Making a Growth Miracle

Historical Persistence and the Dynamics of Development

[Click here for most recent version]

Job Market Paper

Jonas Gathen[*], Oscar Fentanes[‡]

November 25, 2023

Abstract

What explains growth miracles? We argue that growth miracles are driven by a fundamental race: as the economy tries to catch-up to its steady state, changes in the economic environment move the steady state itself and provide new potential for catch-up growth. We quantify this race over the course of development using 40 years of plant-level manufacturing panel data from Indonesia and a structural model of plant dynamics. We estimate the model on the micro data along the observed growth path without assuming that the economy is ever at a steady state. While catch-up growth starting from initial conditions in 1975 accounts for 42% of Indonesia's subsequent industrialization, new changes in the economy induce new catch-up growth. In the end, the economy never catches up.

[*]Toulouse School of Economics (TSE), Email: jonas.gathen@tse-fr.eu. Corresponding author.

[†]Toulouse School of Economics (TSE), Email: oscar.fentanes@tse-fr.eu

[‡]We are grateful for the continued support and many useful comments by our advisors: Fabrice Collard, Matteo Bobba, Christian Hellwig and Stéphane Straub. We further thank in no particular order: Ben Moll, Yongseok Shin, Francisco Buera, Ezra Oberfield, Richard Rogerson, Jihyun Kim, Eugenia Gonzalez-Aguado, Tuuli Vanhapelto, Christophe Bruneel-Zupanc, Miguel Zerecero, Alexandre Gaillard, Andreas Schaab, Gyung-Mo Kim and all participants of the TSE Macro workshop. We want to thank the World Bank Jakarta team, especially Massimiliano Calí, for gratefully sharing the manufacturing plant-level data for Indonesia. We further thank feedback during presentations at Princeton, WashU St Louis, EWMES 2022, the Annual STEG conference 2022, EEA/ESEM 2022, ENTER jamboree 2021 as well as the PhD and Macro workshops at TSE and LSE. All remaining errors are our own.

1 Introduction

Over the past 50 years, extended periods of rapid economic growth in China, India and Indonesia alone lifted roughly 1 billion people out of extreme poverty (World Bank 2023). What drives such growth miracles? A common view is that growth miracles capture a transition process: a permanent policy regime change sets an initially poor and highly misallocated economy on a transition path towards a better long-run equilibrium (Buera and Shin 2013; Asturias et al. 2023). These transitions take time because frictions prevent labor and capital to quickly reallocate across firms and sectors. But if transitions take time, how does this view account for new changes in policies while the economy is still transitioning? And how can we identify the aggregate effects of new policies if the economy is still adjusting to previous policies?

In this paper, we show evidence that instead of a one-time transition, growth miracles are driven by a never-ending race: as the economy tries to catch-up to its steady state, changes in the economic environment move the steady state itself and induce new transition growth. We do so by looking at Indonesia, the fourth most populous country in the world. Indonesia provides an ideal case to study this race since we can draw on almost half a century of manufacturing plant-level panel data during which the Indonesian economy completely transformed: GDP per worker increased five-fold, the working population tripled, output in manufacturing grew 30-fold and the manufacturing employment share doubled.

Drawing on the micro data, we provide empirical evidence that motivate a model of plant dynamics in which growth is driven both by changes in the economic environment and transition phases. Empirically distinguishing these two drivers of growth is crucial because one might otherwise falsely attribute current growth to current policy changes. To do so, we estimate the model on the micro data along the observed growth path without assuming that the economy is at a steady state at any point in time. Intuitively, observed conditional choices of plants identify changes in the economic environment, while the distribution over these choices summarizes the past and reveals the potential for transition growth.

In line with the common view of growth miracles, we find that initial transition growth is important: letting the economy in 1975 transition while shutting down all future changes in the economic environment explains 42% of the manufacturing growth between 1975 and 2015. However, we also find that the Indonesian economy in 2015 is not closer to its steady state than it was in 1975, precisely because new changes in the economic environment in the meantime moved the steady state itself. As the economy is always undergoing important transition processes, one key implication is that evaluating policies without considering these adjustments is highly misleading.

We now provide further details and results. Drawing on our data, we document four main facts that help us disentangle transition growth from changes in the economic environment and motivate our subsequent model:

Fact 1: Rapid economic growth coincided with changes in the plant distribution. Average plant size doubled, the mass of plants increased four-fold and the right tail of the plant size distribution thickened.

Fact 2: Adjustment processes can account for changes in the plant distribution. These are respectively: An aging of the plant distribution together with the fact that plants enter small and grow over their life cycle, slow entry and exit dynamics, and the fact that it takes time to grow large plants.

Fact 3: The drivers of aggregate productivity differed markedly before and after the Asian Financial Crisis. Before the crisis, productivity across plants was driven entirely by the selection of more productive plants, informed by meager within-plant productivity growth. In contrast, the post-crisis growth period was characterized by strong within-plant productivity growth and little growth due to selection.

Fact 4: The allocation of resources did not improve systematically over time. This is robust to different measures of misallocation. We find evidence for volatile productivity dynamics at the plant-level, large changes in entry, and plant-level adjustment frictions – particularly in labor – from an event study design that can jointly account for this.

To quantify the race between transition growth and changes in the economic environment, we draw on a model of plant dynamics in the tradition of Hopenhayn (1992) that is motivated by the previous empirical facts. In the model, firms face risk regarding their productivity, choose to enter and exit and hire labor and capital subject to adjustment frictions that lead to the slow accumulation and reallocation of resources across sectors and firms. We embed these plant dynamics into a two-sector economy to capture the endogenous reallocation of workers across manufacturing and the rest of the economy. The main frictions in manufacturing are labor adjustment costs, in particular convex costs, that prevent plants from growing a large workforce quickly (*Facts 2 + 4*). Plants also endogenously enter and exit based on drawing entry costs and fixed costs of production. The level and dispersion of costs in turn rationalize the observed speed of entry and exit dynamics (*Facts 2*). These features imply that with an initial distribution characterized by few but productive young plants, the economy goes through a process of transition growth as plants gradually grow, more plants gradually enter over time and unproductive plants gradually exit. At any given point in time, the model economy is characterized by a set of exogenous model fundamentals and the state of the current economy as captured by the distribution of plants. Model fundamentals include all cost parameters as well as time-varying aggregates such as aggregate labor supply and technological changes in manufacturing (*Fact 3*) and the rest-of-the-economy. Policy affects growth through driving part of the changes in model fundamentals. Technically, we assume that plants make dynamic choices forming rational expectations over their future idiosyncratic risk but have perfect foresight over future aggregate changes in the economy. This introduces a computationally difficult fixed point problem: plants' dynamic choices depend on expectations over the future path of market-clearing prices, which in turn depend on the endogenous evolution of the entire distribution of plants (as in Buera and Shin 2013).

The key methodological contribution of the paper is to propose a tractable estimation strategy that allows to estimate this model economy on standard plant-level micro data along the growth path in the data without assuming that the observed economy is at a steady state at any point in time. Importantly, our model estimation allows model fundamentals to vary flexibly over time, making it particularly suited to study fast-changing economies and markets. The model estimation proceeds in three main steps that allow to distinguish transition growth from changes in fundamentals and make the computational costs of the estimation independent of the computational costs of solving for a path of model equilibria. In the first step, we identify the path of time-varying equilibrium prices – only wages in our case – along the observed growth path (e.g. as in: Gopinath et al. 2017). Given that our model can account for this equilibrium path, we can treat the path of wages as fixed throughout the estimation and thereby avoid solving for the computationally costly fixed point in the path of equilibria. In the second estimation step, we identify the distribution of plants over the state space of the model, summarizing the history of the economy. In this step, we estimate plant production functions (Ackerberg, Caves, and Frazer 2015; Demirer 2020) and propose a novel strategy to separate plant-level productivity into an idiosyncratic and a common aggregate technology component, allowing us to distinguish selection-driven from technology-driven productivity growth. In the third and last estimation step, we estimate the model parameters that govern plants' adjustment frictions. We do so by drawing on Euler equation Continuous Conditional Choice (CCC) estimation, exploiting observed conditional input and exit choices of plants and avoiding to solve a dynamic programming problem to compute model-based dynamic input choices (Hotz and Miller 1993; Bajari, Benkard, and Levin 2007). In this last step, we estimate sizable convex adjustment costs in our model, which are identified from the empirical pattern that even small but highly productive plants grow their labor force gradually over time.

Using the estimated model, we find that transition growth from starting the economy with initial conditions in 1975 and shutting down all future changes in model fundamentals explains 42% of subsequent manufacturing output growth and all of the aggregate welfare increases that are due to changes in manufacturing over time. Given an initial distribution that features young and small plants, sizable labor adjustment frictions and slow entry and exit dynamics, it takes the economy 26 years to reach 90% of the steady state manufacturing output. Importantly, transition growth remains an important driver of growth precisely because the economy's fundamentals continue to change. To quantify this point, we repeat the previous exercise to compute the transition path for each year, starting from each year's initial distribution and model fundamentals. We find that the economy does not get closer to its (time-varying) steady state. It takes on average 20 years to reach 90% of the steady state manufacturing output and – if anything – the time it takes increases over time. Based on our results, we can thus strongly reject the idea that transition growth is a transitory phenomenon.

Large changes in fundamentals are key to explain the continuing importance of transition growth. The structural model allows us to quantify the role of changes in fundamentals in Indonesia's manufacturing growth miracle and quantify how much of their effect can be explained by changes in observed government policy. We do so by focusing on two important changes in fundamentals that can be linked to development policies that the Indonesian government also pursued to varying degrees over the 40 years we study: (1) large-scale investments in education that raise the pool of skilled (and cheap) labor, and (2) the active use of FDI policy to attract manufacturing plants under foreign ownership.

We find that the manufacturing growth miracle would not have happened in the absence of the estimated doubling in human capital per worker, because labor would have been more expensive in this economy and manufacturing plants are far more sensitive to higher wages than the rest of the economy. To gauge the importance of policy in driving overall human capital increases, we then evaluate Indonesia's largest school construction program (INPRES) through the lens of the model. Building on micro-empirical evidence on the wage effects (Duflo 2001, 2004), the scale of the program (Akresh, Halim, and Kleemans 2023), and the slow labor market integration of treated cohorts, we show that by 2015, the program accounts for roughly 10% of the overall manufacturing output growth that is due to human capital per worker increases in the economy.

In contrast, for FDI, we find that manufacturing output in 2015 would have only been 8% lower in the complete absence of foreign-owned entrants, while we find that regulatory changes in FDI policy in the late 1980s may potentially account for 85% of the overall effect of FDI on manufacturing. Taken together, this still means that most growth by far stems from structural forces related to demographics. Thus, a somber conclusion based on these results – partly resonating related work on the Indian growth miracle (Bollard, Klenow, and Sharma 2013) – is that policy matters less for growth than we might think.

Related literature

We contribute to four main strands of literature. First and most importantly, we complement the growth and macro development literature by studying a model of firm dynamics where growth is driven by the combination of changes in exogenous fundamentals and transition growth. This is in contrast to a firm dynamics literature that has mostly analyzed development differences through differences in steady states.[1] Much fewer papers study transition growth with firm dynamics (e.g. Buera and Shin 2013; Moll 2014; Akcigit, Alp, and Peters 2021; Ruggieri 2022; Asturias et al. 2023; Lanteri, Medina, and Tan 2023). We add to this literature by (1) allowing for further changes in fundamentals along the transition and by (2) estimating this model on the micro data along the transition without assuming that the economy is at a steady state at any point in time. Quantitatively, we find that the combination of both sources of growth matters. Not only are initial transition growth and changes in fundamentals important for growth, but the economy is always far away from its (time-varying) steady state, questioning the usefulness of either comparing steady states or focussing on transitions in the absence of further changes in fundamentals. Apart from the methodological differences, Buera and Shin (2013) and Asturias et al. (2023) are the most closely related in their focus on understanding growth miracles. Our results mainly differ from Buera and Shin (2013) in that we find no role for reductions in frictions and misallocation – their main driver of transition growth – but rather a key role for plant and worker demographics in driving transition growth. We return to Asturias et al. (2023) below.

Second, both modeling and key results in this paper relate to the recent quantitative spatial, trade and migration literature, which focus on frictional worker mobility and trade while abstracting from firm dynamics. For example, the idea that transitions take long and the economy is persistently far away from its steady state resonates with recent findings from Allen and Donaldson (2020) and Kleinman, Liu, and Redding (2023). This literature relates to and builds on the seminal work of Caliendo, Dvorkin, and Parro (2019) who also study the combination of changes in exogenous fundamentals and transition growth. In contrast to Caliendo, Dvorkin, and Parro (2019), tractability in our case does not come from dynamic hat algebra techniques and we estimate all time-varying model fundamentals. Slow transitions in our paper and this literature share common causes: low and highly dispersed exit (moving) probabilities and slow input adjustments.

Third, we contribute more generally to the Quantitative Macroeconomics literature by showing how to tractably estimate the model directly on the observed transition path in the data, identify the model entirely on plant-level data and only use macro moments for model validation, an approach that we see as closely aligned to a growing literature that moves "from micro to macro" (see the overview in: Buera, Kaboski, and Townsend 2023). The "equilibrium estimation" methods that ensure tractable estimation – enforcing the observed path of equilibria throughout the estimation and Euler CCC estimation – are used in other literatures, but have not yet seen wider application in the Macroeconomic literature.[2] We

[1] While the literature on firm dynamics with a development focus is too vast to cite, overviews are for example given by Hopenhayn (2014) and Restuccia and Rogerson (2017) for misallocation, Ulyssea (2020) for informality and Alessandria, Arkolakis, and Ruhl (2021) for trade.

[2] The idea of estimating models on the observed path of equilibria in Macroeconomics dates at least back to Hansen and Singleton (1982). More recent research that conditions estimation on the observed path of equilibrium prices can for example be found in the literature following Caliendo, Dvorkin, and Parro (2019).

find the equilibrium estimation approach to be particularly suited for studying a path of time-varying equilibria, since we can also tractably estimate entire paths of time-varying parameters. While estimation methods that require to first solve the model may offer more flexibility on the choice of moments that identify parameters, they often have to strongly restrict the parameter space.

At last, our paper also relates to the literature on growth and productivity dynamics. Our results of selection-driven aggregate productivity growth in Indonesian manufacturing mirrors similar results in Brandt, Van Biesebroeck, and Zhang (2012) for China and Asturias et al. (2023) for Chile and Korea. We add to these papers by establishing the result of selection-driven productivity growth in a non-parametric setting that nests a larger class of growth models including various endogenous growth models.

The rest of the paper is structured as follows. The next section presents the main empirical evidence. Section 3 develops the model and discusses identification, estimation and model validation. In Section 4, we quantify the main drivers of growth. The last section concludes.

2 Empirical evidence

In this section, we introduce the Indonesian data and key facts about the Indonesian growth experience that motivate the subsequent model.

2.1 Data

Our primary data comes from the plant-level Annual Manufacturing Survey, collected by Indonesia's Central Bureau of Statistics. It covers only medium- to large-sized manufacturing plants by surveying all formal manufacturing establishments with more than 20 employees. The survey contains detailed and consistent annual information on standard plant-level characteristics from 1975 to 2015, a period of 41 years. It covers between roughly 7,500 to 30,000 plants per year. Throughout, we draw on reported information on plants' age (based on birth year), industry (up to 5-digit) and ownership (including foreign ownership). On the production side, we draw on plants' capital stock, value-added revenue, and the number of workers (including paid and unpaid workers) and wage bill (including contributions and in-kind compensation) which are separately reported for production and non-production workers. Unfortunately, capital is only reported starting in 1990. All variables denoted in Indonesian Rupiah are deflated to real values using the aggregate CPI and normalized to

In a recent working paper, Humlum (2022) exploits these two estimation steps in a similar general model framework of growth and firm dynamics, although entirely different context of industrial robot adoption in Denmark.

the year 2010. For aggregate data, we further combine the GGDC 10-sector and Economic Transformation databases for Indonesia, which capture time-consistent aggregate sectoral employment and output series over the time period 1960-2012 (Timmer, de Vries, and De Vries 2015) and 1990-2018 (Kruse et al. 2023) respectively. We refer to this data as the *GGDC* series throughout. In Appendix A.1, we discuss in detail the data cleaning and homogenization steps we take to ensure consistency of all datasets over time.

