply after 2005 - coincides with two major institutional reforms that fundamentally altered Chinese state-owned enterprises (SOEs). While we cannot establish direct causation, these reforms provide a compelling framework for understanding the productivity transformation we observe.

The State-owned Assets Supervision and Administration Commission (SASAC)'s establishment in 2003 consolidated fragmented state ownership that previously had scattered authority across multiple bureaucracies. This consolidation enabled operational changes that could plausibly drive productivity improvements in the state sector.

SASAC enabled SOEs to simultaneously consolidate for efficiency and expand for scale through three complementary mechanisms. First, SASAC facilitated consolidation for operational efficiency as central SOEs underwent strategic restructuring to eliminate redundancies and clarify business mandates.[33] From 2003 to 2007, the number of central SOEs under SASAC was reduced from 186 to 151, while the average number of subsidiaries in these central SOEs increased from 133.0 to 187.8 (Brandt et al. (2023)). According to SASAC's own data, return on assets of these firms increased steadily from 6% in 2005 to 6.7% in 2007 (Lardy (2019)).[34] Second, SASAC actively supported expansion for competitive scale by promoting

[33]For example, China Telecom underwent major restructuring during this period, separating its fixed-line and mobile telecommunications services into distinct operational units with clearer strategic mandates. This included geographic consolidation to eliminate redundancies and improve operational efficiency, part of the government's broader SOE reform agenda to create more competitive and professionally managed telecommunications companies.

[34]The return on assets has declined since the global financial crisis and stayed low at around 2.6% through 2017. However, this paper is focused on only the period from 1998 to 2007.

SOE growth through public listings and facilitating domestic and international mergers and acquisitions (Song et al. (2011a), Song (2018)).[35] These activities provided SOEs with enhanced capital access, advanced technology, management expertise, and economies of scale. As a result, many SOEs became among the largest corporations globally. For example, 67 of 69 Chinese companies that were present on the 2014 Fortune Global 500 list are state-owned (Hsieh and Song (2015)). Third, SASAC introduced governance reforms for market discipline through the mandatory establishment of boards of directors. Complementary to the SASAC reforms, China Securities Regulatory Commission (CSRC) Regulation No. 102 required independent directors, which brought external expertise and market-oriented decision-making into SOE operations (Bhabra and Li (2011)).

The Split-Share Structure Reform, carried out in 2005, was another major institutional reform during this period that addressed a fundamental structural problem in Chinese capital markets. When SOEs went public before 2005, they transferred only small ownership stakes to private investors while the government retained control through dominant non-tradable shares. This dual structure created an incentive misalignment: state shareholders holding two-thirds of shares received compensation based on book asset values rather than market values, systematically rewarding asset accumulation over productive efficiency. The structure also generated broader market distortions including ineffective boards, profit tunneling through related-party transactions, rampant stock market speculation, blocked merger activities, and stunted development of corporate bond and derivative markets (Liao et al. (2014)).

The Split-Share Structure Reform dismantled this dual system by converting non-tradable state shares into tradable shares, directly linking state shareholder wealth to market performance. This institutional change potentially created three operational shifts that could drive productivity improvements: SOE managers now faced market-based rather than asset-based performance metrics; stock price movements introduced market discipline that could penalize operational inefficiencies; and equity market access shifted investment decisions toward

[35]Key examples include: CITIC's 2006 acquisition of Kazakhstan's Karazhanbas Oil and Gas Field (340+ million barrels of proven reserves); Baosteel's 2007 acquisition of Xinjiang Ba Yi Steel to consolidate domestic steel capacity; and SAIC's acquisition of MG Rover technology (Rover 25, Rover 75, and K-series engines), enabling the launch of the Roewe brand by 2006.

market returns rather than political objectives.

The alignment of the timing between these reforms and the patterns uncovered by our decomposition suggest a plausible causal interpretation. The establishment of SASAC provided the operational capabilities (capital access, strategic focus, external expertise) that could enable productivity improvements, potentially explaining why the within-all effect turned positive after 2003. However, the fundamental incentive misalignment remained until 2005.

The Split-Share Structure Reform addressed this constraint by aligning managerial incentives with productivity rather than asset accumulation. This timing coincides precisely with the sharp acceleration in our within-all effect after 2005. By 2006-2007, when the within-all effect contributed 0.277 log points compared to only 0.035 from exit, SOE managers were operating under fundamentally different institutional constraints that systematically rewarded productivity improvements.

While alternative explanations cannot be ruled out,[36] the specific nature and timing of these reforms provides a coherent framework for understanding the striking shift from exit-driven to within-firm productivity growth that defines this period of state sector development.