There are two important data limitations. First, while the data is in principle a census of plants with 20 or more workers, in practice the census misses plants. This shows up in discontinuous jumps in new plant entry during years of the economy-wide Economic Census in 1975, 1985, 1995 and 2006.In those years, plants are added that were previously missed, either because they were small and made the cutoff, or newly entered. For the subsequent analyses, this means that aggregate changes often show discontinuous drops in census years and should actually look more smoothed out. Given that we observe the census in 1975, the initial distribution is correctly reported. Furthermore, this does not bias results that are based on within-plant variation. Our data also misses plants because of non-reporting, either because plants miss to report in some years or because we are forced to drop a plant-year entry due to misreporting (see Appendix A.1). We treat these missing entries as missing at random and specifically account for missing entries in our structural model. We correct our measurements of plant entry and exit by denoting plant entry as the first time when a plant identifier enters the panel and plant exit if we do not observe a given plant identifier at any future time period (see: Appendix A.1.6).

The second main data limitation is with respect to the coverage of our data. As we show in Appendix A.2, while our dataset misses the approximately 99% of Indonesian manufacturing plants that have less than 20 workers, most of these plants are characterized by self-employment with a modal plant size of one to two workers. After cleaning, our manufacturing data captures between 25-30% of total manufacturing employment and value-added output as based on the *GGDC* data, with shares increasing over time (Figure A.1). We think of small scale manufacturing as a separate sector given robust evidence that there are few transitions between small and larger scale manufacturing (e.g. Poschke 2013; Van Biesebroeck 2005; Schoar 2010) and most plants that enter our panel have only recently been established. For example, the median age of plants that newly enter our plant panel is only two years. The focus in this paper is thus on how relatively large plants and their dynamics drive aggregate economic growth. In the model and results parts, we explicitly model the entire economy, taking into account that our data only captures a time-varying share of output in the overall economy.

Figure 1: Evolution of aggregate and sectoral employment and output

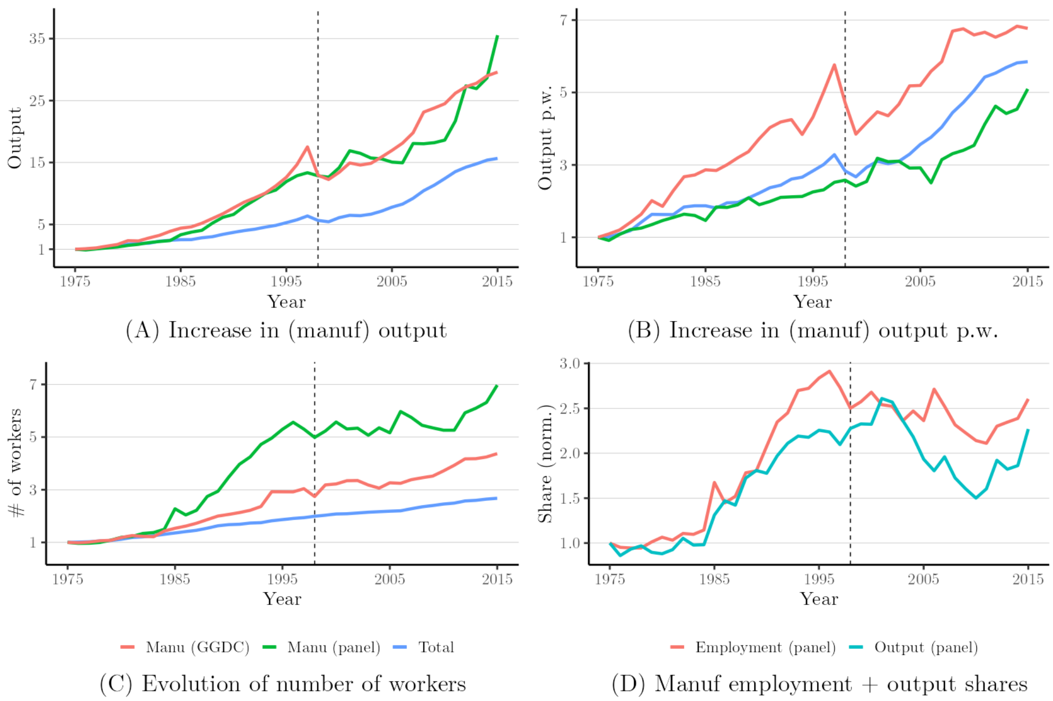

(A) Increase in (manuf) output

(B) Increase in (manuf) output p.w.

(C) Evolution of number of workers

(D) Manuf employment + output shares

Notes: (Economy-wide) Total gives the aggregate of the GGDC data. Panel refers to the Indonesian manufacturing plant census (1975-2015, 20+ workers). All series are normalized by their respective value in the first year. (A) and (B) use value-added output. Dashed vertical lines denote the onset of the Asian Financial Crisis in 1997.

2.2 Four key facts of growth

We now highlight four key facts that shed light on the Indonesian growth experience. The first two facts relate to changes in the entire plant distribution and the importance of slow adjustment processes. Fact 3 and 4 document changes in the drivers of productivity growth and the absence of improvements in the allocation of resources over time.

2.2.1 Rapid economic growth coincided with changes in the plant distribution

As the 4th most populous country in the world, Indonesia underwent a dramatic process of economic development over the past 50 years, a few key features of which are reported in Figure 1. Between 1975-2015, GDP per worker increased more than five-fold (Panel B), driven by a 15-fold increase in output (Panel A) and roughly a tripling of the working population (Panel C). Manufacturing contributed importantly to this aggregate growth: output grew 30-fold and the manufacturing employment share more than doubled (Panel D).

The period 1975-2015 can be divided into two main growth regimes that are separated by the Asian Financial Crisis in 1997. The pre-crisis period captures a period of rapid labor-intensive

Figure 2: Evolution of average plant size and number of plants

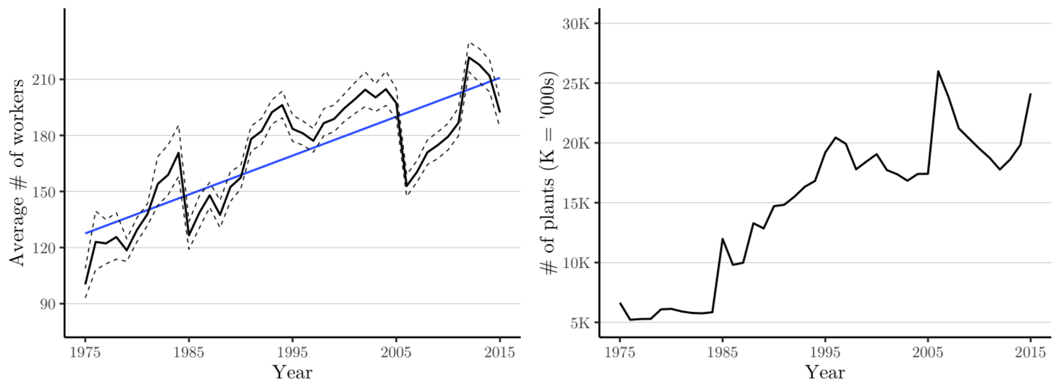

(A) Avg number of workers over time

(B) Number of plants over time

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. Panel A: Workers include paid and unpaid workers. Dotted lines give 95% bootstrapped confidence intervals and solid blue line gives best linear fit. Panel B: Jumps in 1985, 1995 & 2006 are explained by Economic Census years.

industrialization, including the period 1987-1994 that Hausmann, Pritchett, and Rodrik (2005) characterize as a growth acceleration. Most of the total worker flows into manufacturing happen before 1997 and manufacturing grows far more rapidly than the rest of the economy. Fast growth in the aggregate working population is also a defining feature of the pre-crisis period with an average annual growth rate of 3%, 70% higher than in the post-crisis period. Based on our census of medium- and large-sized plants, the rise of manufacturing rapidly takes off in the first half of the 1980s and industrialization peaks with the Asian Financial Crisis as evidenced by manufacturing employment and output shares (Panel D). After the Asian Financial Crisis, which started in 1997 and hit manufacturing mostly in 1998, the economy experienced lower total output growth that was due to substantially lower growth in plant entry and employment and – as we will show further below – by higher plant-level productivity growth.

These aggregate changes went in hand with systematic changes among manufacturing plants. As evidenced in Figure 2, the rapid increase in the total number of workers in manufacturing is met with a doubling of the average number of workers in manufacturing plants (Panel A) and a four-fold increase in the number of manufacturing plants. In both cases, most of the gains were already reached by 1997. Importantly, the entire plant distribution changed systematically over time. Specifically, Figure 3 shows that the right tail of the plant employment distribution systematically thickened over time – a key feature of the development process that has been highlighted for firms (rather than establishments) across and within countries (Chen 2022; Choi et al. 2023; Poschke 2018). Panel B shows that the employment share in plants with more than 50 workers increased by roughly 15%, while the employment share in plants with more than 500 workers increased by more than 50%.[3] The increase in the right tail of the

[3]We report this metric as it is a simple transformation of the Pareto tail, which is also robust to left-

Figure 3: Evolution of employment distribution over time

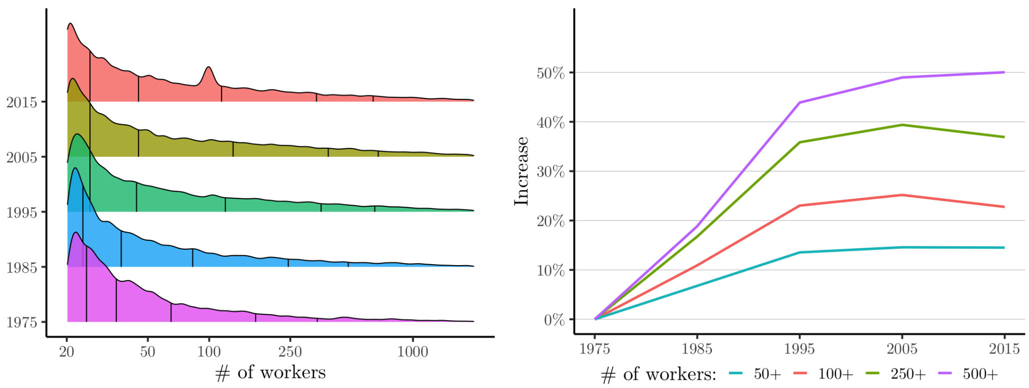

(A) Size distribution over time (B) Employment shares over time (1975 = 1)

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. Only showing years 1975, 1985, 1995, 2005 and 2015. Panel A: Vertical lines give 25th, 50th, 75th, 90th and 95th percentiles. Bunching at 99 workers starts in 2013, after the passing of the 2012 Worker Safety Law that binds for establishments with 100+ workers. Panel B: Share of employment in plants with more than X workers.

employment distribution is a main driver of the increase in the average plant size over time.

2.2.2 Adjustment processes can account for changes in the plant distribution

We now show evidence for three slow adjustment processes that can respectively account for increases in the plant size, the mass of plants and the right tail of the plant size distribution.

The aging of the plant distribution We start by showing that the slow aging of the plant distribution – average plant age increased by 40% since 1975 – can explain average plant size increases. The reason is that plants enter small and grow over their life cycle. Figure 4 plots life-cycle growth profiles across different cohorts of surviving manufacturing plants. Plants enter roughly with a similar average number of workers, which grows with plant age. Plants that survive at least 20 years have about twice as many workers as new entrants; in comparison and as documented in Hsieh and Klenow (2014), manufacturing plants in the US that survive for that long are about six times as large as new entrants.[4]

A benefit of our panel data is that we can show that the increase in the average size of surviving plants is mostly driven by within-plant growth rather than selection (larger plants being more likely to survive). Figure 5 Panel A shows average year-to-year within-plant

censured data. We report the corresponding secular decline in Pareto tail coefficients in Appendix A.2. We find systematic changes in the Pareto coefficient, both in the cross-section (which is not in line with a common Pareto distribution) and over time (which is not in line with traveling wave equilibria).

[4]The numbers are not perfectly comparable, because of the cutoff of 20 workers in the Indonesian data. If young plants in the US data are smaller, this overestimates the difference across countries, while plants that stay below 20 workers in the US data bias downward.

Figure 4: Plant life cycle growth by birth cohort

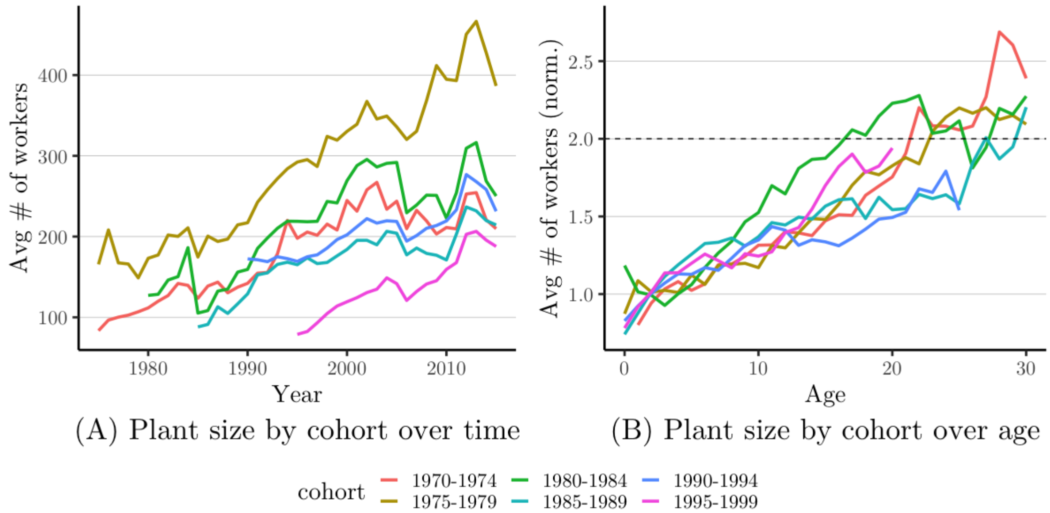

(A) Plant size by cohort over time (B) Plant size by cohort over age

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. Workers include paid and unpaid workers. Plant age is based on reported year of establishment. Panel B normalizes each entry by the cohort-specific average plant size of plants below age 5 (as in Hsieh & Klenow 2014). Note that each cohort over time is an unbalanced panel as only surviving plants stay in the panel and there is (limited) plant entry as plants make the cutoff of 20 workers to be included in the census.

growth by age. Young plants grow their employment quickly, with growth declining slowly as plants become older, average growth running out after around 20 years. Panel B reports the relative contribution of survivor growth to average plant size increases by age using the following accounting identity:

$$\underbrace{\bar{L}_a - \bar{L}_{a-1}}_{\Delta\text{avg plant size}} \equiv \underbrace{\frac{N_a^S}{N_a}(\bar{L}_a^S - \bar{L}_{a-1}^S)}_{\text{Survivor contribution}} + \underbrace{\frac{N_a^E}{N_a}\bar{L}_a^E}_{\text{Entry}} - \underbrace{\frac{N_{a-1}^X}{N_{a-1}}\bar{L}_{a-1}^X}_{\text{Exit}} + \underbrace{\bar{L}_{a-1}^S\left(\frac{1}{N_a} - \frac{1}{N_{a-1}}\right)}_{\text{Net reallocation}} \quad (1)$$

where S refers to the set of surviving plants from age $a-1$ to age a, E refers to entering plants (which exist because of the size threshold in our data) and X refers to exiting plants. The contribution of survivors, entry and exit respectively measure their average size weighted by their share in the population of plants. The net reallocation effect is driven by changes in the total number of plants over age: if exit outweighs entry (as in our case) then workers are reallocated towards fewer plants, mechanically increasing average plant size. We find that for young plants, growth by survivors explains all of the increase in average employment across plants, while selection as given by the remainder dominates the total effect after age 10-15.

Together, this evidence implies that average plant size crucially depends on where the age distribution of plants is; since the initial distribution of plants in the newly emerging manufacturing sector in 1975 featured mostly young plants, average plant size was small. Despite the rapid entry of new plants, plants became older and hence larger over time. For

Figure 5: Within-plant (survivor) growth in employment

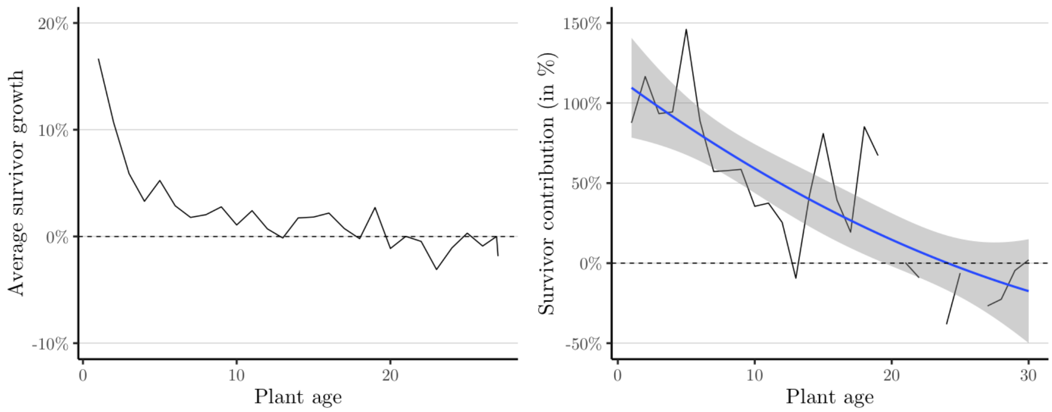

(A) Avg employment growth by age (B) Relative contribution of survivors

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. Sample restricted to survivors and plants for which plant age is defined (N = 543,586). Panel A: Average annual growth in total number of workers (paid + unpaid) across plants by age, weighted by a plant's previous employment. Panel B: Relative contribution of survivor growth to average plant size growth based on accounting identity in Equation 1 and dividing by the left-hand side to obtain relative contributions.

example, average plant age increased by 40% between 1975 to 2006 (the last year in which plant age was reported in the survey). Figure A.5 also shows how the entire age distribution shifted right over time.

It takes time for entry and exit dynamics to play out Next, we highlight a basic driver of slow entry and exit dynamics that can partially account for the observed four-fold increase in the mass of plants: low entry and exit rates. The idea is simple: the young manufacturing sector in 1975 features few plants and if only few potential entrepreneurs move into and out of entrepreneurship, it takes time to build up a mass of plants. For example, suppose there is a fixed mass of potential entrepreneurs and entry and exit rates into entrepreneurship are the same. Then the long-run (steady state) share of entrepreneurs is 1/2, but if the economy starts with no entrepreneurs, it can take many years to get close to the steady state. For entry and exit rates at 7.9% – equal to the average exit rate across all years and plants in our data[5] – it already takes the economy 14 years to reach just 90% of the long-run steady state. While we do not generally observe potential entrants and thus cannot study entry rates without further assumptions, in the Appendix we show additional evidence for slow exit processes. Specifically, Figure A.8 shows that exit rates only slowly

[5] This exit rate is substantially lower than the 14-18% documented for informal establishments in Vietnam (McCaig and Pavcnik 2021) and slightly smaller than the 8.3% documented for small establishments across 12 developing countries (McKenzie and Paffhausen 2019). Within manufacturing, exit rates also seem to decline with plant size, explaining a lower exit rate of 6.2% across all US manufacturing (e.g. Clementi and Palazzo 2016) and broadly similar exit rates for all manufacturing plants in India and Mexico (Hsieh and Klenow 2014).

Figure 6: Reduced-form transition dynamics implied by initial plant size distribution

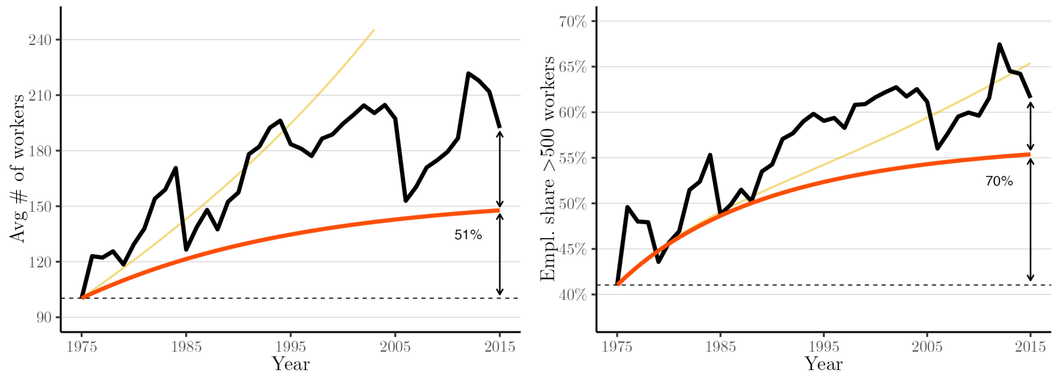

(A) Avg size increase

(B) Increase in right tail

Notes: Predicted changes in distribution based on exercise taking discretized initial plant size distribution (# of workers) in 1975 and transition matrix giving conditional probabilities of moving from one plant size bin to another for years 1975-1976. Predictions iterate on initial discretized distribution with fixed transition matrix. Predicted (constant growth) instead enforces transition matrix incorrectly enforcing constant growth taking average plant size growth for 1975-1976.

decrease with plant productivity and that exit rates do not increase with aggregate shocks such as the Asian Financial Crisis. This is in line with growing evidence that stagnant firms in developing countries tend to survive longer compared to firms in developed countries (e.g. Hsieh and Klenow 2014; Akcigit, Alp, and Peters 2021; Eslava, Haltiwanger, and Pinzon 2022).

It takes time to grow large plants At last, we show that a lack of large plants in the 1970s and the fact that it takes time to grow large plants can jointly explain the slow fattening of the right tail of the employment distribution. We do so by considering the following exercise. Take as the starting point the initial employment distribution of plants Φ_0 in 1975 and discretized in $X = 10$ different size bins. Each bin captures the fraction of plants with this number of workers. We now follow individual plants and compute the probability of moving from one bin to the other between 1975 and 1976, which we summarize in the transition matrix P of dimension X^2. We predict changes in the distribution by iterating on the initial distribution using the fixed transition matrix: $\hat{\Phi}_{t+\tau} = \Phi_0 \cdot P^\tau$. Figure 6 shows that the exercise explains 50% of average plant size increases (A) and 70% of the increases in the employment share of plants with more than 500 workers (B) over time. The reason is that in 1975, the distribution lacks large plants in comparison to the stationary distribution implied by the employment growth observed between 1975 and 1976 and it takes time to grow large plants. The exercise predicts that it takes 25 years to reach 90% of the steady state average plant size, broadly capturing the speed at which plant growth plays out over time.

The exercise is robust to a number of concerns which we address in Appendix A.3.[6] Importantly, the exercise cannot distinguish the drivers of employment growth. The transition matrix $P_{t,t+1}$ only gives a reduced-form summary of plant employment growth subject to any frictions and growth drivers that are present between time t and $t+1$, which may include adjustment frictions, changes in wages or productivity growth. The next two subsections thus focus on two key determinants of within-plant employment growth: productivity growth and plant-level hiring frictions.

2.2.3 Productivity growth and selection

What is the role of productivity growth in the Indonesian growth experience and how much productivity growth is explained by the better selection of plants? In this section, we show that aggregate productivity increased roughly 3.5-fold between 1975 to 2015. However, the underlying drivers of this productivity growth differ fundamentally across Indonesia's two main growth periods: during the period of rapid labor-intensive industrialization (1975-1997), all of the aggregate productivity gains are driven by the selection of more productive plants, while aggregate productivity gains during the period after the Asian Financial Crisis are almost entirely driven by within-plant productivity gains. We obtain these results by standard production function estimation (Ackerberg, Caves, and Frazer 2015; Demirer 2020) and then separately identifying the productivity improvements that come from the better selection of plants versus within-plant productivity improvements.

Estimating productivity Following the literature, we estimate a standard value-added Cobb-Douglas production function in capital k and efficiency units of labor h:

$$y_{it} = x_{it} h_{it}^{\theta_{jt}} k_{it}^{\alpha_{jt}} \tag{2}$$

with $\theta_{jt} + \alpha_{jt} \leq 1$ giving the output elasticities of labor and capital in sector j at time t and x_{it} capturing plant-level productivity. As a baseline and in line with our subsequent dynamic model, we start with common output elasticities across manufacturing industries, but discuss further industry variation below. We estimate the labor and capital output elasticities separately for each year allowing for both inputs to be potentially fully dynamically chosen, which means that at this point we can remain agnostic about the frictions that determine

[6] Specifically, we show similar results when allowing for entry and exit, taking any other starting years in the 1970s and that the slow filling up of the right tail of the employment distribution holds for any other year-to-year transition matrix that can be constructed between 1975 and 2015. An important caveat that we cannot address is that due to the discretization, the exercise would also give an increase in the plant size if all plants were simply growing at a constant rate. The red lines in Figure 6 report changes with a counterfactual transition matrix enforcing average employment growth rates between 1975 to 1976 across all plants. While observed within-plant employment growth is far from constant, the alternative exercise still shows that it is easy to overestimate the growth in the right tail.

plant input choices and changes in these frictions over time. Specifically, we draw on the control function approach in Demirer (2020), which is close in spirit to the standard control function value-added production function estimation based on Ackerberg, Caves, and Frazer (2015), but does not require intermediate inputs (also see: Gandhi, Navarro, and Rivers 2017). We provide a discussion and an identification proof adapted to our setting in Appendix A.5, but the intuition of identification is as follows: exploiting the assumption that productivity follows a first-order Markov process, conditional on previous input choices and output, the ranking of current inputs identifies the ranking of productivity innovations, which can be used to construct a control function for (unobserved) productivity in the output regression.

Results and more details on the estimation are shown in Appendix A.5. Before we discuss the results, it is important to highlight that we estimate the production function on available revenue data. This comes with the standard limitations that we only identify revenue elasticities and revenue-based productivity and cannot distinguish between productivity and demand nor identify changes in markups.[7]

We find no systematic changes in estimated output elasticities over time and very standard values for the output elasticity of labor close to 2/3. Importantly, the estimated output elasticity of labor is substantially larger than the median plant-level labor share (≈ 0.45) and the aggregate labor share in manufacturing (≈ 0.25). In a frictionless model with Cobb-Douglas production, plants would equalize cost shares and revenue elasticities. In the next section, we provide evidence on frictions for labor choices that could rationalize this difference. Our subsequent model then quantitatively accounts for these large differences. We find smaller estimates for the capital output elasticity than generally found in the literature, which – as we discuss in Appendix A.5 – is likely due to attenuation bias from measurement error in observed capital. A lower capital output elasticity means that observed (mismeasured) variation in capital has smaller effects on output and we further show that mismeasurement of the capital elasticity is not biasing our estimates for labor.

Selection versus plant-level productivity growth Next, we quantify how much of the productivity improvements across plants are driven by the selection of more productive plants versus within-plant productivity growth. For this, we assume that plant-level productivity is

[7]Using separate information on prices and quantities, the previous literature has highlighted the important role of demand for driving firm growth (Hottman, Redding, and Weinstein 2016; Foster, Haltiwanger, and Syverson 2016; Eslava and Haltiwanger 2020) and an important part of what we subsequently call "productivity" likely captures demand. We return to this issue when discussing model counterfactuals where the distinction between demand and productivity is key. Also, revenue-based productivity measures may be preferred in the Indonesian context where large changes in product quality bias quantity-based productivity estimates (see Atkin, Khandelwal, and Osman (2019) for the argument and Hill (2000) for a discussion of strong quality improvements in Indonesian manufacturing). Relatedly, we also do not identify changes in mark-ups, which usually requires separate information on prices and quantities (Bond et al. 2021). Studying changes in markups and its drivers over the course of development is an exciting direction of future research, but beyond the scope of this paper.

the product of a common, potentially endogenous, aggregate technology component z_t that improves the productivity of all plants and an idiosyncratic productivity shock s_{it}: $x_{it} \equiv z_t s_{it}$. This setup allows us to separate shared technology growth in z_t, selection on idiosyncratic productivity s_{it} and within-plant growth in s_{it}, and nests the productivity side of many exogenous as well as endogenous growth models in the literature.[8] In this setup, we provide a novel non-parametric identification approach that separates the path of z_t from s_{it} (up to a normalization of z_0).

To understand why separate identification of selection and technology growth is difficult in the first place, let us start by looking at log changes in average productivity over time:

$$\frac{1}{N_t}\sum_{i\in\mathcal{N}_t}\tilde{x}_{it} - \frac{1}{N_{t-1}}\sum_{i\in\mathcal{N}_{t-1}}\tilde{x}_{it-1} = \underbrace{\tilde{z}_t - \tilde{z}_{t-1}}_{\Delta log(z)} + \underbrace{\frac{1}{N^S_{t,t-1}}\sum_{i\in\mathcal{N}^S_{t,t-1}}\Delta\tilde{s}_{it}}_{\text{Survivor }\Delta log(s)} + \underbrace{\frac{1}{N^E_t}\sum_{i\in\mathcal{N}^E_t}\tilde{s}_{it}}_{\text{Entry }\overline{log(s)}} - \underbrace{\frac{1}{N^X_t}\sum_{i\in\mathcal{N}^X_t}\tilde{s}_{it-1}}_{\text{Exit }\overline{log(s)}}$$

where $log(x) \equiv \tilde{x}$ denotes variables in logs. Changes in average productivity $\tilde{x}_{it}$ only identify growth in technology z under the special case that average changes in productivity $\tilde{s}$ among survivors as well as changes in average productivity $\tilde{s}$ between exiting and entering plants exactly cancel out. These terms capture two important selection biases. The entry and exit terms capture a "static" compositional selection bias that leads to overestimates of aggregate technology changes as long as less productive plants are more likely to exit and entering plants are positively selected on productivity. Focusing instead on within-plant changes in productivity deals with the "static" selection bias, but still leaves a term capturing average changes in idiosyncratic productivity among surviving plants. We call this term the "dynamic" selection bias, which generally biases estimated aggregate technology changes downwards if productivity s_{it} is persistent. Intuitively, if surviving plants are selected based on good histories of productivity realizations s_{it}, then they are more likely to mean revert in the future.

Non-parametric identification in our setup means that we make no functional form assumptions on the arbitrarily time-varying path of aggregate technology z_t, the productivity shock process s_{it}, nor on the plant entry and exit processes that drive endogenous selection. For expositional purposes, we provide an idea of the identification and estimation approach and relegate the precise technical assumptions, a detailed identification proof and further estimation details to Appendix A.6. Identification of changes in z_t relies on two sets of assumptions. The first assumption ensures that the productivity shock process s_{it} has a stationary distribution. Technically, we assume that s_{it} follows the same underlying general first-order, ergodic Markov

[8] Specifically, the setup nests standard neoclassical growth models that feature exogenous aggregate productivity growth and firm selection (e.g. Luttmer 2007; Clementi and Palazzo 2016) as well as endogenous growth models that feature a common endogenous growth component such as Romer (1990) or models where Gibrat's law holds and productivity growth is independent of firm size (Klette and Kortum 2004; Atkeson and Burstein 2010; Restuccia and Bento 2015; Peters 2020).

Figure 7: Aggregate manufacturing technology estimates

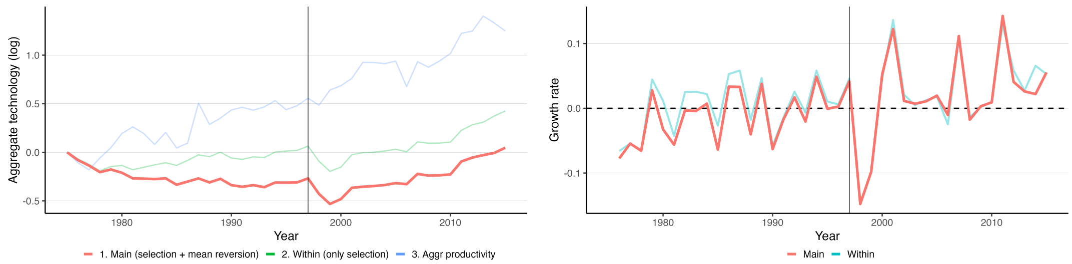

Notes: Panel A: Aggregate technology estimates, showing the main estimator (explained in the data) as well as the within estimator (that only controls for selection, but not mean reversion) and aggregate productivity (measured as value-added-weighted average productivity). Panel B: Corresponding growth rates in aggregate technology. The main and within estimators both use (weighted) median changes in plant-level productivity. Further details in the text.

process across plants and time, allowing for flexible forms of error dependence. The stationary distribution is useful because if we could reweight changes in observed productivity x_{it} among the – potentially highly selected – set of surviving plants based on the stationary distribution of s_{it}, the "dynamic" selection bias exactly cancels out. That is, $\mathbb{E}_i \Delta \tilde{s}_{it}$ is exactly equal to zero at the stationary distribution of s. The second set of assumptions ensures that such a stationary distribution can always be constructed using the observed data, restricting the degree of selection at exit. Specifically, we require that (1) plants' exit decisions are not based on future productivity shock realizations, and (2) there is no sharp productivity cutoff at which all plants would exit, so that there is always common support that allows an appropriate reweighting of the distribution. The latter can be empirically tested and – as shown in Figure A.8 Panel B – finds strong support in the Indonesian data.