7.6 Method Comparison: BHC-FHK vs. Selection-Corrected Decomposition

To highlight the importance of accounting for selection effects, we compare our results with the widely-used Baily et al. (1992) and Foster et al. (2001) (BHC-FHK) decomposition. Table 1 presents the BHC-FHK decomposition results for state-owned enterprises (SOEs) during 1998-2007.

Input-cost-weighted firm-level TFP growth during 1998-2007 was 1.51 log points. The BHC-FHK decomposition attributes 33.7% of this growth to within-firm improvement, 8.0% to between-firm reallocation, 12.4% to the cross effect, 33.9% to entry, and 12.0% to exit.

[36]For example, productivity growth among state firms could be driven by the greater competition from private firms (Hsieh and Song (2015)).

The temporal patterns show declining exit effects over time: from 54.4% (1999-2001) to 24.2% (2001-2003), then stabilizing around 31% in later periods. Within-firm effects follow the opposite pattern, increasing over time from -9.9% (1999-2001) to 77.2% (2005-2007), turning from a negative to a strongly positive contribution.

Both our decomposition and the BHC-FHK approach reveal similar temporal patterns. Exit effects decline over the reform decade, while within-firm effects increase over time and transition from negative to positive contributions during 2001 and 2003.

The two methods disagree on the primary drivers of aggregate productivity growth over the full 1998-2007 period. BHC-FHK attributes productivity growth primarily to within-firm improvements (33.7%) and entry effects (33.9%), with exit playing a minor role (12.0%). Our selection-corrected decomposition reveals the opposite: exit effects dominate (84.7%), while within-firm effects contribute modestly (14.7%).

Table 1: BHC-FHK Decomposition Results

Interval	ΔTFP	within (%)	between (%)	cross (%)	entry (%)	exit (%)
1998 - 2007	1.51	33.71	7.99	12.43	33.91	11.96
1999 - 2001	0.20	-9.86	18.03	25.71	11.70	54.42
2001 - 2003	0.54	53.75	21.29	-5.65	6.39	24.23
2003 - 2005	0.29	48.06	48.71	-3.31	-24.94	31.47
2005 - 2007	0.38	77.17	11.59	-4.68	-15.64	31.56

This dramatic discrepancy stems from how each method treats firm productivity trajectories and exit decisions. BHC-FHK measures exiting firms' productivity only in their final observed year, capturing a snapshot rather than the full decline process. When a firm exits after experiencing substantial productivity deterioration, BHC-FHK sees only the firm's low (compared with survivors), final-year productivity and calculates a small productivity gain from its removal.

Our selection-corrected approach captures the counterfactual productivity path - what would have happened to these firms' productivities had they continued operating. Many exiting firms were on steep decline trajectories. By removing these declining firms, the state sector avoided substantial productivity losses. BHC-FHK misses this dynamic entirely.

The methods also differ fundamentally in what constitutes "within-firm improvement."

BHC-FHK measures within-firm effects for only surviving firms - how much did firms that remained active improve their productivity? Our approach measures within-firm effects for all firms, including those that exit, by using a selection correction to estimate what exiting firms' productivity trajectory would have been had they continued operating.

This conceptual difference has important implications. If exiting firms were experiencing productivity decline before exit (as our selection correction reveals), then including them in the within-firm calculation naturally reduces the measured within-firm effect compared to looking at only survivors. BHC-FHK's larger within-firm effects (33.7% vs. 14.7%) partly reflect this definitional difference.

The fundamental issue is that BHC-FHK treats firm exit as exogenous to productivity dynamics, while our approach recognizes that exit decisions are possibly endogenous responses to productivity decline. This endogeneity creates systematic measurement bias that leads their approach to understate selection effects and overstate other sources of productivity growth. Using our approach instead fundamentally alters our understanding of Chinese SOE reform dynamics.

8 Conclusion

Our paper provides an accounting of how different components governing the underlying process of the evolution of the productivity distribution of China's state-owned enterprises (SOEs) evolved during the period from 1998 to 2007. Our results reveal that exit effects dominated throughout most of the reform period (1998-2005) before giving way to within-firm improvements. While we cannot establish causal links between specific institutional changes and these patterns, our findings raise important questions for future research. The dominant role of exit effects suggests that allowing firms to fail under "let go of the small," hardening budget constraints, and more market competition from private firms may have been critical in enabling market selection, but determining which specific policy changes triggered this massive firm exit requires further investigation. Equally puzzling is the source of the positive within-all effects that emerged after 2003: did productivity improvements stem from SOE consolidation and restructuring activities under SASAC supervision, from better corporate