With the assumptions in hand, the only remaining difficulty is how to construct the weights of the stationary distribution and solve for the time path of changes in z_t. Here, we first solve "forward" for the stationary distribution by starting with equal weights over the initial distribution. Whenever a plant selectively exits, we pass on their weight to plants with similar productivity who did not exit using a standard Kernel estimator, creating a synthetic panel among surviving plants that is "representative" of the underlying process of s. For time growing large – no matter how selected the initial distribution is – one can thereby identify appropriate weights over the selected set of producing plants that recovers the stationary distribution of s (up to a common scaler z). We then move "backwards" from the last period T to identify the path of z_t: initially normalizing z_T, we start with a guess over z_{T-1}, compute the weighted changes in productivity between $T-1$ and T based on the stationary distribution of s in $T-1$ and solve for the implied z_T. This implies finding a root in z_{T-1}. Iterating on this procedure until z_0 recovers the entire path.

Figure 7 shows the estimated path of technology using the full sample and baseline productivity estimates. Over the entire 40-year period, technology improves little, being less than 5%

higher in 2015 than in 1975. However, this masks important changes over time. Specifically, technology actually declined strongly throughout the 1970s, remained almost constant throughout the 1980s and 1990s and then saw rapid growth of roughly 4% per year since the year 2000. At the same time, aggregate productivity – measured as the value-added-weighted average productivity across plants – increased roughly 3.5-fold by 2015 and increased by 75% by the time of the Asian Financial Crisis in 1997. Together, this means that the sources of productivity growth fundamentally differed over the two Indonesian growth periods: plant selection rather than shared technology growth drove more than all of the productivity gains during Indonesia's period of rapid labor-intensive industrialization (1975-1997), while the pattern reversed after the year 2000, with common technology growth explaining more than 90% of the aggregate productivity gains. These estimates – especially after the Asian Financial Crisis – are driven by within-plant productivity growth of surviving plants. Pre-1997, true technology growth is lower than the unweighted within-plant productivity growth because of positive mean reversion that can be explained by the presence of many young and low productive plants that mean revert upwards in their productivity shocks. Post-1997, positive and negative mean reversion roughly balance out, explaining the similar growth paths of the main and within estimators. In Appendix A.6 we further discuss what could be driving these large changes in the role of aggregate technology over time.

2.2.4 The allocation of resources did not improve over time

At last, we look at measures of the misallocation of resources over time. We start with an accounting-based decomposition of growth in Indonesian manufacturing. Using the previously assumed production structure and separation of plant-level TFP into an aggregate technology and idiosyncratic productivity component, we can write growth in manufacturing output as:[9]

$$\Delta\, ln\, Y_t \equiv \underbrace{\Delta\, ln \sum_i f(h_{it}, k_{it})}_{\text{input growth}} + \underbrace{\Delta\, ln\, z_t}_{\text{aggr techn.}} + \underbrace{\Delta\, ln \left[\bar{s}_t + N_t cov\left(s_{it}, \frac{f(h_{it}, k_{it})}{\sum_i f(h_{it}, k_{it})}\right)\right]}_{\text{selection + reallocation effect}} \quad (3)$$

where N_t tracks the number of active plants. Total output is the combination of the state of factor accumulation and aggregate TFP. Aggregate TFP, in turn, can be further decomposed into aggregate technology and – following Olley and Pakes (1996) – a combination of average productivity and a covariance term that captures whether resources in the economy are allocated towards the most productive plants. Since the covariance is affected by common trends in its inputs such as changes in the sample size, Figure 8) plots the correlation of plant-level productivity and resource shares, which is robust to common trends and simply the normalized covariance. Panel A and B show this measure of the allocation of resources for the

[9] The proof can be found in Appendix B.4.

Figure 8: Evolution of cross-sectional correlation of plant productivity and input share

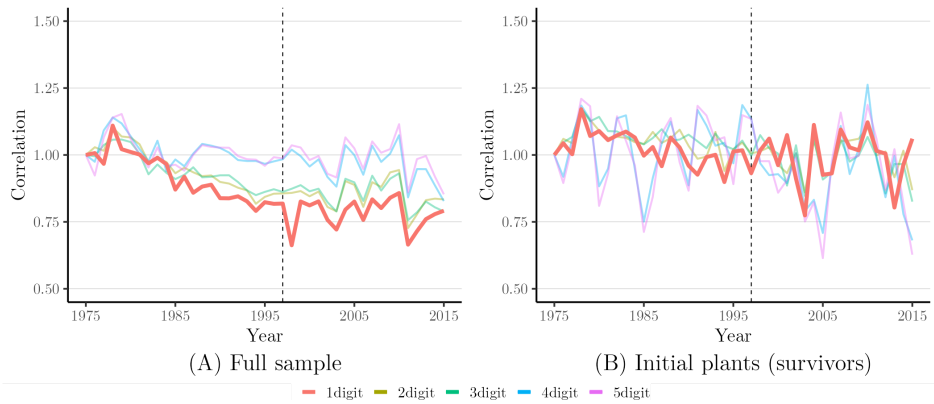

Notes: Input shares are computed based on Cobb-Douglas aggregator. Within-industry results based on first estimating the correlation across plants in a given industry and year and then constructing the weighted average correlation across industries using the industry's average share in manufacturing value added as a time-invariant weight. All series are normalized by the first year.

full sample of plants and only for surviving plants that already operated in 1975, normalizing each by the first year. Each panel additionally plots the cross-sectional correlation within industries.[10] Across all plants, the allocation of resources actually deteriorated over time, while it remained stable for survivors, indicating that the deterioration is driven by the entry of small plants. As further documented in Appendix A.7, this result also holds within a balanced panel and separately within each cohort of entering plants between 1970 to 1999.

In Appendix A.7, we also report changes in the dispersion of marginal revenue products of capital and labor over time, which maps to changes in misallocation in the literature building on Hsieh and Klenow (2009).[11] Again, we find that, if anything, the dispersion of measured marginal revenue products (even within 5-digit industries) tends to increase over time. Thus, based on our estimates and in contrast to Buera and Shin (2013), we find little evidence for an undoing of misallocation being a feature of growth in Indonesia.

The misallocation dynamics shown in Figure 8) are a function of changes in plants' productivities and how inputs reallocate across plants. To shed light on these dynamics at the micro-level, we are interested in how plants' inputs respond to changes in productivity. Figure 9 shows how plants' labor and capital shares evolve as a response to a positive and permanent

[10] Specifically, we proceed similar as in Gopinath et al. (2017): we first estimate the correlation across plants in a given industry and year and then construct the weighted average correlation across industries using the industry's average share in manufacturing value added as an industry-specific time-invariant weight. Using the same weights when aggregating across industries ensures that within-industry estimates reflect purely variation within industries over time.

[11] There are important limitations to interpreting changes in the covariance between measured revenue productivity and input shares as changes in misallocation. For example, in Hsieh and Klenow (2009), an efficient allocation implies zero correlation between TFPR and the input share. We note that revenue productivity in our model is not equal to TFPR in Hsieh and Klenow (2009), but instead much more closely correlated with their measure of TFPQ.

Figure 9: Event study: Input share responses to a permanent productivity shock

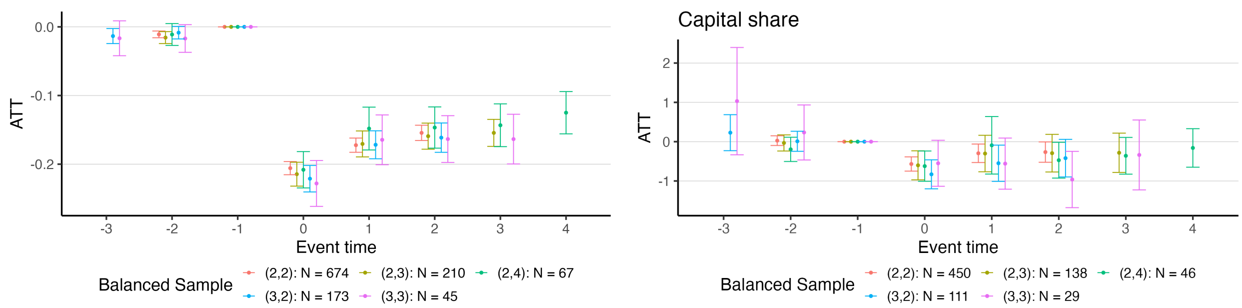

Notes: Treatment defined by permanent productivity shock of 20 percent change that does not revert back, neither before nor after treatment. Event study following Callaway & Sant'Anna (2021) estimated on balanced panel of treated plants and with all non-treated plants as control, event time zero gives first period of treatment, effects at -1 normalized to zero by assumption. Sample for capital share is restricted to post 1990 due to data availability. Details in the text.

productivity shock of at least 20% – roughly equal to the 75th percentile of within-plant productivity changes. We use a standard staggered differences-in-differences design (Callaway and Sant'Anna 2021) that is particularly suited here, because it captures plant-level dynamic responses while controlling for time and plant fixed effects that ensure that results are neither driven by aggregate shocks in specific years nor by fixed differences across plants such as differences across industries or even plant-specific production functions. Given selective plant entry and exit and resulting composition biases from estimating treatment effects on unbalanced panels, Figure 9 reports estimated dynamic treatment effects for different balanced panels.[12]

Plants' labor shares drop by more than 20 percentage points in the first year of treatment and recover slowly over time as plants respond by increasing hiring. In Appendix A.8, we further show that the slow labor share adjustments are indeed driven by slow hiring rather than wage increases. In a model with the above production side but where plants can adjust employment without constraints, there should be no response in the labor share, indicating frictions in adjusting labor. Quantitatively, these frictions are large: For example, the average "treatment" shock is roughly 40%, so that a plant with an initial labor share of 0.7 would not adjust labor at all in the first year of a large productivity increase. In contrast, estimated dynamic responses for capital are much noisier and we cannot reject that capital shares

[12]The treatment effects are estimated with all non-treated plants as control. Event time zero gives the first period of treatment, while treatment effects in event time -1 are normalized to zero by assumption. Apart from the 20% cutoff, the treatment definition also ensures we identify a permanent productivity shock by ruling out shocks followed and preceded by productivity changes of more than 10% per year and ruling out an accumulation of shocks over multiple periods that undo the shock at period 0. E.g. if we look at an event window from -3 to 2, then productivity at -3 cannot be more than 10% apart from the pre-level at -1 and productivity at 2 cannot be more than 10% apart from productivity at event 0. This does not restrict input shares. The rapid observed declines in the sample sizes are due to the high observed volatility of productivity, making it difficult to identify permanent shocks in the data.

immediately recover after the first year. For both inputs, we find no evidence for pre-trends, pointing away from an anticipation of productivity shocks. Taken together, volatile plant-level productivity and slow labor adjustments thus offer an explanation as to why the allocation of resources did not improve over time.

3 Structural model

While the four empirical facts document important changes in the economy and distribution of manufacturing plants over time (Fact 1), entry and exit dynamics that shape the selection of plants (Facts 2, 3 & 4), and the importance of slow plant-level adjustment processes (Facts 2 & 4), the empirical evidence is not sufficient to quantify the drivers of the Indonesian growth miracle. In particular, it does not allow us to quantify the aggregate effects of policy changes separately from transition growth. For this, we now build a model of plant dynamics and growth in the tradition of Hopenhayn (1992). In the model, plants face idiosyncratic risk in their productivity and choose capital and labor inputs subject to labor adjustment costs and a simple financing constraint that rationalize slow plant-level adjustments (Facts 2 & 4). Plants face fixed costs that drive endogenous entry and exit, which in turn drives aggregate selection dynamics (Fact 3).

The model features a time-varying growth path which is driven by three endogenous forces: changes in the input distribution, changes in the productivity distribution due to a combination of exogenous technology growth and plant selection, and changes in (mis)allocation as given by their joint distribution. All three forces are driven by the race between transition growth and by changes in model fundamentals that induce new transition growth. Changes in model fundamentals include changes in labor supply, potential entrants, technology, adjustment frictions and taxes. We further embed the model of plant heterogeneity into a two-sector general equilibrium model to capture changes in the rest of the economy and the endogenous reallocation of labor across sectors over time. The potential of transition growth at any point in time is given by the current distribution of plants encoding the history of the economy and future growth potential as given by current model fundamentals and expectations over the future. We follow the growth literature in treating the growth path as deterministic and as agents having perfect foresight over aggregate changes in the economy.

3.1 Model Setup

The model economy is set in discrete time indexed by $t = 1, 2, ...$ We assume that Indonesia is a small open economy vis-a-vis the rest of the world and has access to world capital markets at interest rate r^*. There are two sectors of production: Manufacturing (M) and the rest-of-the-economy (R). Both sectors produce the same homogeneous, perfectly substitutable good, which serves as numeraire. Manufacturing features heterogeneous plants whose endogenous mass and distribution are time-varying, while we model the rest-of-the-economy as a simple representative firm whose exogenous technology and endogenous labor demand change over time. Labor is inelastically supplied by households which choose in which sector to work. Labor markets in both sectors are fully competitive. There is a government that levies a valued added and a corporate income tax, the two main corporate tax instruments in Indonesia. We assume the government levies these taxes and redistributes revenue back to households. Similarly, we assume that all plant profits are simply transferred back to households.

Manufacturing

The manufacturing sector is composed of risk-neutral plants that are heterogeneous in their productivity s_{it}. Each period, plants choose capital k and labor h on spot markets to produce output, while facing idiosyncratic risk over their future productivity s_{it} and time-varying changes in the economy. We denote plants' payoff-relevant aggregate state of the economy, including perfect anticipation of future changes in aggregates by Ω_t and make its components more explicit below. A plant's output y_{it} and taxable profits π_{it} at time t are given by:

$$y_{it}(s_{it}, z_t, h_{i,t}, k_{i,t}) = z_t s_{it} h_{it}^{\theta} k_{it}^{\alpha} \tag{4}$$

$$\pi_t(s_{it}, h_{i,t}, k_{i,t}; z_t, w_t) = (1 - \tau_t^C)\Big((1 - \tau_t^{VAT}) y_{it}(s_{it}, z_t, h_{i,t}, k_{i,t}) - w_t h_{i,t} - R_t k_{i,t}\Big) \tag{5}$$

where the production function is as in Section 2.2.3, τ_t^{VAT} gives the Indonesian value-added tax and τ_t^C the corporate income tax that is levied on taxable profits. Given the frictions in this economy, both tax instruments generally distort input choices. R_t gives plants' capital borrowing rate, which is equal to the deposit rate r_t plus depreciation δ assuming competitive rental markets. Idiosyncratic risk s_{it} follows a Markov process of order one. We further assume that s_{it} is exogenous and independent of aggregate technology z_t. For z_t, we leave the path unrestricted, but assume it is exogenous.

What drives slow plant adjustments? We assume that plants face labor adjustment costs. These capture, for example, managerial time constraints that arise from the time it takes to hire, fire and reorganize production tasks, a key constraint for plant growth in developing countries (Bloom et al. 2013, 2020).[13] Following the literature, we model them as follows

[13] In Appendix B.1, we provide a simple microfoundation in terms of the costs of scarce managerial time to show how organizational changes induce convex costs. An alternative interpretation of convex adjustment costs is given in labor search models where they are rationalized via convex (reduced-form) hiring or vacancy posting costs (e.g. Bilal et al. 2022; Coşar, Guner, and Tybout 2016)). The key difference is that adjustment costs in search models become partly functions of equilibrium outcomes such as market tightness. We abstract from this general equilibrium mechanism here given that the primary focus of the paper is on longer run

(e.g. Cooper, Gong, and Yan 2018; Cooper, Haltiwanger, and Willis 2015):

$$AC(h_{i,t-1}, h_{i,t}) = \begin{cases} F^{+} + c_0^{+}(h_{i,t} - h_{i,t-1}) + \frac{c_1^{+}}{2}\left(\frac{h_{i,t}-h_{i,t-1}}{h_{i,t-1}}\right)^2 h_{i,t-1} & \text{if } h_t > h_{t-1} \\ 0 & \text{if } h_t = h_{t-1} \\ F^{-} + c_0^{-}(h_{i,t-1} - h_{i,t}) + \frac{c_1^{-}}{2}\left(\frac{h_{i,t}-h_{i,t-1}}{h_{i,t-1}}\right)^2 h_{i,t-1} & \text{if } h_t < h_{t-1} \end{cases} \tag{6}$$

where F are fixed adjustment costs that capture overhead in dealing with hiring (F^{+}) or firing (F^{-}) and c_0 captures per worker hiring and firing costs. Importantly, there are convex adjustment costs whose importance is captured by c_1 and which capture costs of growing (c_1^{+}) or shrinking (c_1^{-}) plants quickly. Convex adjustment costs are key to explain the slow growth of plants over time and are a key determinant of the speed of transition growth. We allow all costs to be asymmetric to accommodate that firing and hiring is often regulated differently. At last, we index all adjustment costs in terms of wages since wage indexation provides a simple way to let costs grow with the economy. Besides this indexation, in the baseline model, we assume that adjustment cost parameters are fixed over time, but the model and estimation can accommodate time-varying parameters and we return to this point in later sections.

On top of adjustment costs, we assume that plants face simple financing constraints. With convex adjustment costs it is costly to build up a plant's workforce quickly. In the absence of further constraints, this can even induce unproductive plants to build up a large workforce while running consistently negative profits, hoping for good future productivity realizations. To rationalize why plants in the data seldomly run negative profits and may sometimes be forced to lay off workers, we introduce two simple and intuitive financing constraints:

$$h_{it} \leq h_{it-1} \ \text{ if } \ \pi_{it} \leq 0 \ \ \text{(Dynamic growth constraint)} \tag{7}$$

$$w_t h_{it} \leq \kappa_t (1 - \tau_t^{VAT}) y_{it} \ \ \text{(Working capital constraint)} \tag{8}$$

The first constraint is a dynamic growth constraint: in case plants already run negative profits, they cannot further grow their size. One can think of this as an absence of risky venture capital in Indonesia. The second constraint captures a working capital constraint that may force plants to downsize, allowing them only to run wage bills that are a multiple κ_t higher than their output. Both constraints give a simple and parsimonious way to capture moral hazard problems that likely prevent plants from running large deficits.[14]

The timing of manufacturing production is summarized in Figure 10. At the beginning of a period, incumbents observe their current productivity and make production decisions. Plants' payoff-relevant aggregate state is given by: $\Omega_t \equiv \{z_t, w_t, r^*\}_t^{\infty}$. After production takes place, incumbent plants incur a fixed cost of production $c_{i,t}^F$, upon which plants decide whether they

growth dynamics and not business cycle variation in unemployment and market tightness.

[14] We provide a simple microfoundation of the working capital constraint in Appendix B.2.

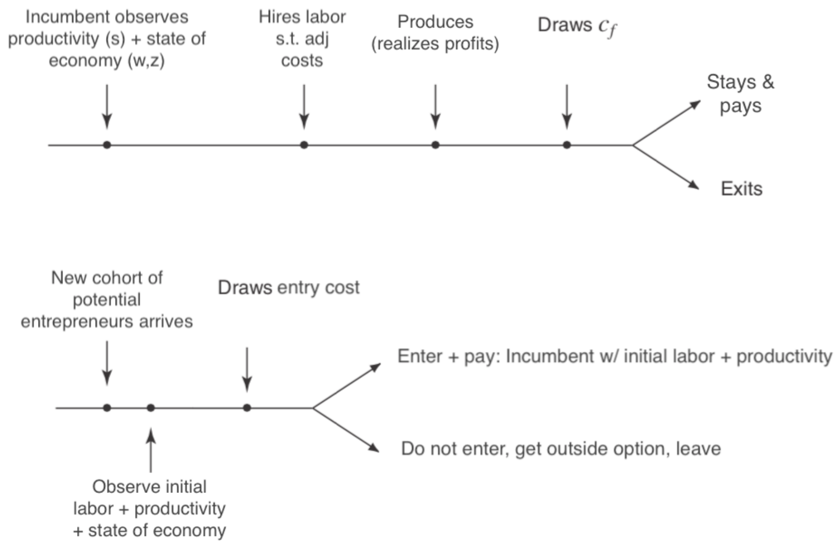

Figure 10: Timing in time t for manufacturing firms.

want to continue producing (and pay $c^F_{i,t}$) or permanently exit (as in Clementi and Palazzo 2016). The fixed cost is drawn from a distribution G, which we assume to be Gumbel with scale and variance parameters (μ^X_t, σ^X_t). A larger variance in the fixed costs rationalizes more overlap in the labor and productivity distributions of surviving and exiting plants. The exit decision of the plant depends on the plant's expected future value, the cost shock as well as the costs of closing down the plant (as in Hopenhayn and Rogerson 1993):

$$\max\Big\{\beta\mathbb{E}[V^M(s_{i,t+1}, h_{i,t}; \Omega_{t+1})|s_{i,t}, h_{i,t}; \Omega_t] - c^F_{i,t}, -C_E(h_{i,t})\Big\} \tag{9}$$

where V^M gives the continuation value of an incumbent plant and $C_E(h_{i,t})$ the costs of closing down the plant. The above maximization problem implicitly defines plant i's survival probability that the operating cost draw $c^F_{i,t}$ is lower than its future expected continuation value: $\lambda(s_{i,t}, h_{i,t}; \Omega_t) \equiv \mathbb{P}\Big(x \geq c^F_{i,t}\Big) = G\Big(x\Big)$ where $x \equiv \beta\mathbb{E}[V^M(s_{i,t+1}, h_{i,t}; \Omega_{t+1})|s_{i,t}, h_{i,t}; \Omega_t] + C_E(h_{i,t})$.

The ex-ante value of an incumbent manufacturing plant can be written in recursive form according to:

$$\begin{aligned} V^M(s_{i,t}, h_{i,t-1}; \Omega_t) = \max_{h_{i,t} \leq \bar{h}, k_{i,t}} \Bigg\{ &\pi(s_{it}, h_{i,t}, k_{i,t}; z_t, w_t) - w_t AC_t(h_{i,t-1}, h_{i,t}) + \lambda(s_{i,t}, h_{i,t}; \Omega_t)\Big\{ \\ &- \mathbb{E}_c[c^F_{it}|\text{stay}(s_{i,t}, h_{i,t}; \Omega_t)] + \beta\mathbb{E}\Big[V^M(s_{i,t+1}, h_{i,t}, \Omega_{t+1})|s_{i,t}, h_{i,t}, \Omega_{t+1}\Big]\Big\}\Bigg\} \end{aligned} \tag{10}$$

Plants have a common discount factor $\beta = 1/(1 + r^*)$, which is pinned down by the world interest rate. The presence of adjustment costs and financial constraints in combination with productivity dynamics makes this a dynamic problem since plants take into account that contemporaneous changes in their inputs influence adjustment and financing costs in the future.

Next, we consider endogenous plant entry. As visualized in Figure 10, each period there is a cohort of potential entrants (PE) of measure $|PE|_t$. Each potential entrant draws a random entry cost c^E_{it} from a distribution P, which they need to pay in case they start producing. Again, we assume that entry costs follow a Gumbel distribution with scale and variance parameters (μ^E_t, σ^E_t). Potential entrants differ in their idiosyncratic productivity s_{it} and their initial labor $h_{i,t}$, which they know when making the entry decision for producing in period t. The initial heterogeneous level of labor and productivity is key to capture that there is plant entry of many small as well as some very large plants that matter in the aggregate. It also accounts for the fact that we only observe and model plants with 20 or more workers. The exogenous distribution of potential entrants is given by $PE_t(h_t, s_t)$, which is time-varying due to exogenous reasons such as demographic changes that have been shown to be key for explaining variation in firm creation over time (Bernstein et al. 2022; Karahan, Pugsley, and Şahin 2019; Liang, Wang, and Lazear 2018). Potential entrant i with entry cost shock c^E_{it} enters if its net value is positive:

$$V_{PE}(s_{i,t}, h_{i,t}; \Omega_t) = \max \left\{ V^M(s_{i,t}, h_{i,t}; \Omega_t) - c^E_{it}, 0 \right\} \tag{11}$$

where we have normalized the outside option to zero. Similar to exit, this gives the following mapping $\mathbb{P}\Big(V^M(s_{i,t}, h_{i,t}; \Omega_t) \geq c^E_{it}\Big) = P\Big(V^M(s_{i,t}, h_{i,t}; \Omega_t)\Big)$. Note that in this specification the initial mass and distribution of entrants is endogenous, but entrants only start making input choices the period after they entered. A time-varying distribution of potential entrants also allows us to deal with plant entry jumps in years of the economic census as we further discuss in the estimation section. We denote the endogenous mass of entry for each state (h_t, s_t) in period t by $\mu(h_t, s_t)$, which is a function of Ω_t. Similarly, we define by $m(h_t, s_t)$ the endogenous mass of producing plants for each state (h_t, s_t) in period t. With slight abuse of notation, denote by M_t the set of producing plants at time t.

Rest-of-the-economy

We model the rest-of-the-economy parsimoniously as a representative firm with a decreasing returns to scale (DRS) production function:

$$Y^R_t = A_t \left(H^R_t\right)^{\theta_R} \quad \text{with} \quad \theta_R \in (0, 1) \tag{12}$$

where A_t is time-varying TFP, H_t^R gives labor employed and θ_R gives the output elasticity in the rest-of-the-economy. Decreasing returns to scale ensure that economy-wide wages can be affected by changes in manufacturing. The rest-of-the-economy sector takes as given productivity A_t and the wage rate w_t and chooses optimal labor demand maximizing per period profits: $\pi_t^R(A_t, w_t, \tau_t) = Y_t^R(A_t) - (1+\tau_t^R)w_t H_t^R$, subject to labor demand wedges τ_t^R. Labor demand is then given by:

$$H_t^{R*} = \left(\frac{\theta_R A_t}{(1+\tau_t^R)w_t}\right)^{\frac{1}{1-\theta_R}} \tag{13}$$

Labor demand wedges in the rest-of-the-economy are a simple way to capture observed variation in the labor intensity of output and one can think of them as changes in labor frictions. A_t captures changes in technology of the rest-of-the-economy. We allow both A_t and τ_t^R to change over time, but treat them as deterministic and exogenous paths.

Households

There is a continuum of households j that are characterized by their exogenous household-specific efficiency units of labor h_{jt} and whose exogenous mass at time t is denoted by L_t. Households supply labor inelastically so that the aggregate labor supply is given by $H_t = \int_j h_{jt} dj$. Changes in L_t and $\frac{H_t}{L_t}$ capture changes in the working population and education per worker respectively. We abstract from consumption-savings decisions by assuming that households are hand-to-mouth, simply consuming their labor income y_{jt} net of transfers from the government and plant profits T_{jt}: $c_{jt} = y_{jt} + T_{jt}$.[15] Households allocate their labor supply across both sectors based on maximizing labor income: $y_{jt} = \max\{h_{jt}w_t^M, h_{jt}w_t^R\}$.

Equilibrium

We assume that the observed growth path in the data is characterized by a path of per-period perfect foresight *Recursive Competitive Equilibria*.

Definition 3.1 (Model fundamentals.)**.** Model fundamentals at time t capture all exogenous model parameters, processes and distributions as given by: $\Theta_t^F = \{\theta, \alpha, \delta, F^-, c_0^-, c_1^-, F^+, c_0^+, c_1^+, \{A_t, \tau_t^R, H_t, PE_t, z_t, \tau_t^C, \tau_t^{VAT}, \kappa_t, \mu_t^X, \sigma_t^X, \mu_t^E, \sigma_t^E\}_t^\infty\}$. We further denote by $\bar{\Theta}_t^F$ the modified set of model fundamentals where all fundamentals are

[15] We only make this assumption to fix ideas. Given the small open economy setup, the domestic supply of capital is inelastic to changes in domestic savings behavior such that the production side – which is the focus of this paper – would look exactly the same if heterogeneous households would instead solve a savings and consumption choice. Given inelastic labor supply, the model is also isomorphic to one with uniform labor income taxes in both sectors.

fixed to their value at time t forever. At last, denote by $\Theta_t^F \setminus \{x = \bar{x}\}$ the modified set of model fundamentals where only model fundamental x is changed to $\bar{x}$.

Definition 3.2 (Initial distribution.)**.** The distribution of surviving plants from period $t-1$ gives the initial distribution at time t and is denoted by S_t.

A path of perfect foresight *Recursive Competitive Equilibria* starting at time t is then given by model fundamentals Θ_t^F, an initial distribution S_t, and endogenous sequences of prices $\{w_t, r^*\}_t^\infty$, corresponding quantities and distribution of producing plants $\{m_t\}_t^\infty$ such that each period $\tau \in [t, \infty)$:

1. The rest-of-the-economy sector statically chooses optimal labor demand maximizing profits taking as given productivity A_τ, the wage w_τ and labor demand wedges τ_τ^R.
2. Manufacturing plants choose optimal labor and capital demand.
3. Potential entrants optimally make entry and incumbents optimally make exit decisions.
4. Households inelastically supply total labor H_τ and optimally allocate labor across sectors to maximize labor income.
5. The aggregate wage w_τ adjusts to ensure that the labor market clears: $H_\tau = H_\tau^R(w_\tau, A_\tau, \tau_\tau^R) + \sum_{i \in M_\tau} h(s_{i\tau}, h_{i,\tau-1}; \Omega_\tau)$
6. The government runs a balanced budget by levying a value added and corporate income tax and redistributing revenue back to households.
7. The capital market clears every period such that international capital supply equals domestic capital demand: $\sum_{i \in M_\tau} k(s_{i\tau}, h_{i,\tau-1}; \Omega_\tau) = K_\tau^{INT}$.
8. The mass of active plants in τ and previous aggregate state $\Omega_{\tau-1}$ is equal to surviving plants from $\tau - 1$ plus endogenous new entrants:

$$\forall (s_\tau, h_\tau): \quad m(s_\tau, h_\tau; \Omega_\tau) = \sum_{s_{\tau-1}, h_{\tau-1}} \Big(\mathbb{1}_{h^* = h_\tau} \mathbb{P}[s_\tau | s_{\tau-1}] \lambda(s_{\tau-1}, h_{\tau-1}; \Omega_{\tau-1}) \times m(s_{\tau-1}, h_{\tau-1}; \Omega_{\tau-1}) \Big) + \mu(s_\tau, h_\tau) \tag{14}$$

9. The goods market clears each period such that total production is either consumed or exported: $Y_\tau = \sum_{i \in M_\tau} y_{i,\tau} + Y_\tau^R = C_\tau + NX_\tau$ where $NX_\tau = EXP_\tau - K_\tau^{INT}$ are net exports.[16]

The observed growth path features a combination of changes in model fundamentals that move the economy's *steady state* and transition growth as the economy is trying to catch up to this steady state. We now formalize these concepts.

Definition 3.3 (Balanced Growth Path (BGP) and Steady State (SS))**.** Along a *BGP*, underlying technology in both sectors (A_t, z_t) and the endogenous wage grow at the same constant rate, while all remaining model fundamentals and the endogenous distribution of

[16] Note that we have implicitly treated all entry costs, fixed costs and adjustment costs as shadow costs here, as they neither directly enter labor market clearing nor the aggregate resource constraint.

plants stays constant (see details in Appendix B.3). A *steady-state* is a BGP for which the growth rate is zero. Both *BGP* and *steady-state* are uniquely defined by *model fundamentals* Θ^F that admit a BGP/SS.[17]

Definition 3.4 (Transition path)**.** The unique perfect foresight *transition path* starting at t towards a BGP is defined by an initial distribution S_t and model fundamentals Θ_t^F (which admit a BGP), and gives a path of equilibrium wages over the transition.[18]

In Section 4, we use the model to separately quantify the role of *transition growth* from changes in the *steady state*.

3.2 Estimation

The model captures a race between changes in model fundamentals and the distribution of plants (over the state space) trying to catch up to these changes. We now take this model to the data and show how to disentangle changes in the distribution from changes in model fundamentals. Estimation proceeds in three main steps, while an additional step is needed for model counterfactuals. The first step identifies equilibrium prices – only wages in our case – in the data. We take an equilibrium estimation approach (Hotz and Miller 1993; Bajari, Benkard, and Levin 2007; Caliendo, Dvorkin, and Parro 2019), which means that we treat our model as generating the equilibrium wage path we observe in the data. We can thus treat the path of equilibrium wages as fixed throughout the estimation and only need to solve for changes in the equilibrium wage path for counterfactuals. This greatly simplifies the estimation as it avoids solving for the equilibrium path of the model during the estimation. In the second step, we identify the distribution over the entire state space of the economy over time and use this to back out related model fundamentals such as the initial distribution. The third step then solves for remaining fundamentals that are related to the dynamic input and exit choices of plants drawing on observed choices of plants conditional on the state space. In this step, we also need to make an explicit assumption about the evolution of model fundamentals beyond the time frame of our data. To conduct model counterfactuals, we further back out model fundamentals that are not needed to solve the baseline economy, such as fundamentals of the rest of the economy.