governance, from better managerial incentives following the split-share structure reform, from technology spillovers as trade barriers fell, or from some combination of these factors? While our statistical decomposition cannot disentangle these mechanisms, in this paper we have uncovered that China's industrial transformation involved two distinct dynamics - first a dramatic selection process that eliminated most underperforming firms ("trimming the tree"), then genuine productivity improvements among state firms ("growing the tree"). Understanding why this particular sequence occurred and whether it represents a necessary ordering for successful industrial transitions remains an open question, one with profound implications for other economies contemplating SOE reform. Our conceptual framework provides future researchers with tools to examine these questions across different contexts and time periods. It also helps policymakers distinguish whether productivity gains stem from better performance or better selection.

For future work, we propose extending our selection model to address the double-selection issue of state-owned firms: selecting into privatization and selecting to remain in the market. This extension would allow us to evaluate whether accounting for both selection processes would alter our conclusions about the effectiveness of SOE reforms. Additionally, incorporating alternative performance measures beyond productivity - such as profitability, innovation output, and return on capital - would provide a more comprehensive assessment of how institutional changes affect firm performance and economic outcomes. These extensions would further illuminate the dynamics of state capitalism and offer insights into the long-term sustainability of China's economic growth model.

References

Acemoglu, Daron, Philippe Aghion, and Fabrizio Zilibotti, "Distance to Frontier, Selection, and Economic Growth," *Journal of the European Economic Association*, 2006, *4* (1), 37–74.

Backus, Matthew, "Why is Productivity Correlated with Competition?," *Econometrica*, 2020, *88* (6), 2415–2444.

Bai, Chong-En, Jiangyong Lu, and Zhigang Tao, "Property Rights Protection and Access to Bank Loans: Evidence from Private Enterprises in China," *Economics of transition*, 2006, *14* (4), 611–628.

Baily, Martin Neil, Charles Hulten, and David Campbell, "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity: Microeconomics*, 1992.

Baqaee, David Rezza and Emmanuel Farhi, "Productivity and Misallocation in General Equilibrium," *The Quarterly Journal of Economics*, 2020, *135* (1), 105–163.

Basu, Susanto and John G Fernald, "Aggregate Productivity and Aggregate Technology," *European Economic Review*, 2002, *46* (6), 963–991.

Berkowitz, Daniel, Hong Ma, and Shuichiro Nishioka, "Recasting the Iron Rice Bowl: The Reform of China's State-owned Enterprises," *Review of Economics and Statistics*, 2017, *99* (4), 735–747.

Bhabra, Harjeet S and Tiemei Li, "Independent Directors and Corporate Performance: Evidence from Listed Firms in China," *Corporate Ownership and Control*, 2011, *8* (3), 46.

Brandt, Loren and Xiaodong Zhu, "Soft Budget Constraint and Inflation Cycles: A Positive Model of the Macro-Dynamics in China during Transition," *Journal of Development Economics*, 2001, *64* (2), 437–457.

_ , **Gueorgui Kambourov, and Kjetil Storesletten**, "Barriers to Entry and Regional Economic Growth in China," *Review of Economic Studies*, Forthcoming.

_ , **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), 339–351.

_ , _ , **Luhang Wang, and Yifan Zhang**, "WTO Accession and Performance of Chinese Manufacturing Firms," *American Economic Review*, 2017, *107* (9), 2784–2820.

_ , **John Litwack, Elitza Mileva, Luhang Wang, Yifan Zhang, and Luan Zhao**, "Recent Productivity Trends in China: Evidence from Macro- and Firm-Level Data," *China: An International Journal*, 2022, *20* (1), 93–113.

_ , **Ruochen Dai, and Xiaobo Zhang**, "The Anatomy of SOE in China," Technical Report, Working Paper, University of Toronto 2023.

Chan, Kam Wing, "The Household Registration System and Migrant Labor in China: Notes on a Debate," *Population and development review*, 2010, *36* (2), 357–364.

Chen, Yuyu, Mitsuru Igami, Masayuki Sawada, and Mo Xiao, "Privatization and Productivity in China," *The RAND Journal of Economics*, 2021, *52* (4), 884–916.

Chow, Gregory C, *China's Economic Transformation*, John Wiley & Sons, 2015.

Foster, Lucia, John C Haltiwanger, and Cornell John Krizan, "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in "New Developments in Productivity Analysis," University of Chicago Press, 2001, pp. 303–372.

Gollop, Frank M, Barbara M Fraumeni, and Dale W. Jorgenson, *Productivity and US Economic Growth*, Harvard University Press, 1987.