With each estimation step, we enforce more model structure and assumptions, making parameter identification very transparent. An important benefit of our approach is that we can directly draw on the production function and aggregate technology estimates discussed

[17] Uniqueness depends on the uniqueness of the individual policy functions and the unique mapping between policies and prices. We treat our numeric algorithm as formally defining the equilibrium refinement conditions sufficient for uniqueness.

[18] Uniqueness of the *transition path* can be proven via contraction mapping arguments over the path of price expectations. Again, our numeric algorithm for the transition gives a unique path for perfect foresight equilibria and we treat this as formally defining the equilibrium refinement conditions sufficient for uniqueness.

in Section 2.2.3 whose identifying assumptions nest our model. Table 1 provides an overview of the estimation steps and all model fundamentals and estimates. We now discuss each step in more detail. Given the larger number of model fundamentals, we focus throughout on the most important parts and relegate details to Appendix B.5.

3.2.1 Step 1: Equilibrium wage estimation

In the first step, we estimate the path of equilibrium wages that clear labor markets in each period. While prices could in principle be directly observed in the data, based on our model, we only observe plants' wage bills ($w_t h_{it}$), which are a combination of the wage and the quantity of labor. We are only interested in changes over time and thus normalize the level of initial wages w_0 to unity. Ideally, we would like to capture changes in the wage by looking at wage changes for a worker whose efficiency units of labor remained constant. This would for example avoid any assumptions on how workers with different skills select across plants. In the absence of worker-level data that spans the entire time period, we instead draw on changes in within-plant per worker wages for similar job types, exploiting that the Indonesian data reports wages and the number of workers separately for production and non-production work. Our identification strategy for the wage allows for arbitrary sorting of workers with different skills not only across plants but also across different job types within plants, but restricts changes in the skill sorting within job types over time. Formally, we assume that plant i uses on average the same skills per worker within job types k: $h_{it}^k / l_{it}^k = \alpha_i^k \cdot \varepsilon_{it}^k$. ε_{it}^k allows job types within plants to vary in their skill intensity around α_i^j over time. With standard restrictions on ε_{it}^k, this ensures that wages are identified from:

$$\mathbb{E}_i \left[log(w_{t+1} h_{it+1}^k / l_{it+1}^k) - log(w_t h_{it}^k / l_{it}^k) \right] = log(w_{t+1}) - log(w_t)$$

As our estimate of changes in wages, we use median within-plant-worker-type changes in wages, weighting observations by the average of total workers of type k between t and $t+1$, ensuring that wages are identified from median wage changes of workers (not plants). If anything, we think this estimator overestimates wage increases, because (1) any increase in within-worker human capital (e.g. on-the-job learning) will be attributed to increases in the wage, and (2) the estimates are for surviving plants, which might see more wage growth.

Figure 11 plots the estimated real wage series in the data. Our preferred estimator uses wage bills for production workers only, as production workers are relatively homogeneous and thus the identifying assumption is more likely to hold. For completeness, we also report estimated wage series using non-production workers and pooling all workers. Our estimates show that wages per efficiency unit of labor increased by more than 85% over the 40-year period with important variation over time. Given that the average wage bill per manufacturing worker increased roughly 4-fold in the data, these estimates imply that the average manufacturing

Table 1: Overview of parameter identification and estimation

Object	Description	Type	Identification idea	Value	Details
<u>**Parameterization:**</u>					
r^*	World interest rate	F	Risk-free rate	0.04	
δ	Depreciation rate	F	Standard	0.1	
τ_t^C	Corporate tax	F	Official rate	0.2	Section A.2.4
τ_t^{VAT}	VAT	F	Official rate	0.1	Section A.2.4
<u>**Estimation:**</u>					
Step 1:					
θ	Prod function	F	Control function	0.694	Section 2.3.2
α	Prod function	F	Control function	0.03	Section 2.3.2
w_t	Wage path	E	$\Delta_i(w_t h_{it})/l_{it}$	Fig. 12	Section 3.2.2
κ_t	Borrowing constraint	F	Max labor share	1.7	Section A.2.4
Step 2:					
z_t	Techn path	F	$\Delta_i$productivity$_{it}$	Fig. 8	Section 2.3.2
$\mathbb{P}(s'\|s)$	Transition matrix	F	Obs. transitions		Section 3.2.2
Init distrib		F	Obs. survivors		Section 3.2.2
E_t	Entrants	E	Obs. entry		Section 3.2.2
Step 3:					
Adj costs		F	Euler CCC	Table 2	Section 3.2.2
Cost ratio	Fixed cost	F	Euler CCC	Table 2	Section 3.2.2
Cost level		F	Match mass 2015		Section 3.2.2
For counterfactuals:					
A_t, θ_R, τ_t^R	Rest-of-Economy	F	First-order condition		Section A.2.4
Entry costs		F			Section A.2.4
PE_t	Potential Entrants	F	Entrants + entry proba		Section A.2.4

<u>Details:</u> Types are: F(undamental) and E(quilibrium object). The former stay fixed in counterfactuals, the latter change endogenously. If applicable, reported standard errors correct for multi-step estimation procedure and cluster at the plant-level by using block bootstrap across all estimation steps (This is currently still work in progress).

Figure 11: Evolution of estimated real wage in Indonesian manufacturing

Notes: Based on within-task changes in the wage bill per worker by task (production and non-production). Data: Indonesian manufacuring census (1975-2015, 20+ workers).

worker in 2015 was about 2.2x more efficient than the average manufacturing worker in 1975. We find this a reasonable estimate given the large educational gains of Indonesian workers observed over this period. As external validation, Gathen (2021) also finds a similar wage increase from separate estimates on Indonesian worker data between 1998 and 2015.

3.2.2 Step 2: Mapping the entire distribution over the state space

In the second step, we map any plant in our data to the state space of our model: $(h_{it-1}, s_{it}, w_t, z_t)$. While our data captures a discrete number of plants, we treat this mass as continuous for the estimation. This allows us to identify changes in the entire distribution of plants over the state space over time, which is crucial to determine the potential for transition growth. Specifically, we use this mapping to identify two key model fundamentals – the productivity process and the initial distribution – and an equilibrium object which changes in counterfactuals: the distribution of entrants.

To identify the state space, we draw on the wage estimates from Step 1, which together with a plant's wage bill $(w_t h_{it})$ identifies h_{it-1}. To obtain plant productivity and technology, we draw on the identification and estimation approach from Section 2.2.3, which hold under weak assumptions on exit and input choices. Having identified plant productivity s_{it} (up to a normalization of z_t), we estimate the dynamic process of s_{it} by discretizing s and then estimating the transition matrix $\mathbb{P}(s'|s)$ non-parametrically using (pooled) within-plant productivity changes in the data. The transition matrix is a fundamental of the economy and is identified based on within-plant changes in productivity conditional on previous productivity, and only requires that all productivity states are observed with positive probability at some point.[19]

[19] While an imperfect ex-post test, we check ergodicity of the implied idiosyncratic productivity process s_{it}

Next, we exploit the state space mapping to identify the initial distribution of surviving plants over $(s_{i,t}, h_{i,t-1}; \mathbf{\Omega})$ in 1976. While our data starts in 1975, the first year for which we can identify $h_{i,t-1}$ is 1976. We treat this initial distribution as a model fundamental, implicitly assuming that in any counterfactual, initial survivors do not anticipate any changes to the baseline equilibrium paths of wages and technology prior to 1976. The main benefit of directly taking the initial distribution from the data is that we can remain agnostic about its origins and allow the data to reveal the initial degree of misallocation. A downside is that if the model does not capture all mechanisms of dispersion over the state space, the initial distribution may look more "misallocated" than it actually is; leading to overestimating model-implied transition dynamics.

At last, the state space mapping also allows us to identify time-varying entrant distributions $E(s_t, h_t; \mathbf{\Omega})$, which are equilibrium objects but related to the fundamental potential entrant distributions via: $PE_t(s_t, h_t; \mathbf{\Omega}) = E_t(s_t, h_t; \mathbf{\Omega})/\mathbb{P}_E(s_t, h_t; \mathbf{\Omega})$, where $\mathbb{P}_E(.)$ gives the entry probability, which is a function of the model-implied value of entering as well as the parameters of the entry cost distribution. For the model estimation along the baseline equilibrium path, we treat these equilibrium entrant distributions as fixed. Our approach implies that the baseline model exactly replicates observed plant entry. This is in contrast to plant-level exit and labor demand decisions, for which our estimation approach allows the model to fail.

3.2.3 Step 3: Estimating the dynamics of the model

Step 3 reveals the remaining parameters of the economy that are needed for the baseline model: fixed cost parameters that govern entry and exit decisions as well as adjustment cost parameters that govern how plants make dynamic labor choices. This step enforces more model structure, particularly on how plants make dynamic input, exit and entry choices, and how plants form expectations over the future. We separate this section into two parts with the first part only exploiting optimal plant-level choices across consecutive periods, while the second part also enforces long-run expectations.

Euler equation CCC estimation We identify most remaining parameters by exploiting observed exit and labor input choices conditional on the state space, drawing on standard conditional choice probability (CCP) and continuous conditional choice (CCC) Euler estimation techniques (Hotz and Miller 1993; Bajari, Benkard, and Levin 2007). Taking first-order conditions with respect to labor from the incumbent's value function above and directly

by verifying that all states in the discretized transition matrix can be eventually reached. In Appendix B.5.3, we provide details on the discretization of productivity and labor, which we rely on for numerically solving the model and counterfactuals.

plugging in the envelope condition, we obtain the following Euler equation:

$$
\begin{aligned}
0 = & \underbrace{\frac{\partial \pi(s_{i,t}, k_{it}, h_{i,t}, z_t)}{\partial h_{i,t}}}_{\text{Labor wedge}} - \underbrace{w_t \frac{\partial C_h(h_{i,t}, h_{i,t-1})}{\partial h_{i,t}}}_{\text{Current marginal adj costs}} \\
& + \underbrace{\frac{\partial \lambda(s_{i,t}, h_{i,t}, \Omega_t)}{\partial h_{i,t}} \left\{ -\tilde{g}(s_{i,t}, h_{i,t}, \Omega_t) + \beta \mathbb{E}\Big[V(s_{i,t+1}, h_{i,t}, \Omega_{t+1}) | s_{i,t}, h_{i,t}, \Omega_t \Big] \right\}}_{\text{Marginal benefit on survival}} \\
& + \underbrace{\lambda(s_{i,t}, h_{i,t}, \Omega_t) \left\{ -\frac{\partial \tilde{g}(s_{i,t}, h_{i,t}, \Omega_t)}{\partial h_{i,t}} + \beta \mathbb{E}\left[-w_{t+1} \frac{\partial C_h(h_{i,t+1}, h_{i,t})}{\partial h_{i,t}} | s_{i,t}, h_{i,t}, \Omega_t \right] \right\}}_{\text{Marginal benefits on future costs}}
\end{aligned}
\tag{15}
$$

where we have used $\tilde{g}(s_{i,t}, h_{i,t}, \Omega_t)$ to denote the expected fixed cost conditional on surviving to emphasize that it is a function of the state space.

The Euler equation, which holds for any plant that is optimally adjusting labor, says that plants should equalize today's marginal product of labor with the marginal costs of labor and current as well as future labor adjustments. Adjustment costs give a natural explanation for why there is a "wedge" between the static marginal product and the marginal costs of labor (Hsieh and Klenow 2009). For our estimation purposes, the important features of the Euler equation are that it holds along the transition, and that it gives a nonlinear equation in observable plant-level choices (exit and input choices) and parameters that govern survival probabilities as well as adjustment costs. Specifically, marginal adjustment costs are a function of adjustment cost parameters. The tricky terms are expected and marginal expected fixed costs, marginal survival probabilities and the expected future continuation value. As we show in Appendix B.6, the Gumbel distribution for the fixed costs ensures that we can analytically invert all of these terms as functions of observed survival probabilities and parameters of the Gumbel distribution.

Appendix Table A.1 presents the non-linear least squares (NLS) estimation results and Appendix B.6 gives the exact estimating equation and estimation details. The Euler equation flexibly identifies marginal adjustment costs. Intuitively, linear adjustment costs are identified from the observed labor wedge across plants and the probability of switching between shrinking and growing as determined by the volatility of the estimated productivity process. Convex costs instead scale with the labor growth and are thus identified from the variation in within-plant labor demand growth across periods, again conditioned by the observed volatility of the productivity process. We find sizable adjustment costs – especially convex costs on growing – that rationalize why even productive plants (with a high labor wedge) conditional on previous plant employment do not grow faster. Quantitatively, our estimates imply that growing a plant's workforce by 20% within a year – a growth rate slightly above the 75th percentile – leads to adjustment costs that are about 75% of the previous wage bill. Informed by faster observed shrinking conditional on productivity, convex costs on shrinking are estimated to be less than half as big. We also estimate that a plant pays almost 75% of a new worker's annual wage in the form of hiring costs, which is identified from the high observed wedge between the marginal product and wage and the high volatility of productivity that make any investments in the workforce risky. In Appendix B.6, we also report time-varying estimates of adjustment costs. If anything, we find that convex adjustment costs tend to increase over time beyond what is implied by increases in the wage, pointing away from a reduction in frictions driving Indonesian growth.

Solving the baseline model The Euler equation only identifies the ratio between the level and scale of the fixed cost distribution that determine plant exit. To separately identify the level, we solve the model and match one moment in the data and model: the mass of plants in 2015, assuming that the census is complete for that year. The estimated level and scale of the fixed cost distribution rationalize average exit rates and the low but positive correlation with underlying productivity and size.

Solving the model introduces two issues that are common to equilibrium estimation approaches. First, by requiring to solve for plants' value functions, we need to make an explicit assumption on long-run expectations of plants after the year 2015 when our data ends. We assume that after 2015, plants expect to be on a balanced growth path with manufacturing technology growing at the average rate at which it grew in the preceding ten years.[20] Given that technology in manufacturing grew strongly since 2000, this assumption implies optimistic expectations, in line with low exit and strong observed plant growth in the years prior to 2015.

The second issue for the equilibrium estimation is that enforcing revealed equilibrium prices along the estimation of the baseline economy does not guarantee that these prices actually clear markets in our model over time. We ensure consistency – informed by our specific context and data availability – by treating the observed data as correctly revealing prices, but not necessarily correctly revealing aggregate labor demand and supply. As discussed in Section 2.1, our data does not correctly reveal aggregate labor demand and supply due to mis- and non-reporting and as explicitly taken into account in our measure of plant exit. We thus use the model-implied aggregate labor demand and supply for the baseline growth path and enforce the implied fundamentals that ensure market clearing for all future model-based counterfactuals. This is how our approach also ensures that counterfactuals are consistent with the baseline model economy.

[20] We provide technical details on the BGP and how to solve for the stationarized value function in 2015 in Appendix B.3. An alternative would be to solve for a continued transition towards a long-run BGP by making explicit assumptions on how all fundamentals evolve after 2015. We do not follow this approach, because it adds substantial additional computational costs while requiring similarly strong assumptions.

Figure 12: Baseline model fit

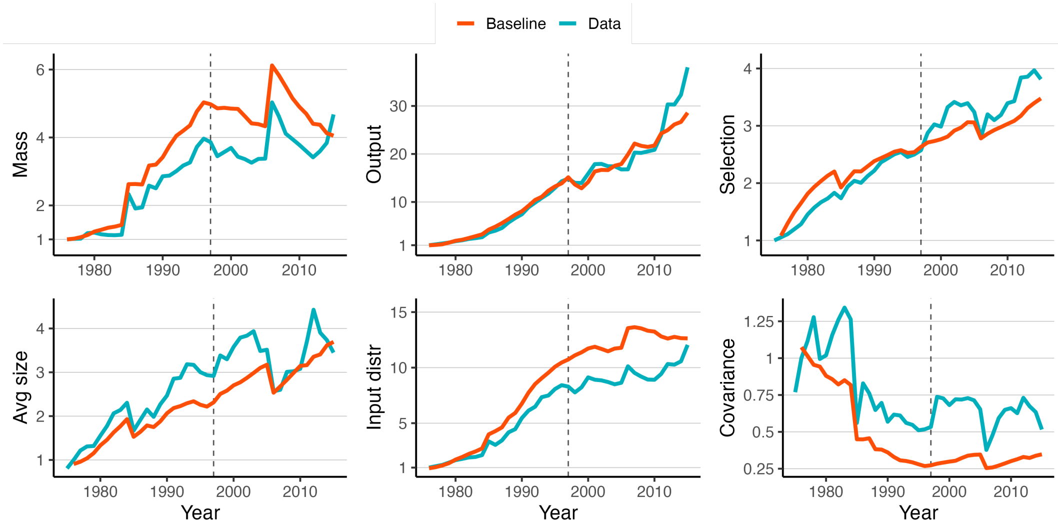

Notes: All graphs report results for manufacturing only and for each panel, both series are normalized to the initial level in the data. Average size is reported in efficiency units of labor. Selection (measuring average productivity), input distribution (the sum of the Cobb-Douglas aggregator of capital and labor) and covariance (the covariance between plant productivity and the share in inputs as given by the Cobb-Douglas aggregator) refer to the three corresponding terms in the growth accounting identity introduced in Section 2.

3.3 Evaluating model fit

To assess how well the model fits the data over time, we revisit three main results from Section 2. Having estimated the model on micro moments, we start out by moving from "micro to macro," evaluating the model's aggregate predictions. We then validate how well the model matches plants' dynamic input choices and changes in the entire plant distribution over time.

Figure 12 shows how the baseline model fits the mass of plants, aggregate output in manufacturing and the three main endogenous components that we used in Section 2.2.4 to formally decompose manufacturing growth. For completeness, we also plot the evolution of the average plant size as measured in efficiency units of labor. The model closely tracks the more than 30-fold increase in manufacturing output over time, including rapid growth in the absence of technology improvements until the Asian Financial Crisis, the decline during the crisis and the fast post-crisis growth. Overall, the model tends to slightly underestimate output at the plant level given the higher model-implied mass of plants.

The accounting identity helps to understand why the model fits the aggregate data well and where the model underperforms. First, as the main component of aggregate output growth, the model closely tracks the distribution of labor and capital across the endogenous plant distribution over time. Neither total labor demand, the total number of plants over time (apart from the first and last year) nor the distribution of inputs is hit by construction.

Figure 13: Model validation: Event study and changes in distribution

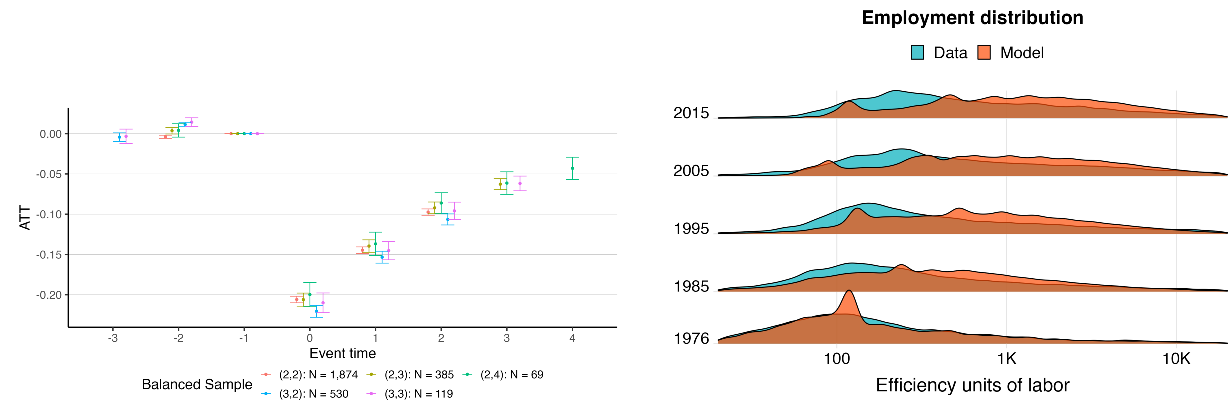

Notes: Left graph replicates event study exactly as in Section 2 using simulated data from model growth path. Right graph gives changes in the employment distribution (using efficiency units of labor) for model versus data. For both data and model, wage bills are used and divided by (same) estimated wage path.

The model captures the right degree of slow labor accumulation across the entire plant size distribution over time, which can be even more clearly seen from looking at the average plant size. Secondly, the model tracks well the evolution of average productivity across plants, capturing well the endogenous selection of plants over time. If anything, the model overpredicts productivity growth in the early years and underpredicts towards the end. This can be in part explained by a too fast productivity convergence implied by the estimated productivity process, stemming from frequent temporary productivity shocks that lead to overestimating productivity transitions. At last, the model also performs reasonably well on the most difficult part: the endogenous evolution of the joint distribution of productivity and inputs, as captured by the covariance term. Here, the model captures the decline in the covariance over time, as many small and productive plants enter and resources only reallocate slowly due to sizable adjustment frictions.

Next, we are interested in whether the model is also in line with the micro-level labor dynamics. Figure 13 Panel A shows the same estimated event study as in the data, using the exact same treatment definition and sample restrictions but now using simulated data from our baseline model from 1976 to 2015. The event study results align well. Both pre-trend, the exact magnitude of the treatment effect at impact and the slow recovery of the labor share follow the data. We also note that the balanced sample restrictions imposed in the data lead to very similar sample sizes in the simulated data indicating that we identify similarly selected sets of plants (which received a rare permanent productivity shock).

At last, Figure 13 Panel B compares model-implied versus observed changes in the entire employment distribution of plants over time. In 1976, the first year that our model predicts plant decisions, the distributions are still largely indistinguishable. Over time, the employment distribution moves strongly to the right with average employment increasing almost 4-fold and the mass of the distribution shifting from a strong left tail towards the right. The model tracks this overall change well, but slightly overpredicts the right tail. Importantly, we do not see a marked deterioration of the distribution even after 40 years of endogenous evolution.

4 Quantifying the drivers of aggregate growth

Using the estimated model, we now quantify the drivers of growth. Specifically, we quantify (1) the importance of initial transition growth, (2) the continuing importance of transition growth over the course of development, and (3) the role of policy. We present each in turn.

4.1 Initial conditions and the role of transition growth

We start by quantifying the importance of transition growth based on the initial economy at the onset of the Indonesian growth miracle in 1976. How much would the 1976 economy have grown purely from transition growth in the absence of any further changes in model fundamentals? For this, we start from the initial economy with the initial distribution M_{1976}, fix initial model fundamentals $\bar{\Theta}^F_{1976}$ to their value in 1976 and solve for the perfect foresight transition path (see Section 3), which includes solving for a counterfactual path of equilibrium wages that clears the labor market over time as the initial distribution of plants transitions towards the steady state distribution defined by $\bar{\Theta}^F_{1976}$.

Figure 14 highlights the resulting counterfactual growth in manufacturing. Overall output in manufacturing increases roughly 12-fold over time, accounting for 42% of the overall output gains compared to the baseline (model) economy. The reason is that young and small plants – which dominate the initial distribution and new entrants – gradually hire more workers and increase their productivity through a combination of productivity convergence and the exit of less productive plants. At the same time, entry consistently exceeds exit and the mass of plants gradually doubles over time. The increase in workers across manufacturing plants is mostly driven by the reallocation of labor from the rest of the economy: in the absence of observed changes in the aggregate labor supply and technology in the rest of the economy, the model predicts that Indonesia would have seen a manufacturing miracle with manufacturing labor and output shares reaching close to 30% over 40 years (in contrast to observed shares of less than 10%). Average plant size increases far more rapidly because in the absence of productivity improvements in the rest of the economy, aggregate wages stay almost 40% lower than in the baseline economy. Hence, initial conditions explain more than all of the increase in the average plant size over time, with cheap labor being the main driving force. Looking at changes in the entire distribution, Figure A.20 highlights that the increase in the average plant size is indeed driven by a movement in the right tail, which forms slowly

Figure 14: Growth from initial conditions

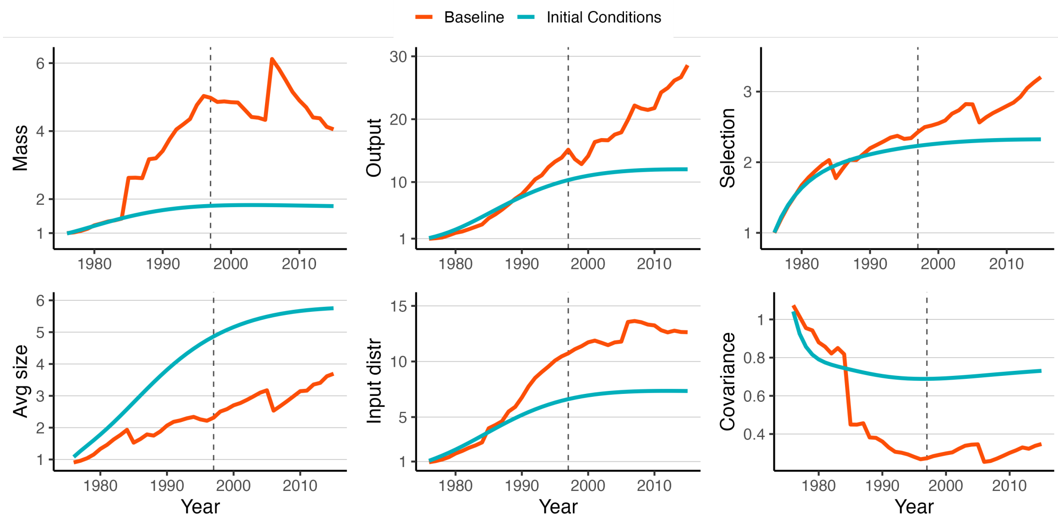

Notes: Results for counterfactual where economy evolves only based on initial conditions (all fundamentals fixed to initial level). All graphs report results for manufacturing only. Average size is reported in efficiency units of labor. Selection (measuring average productivity), input distribution (the sum of the Cobb-Douglas aggregator of capital and labor) and covariance (the covariance between plant productivity and the share in inputs as given by the Cobb-Douglas aggregator) refer to the three corresponding terms in the growth accounting identity introduced in Section 2.

because it takes time to grow large plants and the initial distribution lacks large plants, in line with the empirical evidence in Section 2.2.2.

What about economy-wide effects? Aggregate output per worker increases by roughly 25% by 2015 in this counterfactual economy, explaining 5.2% of the close to 6-fold increase in aggregate output per worker observed between 1976 and 2015. However, this comparison may be unfair given that a large part of the 6-fold increase in aggregate output per worker is driven by changes in the rest of the economy, not by manufacturing. In the end, manufacturing in our data captures less than 10% of aggregate output in the economy. We thus also consider a counterfactual in which only the rest of the economy fundamentals change as observed, but manufacturing fundamentals and the initial distribution of manufacturing plants stay fixed to their values in 1976. Using this counterfactual to "purge" the effects of changes in the rest of the economy and isolate the effects of changes in manufacturing only, we find that initial transition growth accounts for all (117%) of the aggregate output per worker gains that are due to changes in manufacturing by 2015.

4.2 The never-ending race: Transition growth remains important

As evidenced in the previous section, even in the absence of further changes in model fundamentals, transition dynamics – due to slow labor hiring and firing and slow plant entry

Figure 15: Distance to steady state over time

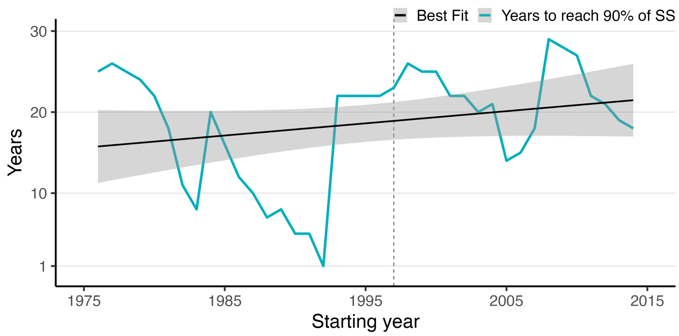

Notes: Years to reach 90% of the steady state manufacturing output for each transition path over 1976-2015. Each year's perfect foresight equilibrium transition path starts from that year's initial distribution and fixes fundamentals of that year over the transition. Best linear fit includes 95% CIs. Jumps are partly driven by census years in which potential entrant distributions change more strongly.

and exit dynamics – take decades to play out. A key question is whether the Indonesian economy runs out of transition growth over time as new changes in model fundamentals provide new potential for transition growth. We find that the quantitative answer to this question is no. We show this result by revisiting the previous exercise but instead of computing the transition path based only on the initial distribution in 1976 and initial model fundamentals $\bar{\Theta}^F_{1976}$, we compute the perfect foresight transition paths for each year between 1976 and 2015 starting from that year's initial distribution and model fundamentals $\bar{\Theta}^F_t$. This gives a total of 40 different counterfactual transition paths with their corresponding counterfactual equilibrium wage paths and plant distributions. As a measure of the transition potential, we then calculate for each transition path the number of years it takes to reach 90% of the (time-varying) steady state manufacturing output. Figure 15 shows that it takes the 1976 economy 26 years to come close to the steady state if fundamentals were to remain constant at $\bar{\Theta}^F_{1976}$. On average, it takes about 20 years and, importantly, the number of years to come close to the steady state does not systematically decline and – if anything – increases over time.

This race between catching up to the steady state and changes in the steady state itself can only be studied in a model that features both transition growth and changes in fundamentals and we find strong quantitative evidence that due to the combination of large and frequent changes in fundamentals and slow transition dynamics, the Indonesian economy does not get closer to its time-varying steady state. Large demographic changes and policy changes pre-1975 also provide a simple explanation for why the initial Indonesian economy in 1976 was far away from its steady state. At last, potential for transition growth may not always provide a positive force for economic growth; in fact, after 40 years of demographic changes, the mass of plants in 2015 is above its steady state and transition growth is now negative as the mass of plants slowly declines along the transition – an important consequence of an aging population.

4.3 The role of policy

While the Indonesian growth experience is driven by a never-ending race of transition growth and changes in model fundamentals that induce new transition growth, the question is how government policy enters. The simple answer is: policy drives part of the changes in model fundamentals. Thus, to evaluate the effect of policy, we need to link changes in policies to changes in model fundamentals. In this section, we show how to do this by focussing on two specific but very important Indonesian government policy changes since 1975: education reform that maps to changes in human capital (and thus aggregate effective labor supply), and changes in Indonesia's foreign direct investment (FDI) policy that map to changes in the distribution of potential foreign entrants over time. In both cases, we first quantify the overall effect of changes in the specific model fundamental and then quantify the (relative) effect that can be attributed to specific policy changes. We show that overall changes in human capital were large and a necessary condition for Indonesia's manufacturing take-off, but observed education policies only explain 5% of this effect. In contrast, we find that the overall growth effects of FDI in Indonesian manufacturing were modest, but that observed changes in FDI policy explain up to 85% of its effects.

4.3.1 The role of cheap labor & the INPRES school construction program

What are the economy-wide and manufacturing growth effects of dramatic increases in human capital in the Indonesian economy? Over the period 1976 to 2015, our estimates suggest that human capital per worker H_t/L_t increased by 220%. To quantify the overall effects of human capital increases, we consider a counterfactual in which the Indonesian economy had not seen any human capital per worker increases over time. That is, we consider a counterfactual growth path where we start from the initial distribution in 1976 and a counterfactual path of model fundamentals with a modified path for the aggregate labor supply: $\Theta^F_{1976} \setminus \{H_t/L_t = H_{1976}/L_{1976}\}^\infty_t$. To quantify the extent to which policy contributed to the overall increases in human capital, we evaluate the effects of a particular educational policy change. Namely, we evaluate the largest school construction program in Indonesia's history and one of the largest in the world: the 1970s INPRES school construction program. The program successfully led to increases in schooling and wages (see: Duflo 2001; Akresh, Halim, and Kleemans 2023). We assume the program only affected the Indonesian economy through its effects on human capital. To quantify the effects of the INPRES program, we consider a counterfactual in which all gains in human capital materialized except the ones that were due to the INPRES program. For this we construct a counterfactual path of aggregate human capital in the absence of the INPRES program, drawing on existing estimates on the wage effects of the program (which map to marginal changes in human capital), the known scale of the program and the increasing share of treated cohorts over time (details in Appendix C.2).