Griliches, Zvi and Haim Regev, "Firm Productivity in Israeli Industry 1979–1988," *Journal of Econometrics*, 1995, *65* (1), 175–203.

Groetsch, C, "The Theory of Tikhonov Regularization for Fredholm Equations," *Boston Pitman Publication*, 1984, p. 104.

Groves, Theodore, Yongmiao Hong, John McMillan, and Barry Naughton, "Autonomy and Incentives in Chinese State Enterprises," *The Quarterly Journal of Economics*, 1994, *109* (1), 183–209.

Hopenhayn, Hugo A, "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica: Journal of the Econometric Society*, 1992, pp. 1127–1150.

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

_ **and Zheng Michael Song**, "Grasp the Large, Let Go of the Small: The Transformation of the State Sector in China," Technical Report, National Bureau of Economic Research 2015.

Jefferson, Gary H, Thomas G Rawski, and Yuxin Zheng, "Chinese Industrial Productivity: Trends, Measurement Issues, and Recent Developments," *Journal of Comparative Economics*, 1996, *23* (2), 146–180.

Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica: Journal of the Econometric Society*, 1982, pp. 649–670.

Khandelwal, Amit K, Peter K Schott, and Shang-Jin Wei, "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese exporters," *American Economic Review*, 2013, *103* (6), 2169–2195.

Kreinin, Alexander and Marina Sidelnikova, "Regularization Algorithms for Transition Matrices," *Algo Research Quarterly*, 2001, *4* (1/2), 23–40.

Lardy, Nicholas R, *Markets over Mao: The Rise of Private Business in China*, Columbia University Press, 2014.

_ , *The State Strikes Back: The End of Economic Reform in China?*, Peterson Institute for International Economics, 2019.

Lau, Lawrence J, Yingyi Qian, and Gerard Roland, "Reform without Losers: An Interpretation of China's Dual-Track Approach to Transition," *Journal of Political Economy*, 2000, *108* (1), 120–143.

Li, Wei, "The Impact of Economic Reform on the Performance of Chinese State Enterprises, 1980–1989," *Journal of Political Economy*, 1997, *105* (5), 1080–1106.

Liao, Li, Bibo Liu, and Hao Wang, "China's Secondary Privatization: Perspectives from the Split-Share Structure Reform," *Journal of Financial Economics*, 2014, *113* (3), 500–518.

Liu, Ernest, "Industrial Policies in Production Networks," *The Quarterly Journal of Economics*, 2019, *134* (4), 1883–1948.

Liu, Zheng, Pengfei Wang, and Zhiwei Xu, "Interest Rate Liberalization and Capital Misallocations," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 373–419.

Melitz, Marc J and Sašo Polanec, "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit," *The RAND Journal of Economics*, 2015, *46* (2), 362–375.

Naughton, Barry, "Claiming Profit for the State: SASAC and the Capital Management Budget," *China Leadership Monitor*, 2006, *18*, 1–9.

Olley, G. Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 1996, *64* (6), 1263–1297.

Petrin, Amil and James Levinsohn, "Measuring Aggregate Productivity Growth using Plant-Level Data," *The RAND Journal of Economics*, 2012, *43* (4), 705–725.

Riedel, James, Jian Gao, and Jing Jin, "How China Grows: Investment, Finance, and Reform," 2020.

Rubens, Michael, "Market Structure, Oligopsony Power, and Productivity," *American Economic Review*, 2023, *113* (9), 2382–2410.

Song, Ligang, "State-Owned Enterprise Reform in China: Past, Present and Prospects," *China's 40 years of Reform and Development*, 2018, *345.*

_ , **Jidong Yang, and Yongsheng Zhang**, "State-Owned Enterprises' Outward Investment and the Structural Reform in China," *China & World Economy*, 2011, *19* (4), 38–53.

Song, Zheng, Kjetil Storesletten, and Fabrizio Zilibotti, "Growing Like China," *American Economic Review*, 2011, *101* (1), 196–233.

Tombe, Trevor and Xiaodong Zhu, "Trade, Migration, and Productivity: A Quantitative Analysis of China," *American Economic Review*, 2019, *109* (5), 1843–1872.

Wang, Ruiting, Xue Wang, Gang Xu, and Tao Zha, "Privatization's Impacts on State-Owned Enterprises: A Tale of Zombie versus Healthy Firms," Technical Report, National Bureau of Economic Research 2024.

Yang, Weidong, "The New Round of Reform of State-owned Enterprises," *Journal of Entrepreneurship in Emerging Economies*, 2015, *7* (1), 55–66.