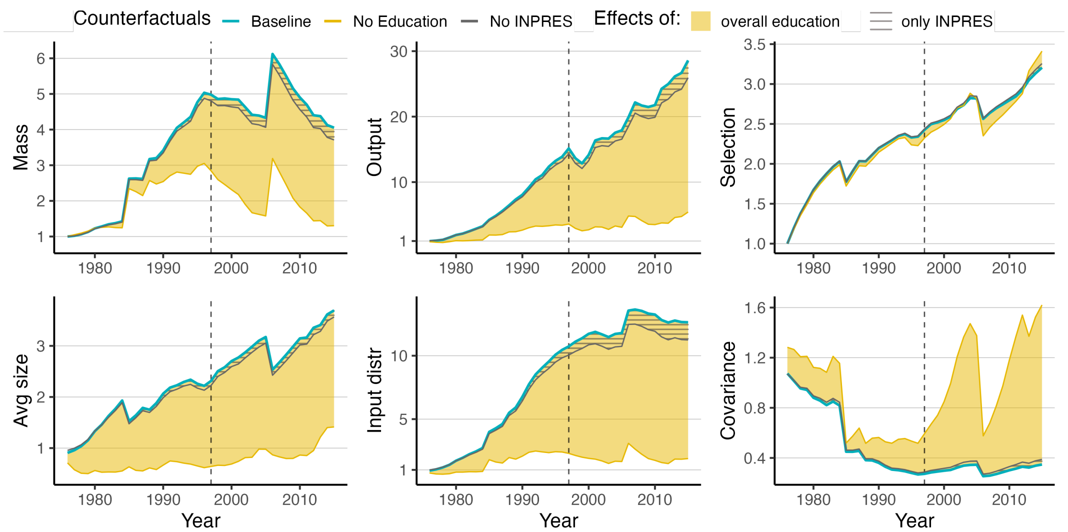

Figure 16: The role of education in Indonesia's manufacturing growth.

Overall, we find that the estimated 220% increase in human capital per worker increased aggregate output per worker by 26.5% in 2015. The seemingly small aggregate effect is explained by output in the rest of the economy being less dependent on labor, as captured by a low estimated labor elasticity. Figure 16 visualizes the quantitative effects that increases in human capital had on manufacturing growth over the period 1976-2015. Manufacturing growth in Indonesia relied heavily on the cheap labor that increases in human capital brought; in the absence of this effective increase in the supply of labor ("No Education"), wages would have roughly doubled and Indonesia would not have developed a successful manufacturing sector, with manufacturing output and its employment share not even reaching 1/4 of their historical level. Average plant size would have only increased marginally over time, far less plants would have entered and far more would have exited. The INPRES program only accounts for a small share of these effects, roughly driving 5% of the aggregate output per worker gains from increases in human capital. Figure 16 shows that the effect on manufacturing output, the mass of producing plants and overall hiring, is approximately twice as big as the aggregate effects. This is because manufacturing is more sensitive to labor costs than the rest of the economy. Furthermore, the positive effects of the INPRES program slowly increase over time as more cohorts of Indonesians that benefited from the new-built schools enter the labor market.

4.3.2 The role of foreign ownership and FDI policy

Next, we look at the role of foreign direct investment in manufacturing, which we define as the foreign ownership of manufacturing plants. Foreign-owned manufacturing plants are quantitatively important, accounting for roughly 30% of manufacturing output in 2015 (see Figure A.21 in Appendix C.3) and the aggregate importance of foreign ownership increased steadily since the late 1980s. FDI policy primarily affects the entry of foreign-owned plants, since ownership shares are highly persistent and most variation in foreign ownership is across, not within plants. We thus assume that FDI only affects the Indonesian growth experience through changing the distribution of potential entrants – a model fundamental that is robust to time variation in the incentives to enter – and consider counterfactual growth paths in which we only change the path of potential entrant distributions. Again, we want to separately quantify the effects of FDI and the relative effect that changes in observed FDI policy had on FDI. For this, we separate the distribution of potential entrants at any point in time into the distribution of potential foreign entrants and potential domestic entrants enforcing model-consistent entry decisions. We then construct a counterfactual path of potential entrant distributions without foreign entrants. To capture the effect of policy, we consider important regulatory changes in FDI policy in 1987. Specifically, we exploit variation in potential foreign entrant distributions right before and after the reform to measure the effect of policy and use the estimated effect to construct a counterfactual path of potential foreign entrant distributions in the absence of the FDI policy change (details in Appendix C.3).

We find that FDI helped manufacturing growth, but did not play a transformative role. Specifically, the entry of foreign-owned manufacturing plants explains 7.5% of the aggregate output per worker gains due to manufacturing growth and we estimate that manufacturing output and the manufacturing employment share would be 8% lower in 2015 in the absence of FDI. The reason for this rather small effect is that given a high estimated supply of domestic potential entrants, the downward pressure on labor demand and wages due to the disappearance of foreign entrants leads to an elastic response of domestic entry in general equilibrium that mitigates some of the negative effects of losing FDI. In contrast to the case of education policy, we find that *changes* in FDI policy potentially explain most of the overall growth effects from FDI. Specifically, changes in FDI policy potentially explain a four- to five-fold increase in potential entry and these changes in FDI policy in turn explain 85% of the overall growth effects from FDI.

5 Conclusion

This paper studied the drivers of growth miracles. Building on 40 years of plant-level manufacturing panel data for Indonesia, we motivated a model in which rapid growth is driven by a combination of transition growth and changes in fundamentals that are dominated by worker and plant demographics. We showed how to tractably estimate this model on the observed growth path using standard plant-level data and without assuming that the observed economy is at a steady state at any point in time. We found that transition growth is key: 42% of the observed manufacturing output growth is simply explained by initial conditions in 1975 – dominated by young and small manufacturing plants – providing ample opportunities for catch-up growth. Transition growth also does not become less important over time because important demographic changes in the economy induce further potential for transition growth.

Since our model and estimation framework maps directly to observed time-varying aggregate growth and its micro-level drivers of endogenous changes in the distribution of plants, it is particularly suited to study the dynamic growth effects of observed policy. This link is important not only to better validate macroeconomic models of growth, but also to study the aggregate growth and general equilibrium effects of policy – effects that are rarely identified in micro-empirical policy evaluations. In this paper, we only started looking at this by showing how to use the estimated model to evaluate the dynamic growth effects of two important Indonesian policies: education reform and changes in FDI policy. Based on our results, a somber conclusion – partly resonating related work on the Indian growth miracle (Bollard, Klenow, and Sharma 2013) – is that observed policy mattered less for growth than we might think. Instead, we find that Indonesian growth was mostly driven by structural forces related to demographics. This does not mean that policy necessarily plays no role. In fact, in Appendix C.4 we consider two sets of reduced-form policies – a reduction in (convex) labor adjustment frictions and an increase in annual technology growth – that both would have doubled Indonesian manufacturing output by 2015, an even more remarkable manufacturing miracle closer to experiences in countries such as China and Malaysia. Future research should further unpack what drives changes in adjustment costs, technology growth or the pool of potential entering plants and link these closer to policy.

6 References

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 (6): 2411–51.

Akcigit, Ufuk, Harun Alp, and Michael Peters. 2021. "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries." *American Economic Review* 111 (1): 231–75.

Akresh, Richard, Daniel Halim, and Marieke Kleemans. 2023. "Long-Term and Intergenerational Effects of Education: Evidence from School Construction in Indonesia." *The Economic Journal* 133 (650): 582–612.

Alessandria, George, Costas Arkolakis, and Kim J. Ruhl. 2021. "Firm Dynamics and Trade." *Annual Review of Economics* 13 (1): 253–80.

Allen, Treb, and Dave Donaldson. 2020. "Persistence and Path Dependence in the Spatial Economy." National Bureau of Economic Research.

Asturias, Jose, Sewon Hur, Timothy J. Kehoe, and Kim J. Ruhl. 2023. "Firm Entry and Exit and Aggregate Growth." *American Economic Journal: Macroeconomics*.

Atkeson, Andrew, and Ariel Tom'as Burstein. 2010. "Innovation, Firm Dynamics, and International Trade." *Journal of Political Economy* 118 (3): 433–84. https://doi.org/10.1086/653690.

Atkin, David, Amit K. Khandelwal, and Adam Osman. 2019. "Measuring Productivity: Lessons from Tailored Surveys and Productivity Benchmarking." In *AEA Papers and Proceedings*, 109:444–49. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203. https://www.aeaweb.org/articles?id=10.1257/pandp.20191005.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica* 75 (5): 1331–70.

Basri, M. Chatib, Mayara Felix, Rema Hanna, and Benjamin A. Olken. 2021. "Tax Administration Versus Tax Rates: Evidence from Corporate Taxation in Indonesia." *American Economic Review* 111 (12): 3827–71. https://doi.org/10.1257/aer.20201237.

Bernstein, Shai, Emanuele Colonnelli, Davide Malacrino, and Tim McQuade. 2022. "Who Creates New Firms When Local Opportunities Arise?" *Journal of Financial Economics* 143 (1): 107–30.

Bilal, Adrien, Niklas Engbom, Simon Mongey, and Giovanni L. Violante. 2022. "Firm and Worker Dynamics in a Frictional Labor Market." *Econometrica* 90 (4): 1425–62.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013.

"Does Management Matter? Evidence from India." *The Quarterly Journal of Economics* 128 (1): 1–51.

Bloom, Nicholas, Aprajit Mahajan, David McKenzie, and John Roberts. 2020. "Do Management Interventions Last? Evidence from India." *American Economic Journal: Applied Economics* 12 (2): 198–219.

Bollard, Albert, Peter J. Klenow, and Gunjan Sharma. 2013. "India's Mysterious Manufacturing Miracle." *Review of Economic Dynamics* 16 (1): 59–85.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. 2021. "Some Unpleasant Markup Arithmetic: Production Function Elasticities and Their Estimation from Production Data." *Journal of Monetary Economics* 121: 1–14.

Brandt, Loren, Johannes Van Biesebroeck, and Yifan Zhang. 2012. "Creative Accounting or Creative Destruction? Firm-Level Productivity Growth in Chinese Manufacturing." *Journal of Development Economics* 97 (2): 339–51.

Buera, Francisco, Joseph Kaboski, and Robert Townsend. 2023. "From Micro to Macro Development." *Journal of Economic Literature* 61 (2): 471–503.

Buera, Francisco, and Yongseok Shin. 2013. "Financial Frictions and the Persistence of History: A Quantitative Exploration." *Journal of Political Economy* 121 (2): 221–72.

Cali, Massimiliano, Marco Le Moglie, and Giorgio Presidente. 2021. "Gain Without Pain? Non-Tariff Measures, Plants' Productivity and Markups."

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro. 2019. "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock." *Econometrica* 87 (3): 741–835.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.

Chen, Zhang. 2022. "Economic Growth and the Rise of Large Firms." *Job Market Paper*. https://zhangc-econ.github.io/files/TailGrowth.pdf.

Choi, Jaedo, Andrei A. Levchenko, Dimitrije Ruzic, and Younghun Shim. 2023. "Superstars or Supervillains? Large Firms in the South Korean Growth Miracle." https://www.freit.org/EIIT/2023/selected/choi.pdf.

Clementi, Gian Luca, and Berardino Palazzo. 2016. "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations." *American Economic Journal: Macroeconomics* 8 (3): 1–41.

Cooper, Russell, Guan Gong, and Ping Yan. 2018. "Costly Labour Adjustment: General Equilibrium Effects of China's Employment Regulations and Financial Reforms." *The Economic Journal* 128 (613): 1879–1922.

Cooper, Russell, John Haltiwanger, and Jonathan L. Willis. 2015. "Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications." *Research in Economics* 69 (1): 37–50.

Coşar, A. Kerem, Nezih Guner, and James Tybout. 2016. "Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy." *American Economic Review* 106 (3): 625–63.

Demirer, Mert. 2020. "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups." *UPenn. Job Market Paper*.

Duflo, Esther. 2001. "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment." *American Economic Review* 91 (4): 795–813.

———. 2004. "The Medium Run Effects of Educational Expansion: Evidence from a Large School Construction Program in Indonesia." *Journal of Development Economics* 74 (1): 163–97.

Eslava, Marcela, and John C. Haltiwanger. 2020. "The Life-Cycle Growth of Plants: The Role of Productivity, Demand and Wedges." National Bureau of Economic Research.

Eslava, Marcela, John Haltiwanger, and Alvaro Pinzon. 2022. "Job Creation in Colombia Versus the USA: Up or Out Dynamics Meet the Life Cycle of Plants." *Economica* 89 (355): 511–39.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83 (329): 91–129.

Gandhi, Amit, Salvador Navarro, and David Rivers. 2017. "How Heterogeneous Is Productivity? A Comparison of Gross Output and Value Added." *Journal of Political Economy*.

Gathen, Jonas. 2021. "Government Jobs for Talents, Cronies or Sale? Estimating Government Worker Skills with an Application to Indonesia." *Working Paper*.

Gillis, Malcolm. 1985. "Micro and Macroeconomics of Tax Reform: Indonesia." *Journal of Development Economics* 19 (3): 221–54.

Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez. 2017. "Capital Allocation and Productivity in South Europe." *The Quarterly Journal of Economics* 132 (4): 1915–67.

Hansen, Lars Peter, and Kenneth J. Singleton. 1982. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica: Journal of the Econometric Society*, 1269–86. http://www.jstor.org/stable/1911873.

Hastie, Trevor J. 2017. "Generalized Additive Models." In *Statistical Models in S*, 249–307. Routledge.

Hausmann, Ricardo, Lant Pritchett, and Dani Rodrik. 2005. "Growth Accelerations." *Journal of Economic Growth* 10: 303–29.

Hicks, Joan Hamory, Marieke Kleemans, Nicholas Y. Li, and Edward Miguel. 2017. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." National Bureau of Economic Research.

Hill, Hal. 2000. *The Indonesian Economy.* Cambridge University Press.

Hopenhayn, Hugo. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, 1127–50.

———. 2014. "Firms, Misallocation, and Aggregate Productivity: A Review." *Annual Review of Economics* 6 (1): 735–70.

Hopenhayn, Hugo, and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy* 101 (5): 915–38.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein. 2016. "Quantifying the Sources of Firm Heterogeneity." *The Quarterly Journal of Economics* 131 (3): 1291–1364.

Hotz, V Joseph, and Robert A Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60 (3): 497–529.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *The Quarterly Journal of Economics* 124 (4): 1403–48.

———. 2014. "The Life Cycle of Plants in India and Mexico." *The Quarterly Journal of Economics* 129 (3): 1035–84.

Hsieh, Chang-Tai, and Benjamin A. Olken. 2014. "The Missing" Missing Middle"." *Journal of Economic Perspectives* 28 (3): 89–108.

Humlum, Anders. 2022. "Robot Adoption and Labor Market Dynamics."

Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin. 2019. "Demographic Origins of the Startup Deficit." National Bureau of Economic Research.

Kleinman, Benny, Ernest Liu, and Stephen J. Redding. 2023. "Dynamic Spatial General Equilibrium." *Econometrica* 91 (2): 385–424.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy* 112 (5): 986–1018. https://doi.org/10.1086/422563.

Kruse, Hagen, Emmanuel Mensah, Kunal Sen, and Gaaitzen de Vries. 2023. "A Manufacturing (Re) Naissance? Industrialization in the Developing World." *IMF Economic Review* 71

(2): 439–73.

Lanteri, Andrea, Pamela Medina, and Eugene Tan. 2023. “Capital-Reallocation Frictions and Trade Shocks.” *American Economic Journal: Macroeconomics* 15 (2): 190–228.

Lent, G., and P. Ojha. 1969. “Sales Taxes in Countries of the Far East.” *IMF Staff Papers* 16.

Liang, James, Hui Wang, and Edward P. Lazear. 2018. “Demographics and Entrepreneurship.” *Journal of Political Economy* 126 (S1): S140–96.

Luttmer, Erzo GJ. 2007. “Selection, Growth, and the Size Distribution of Firms.” *The Quarterly Journal of Economics* 122 (3): 1103–44.

McCaig, Brian, and Nina Pavcnik. 2021. “Entry and Exit of Informal Firms and Development.” *IMF Economic Review* 69: 540–75.

McKenzie, David, and Anna Luisa Paffhausen. 2019. “Small Firm Death in Developing Countries.” *Review of Economics and Statistics* 101 (4): 645–57.

Moll, Benjamin. 2014. “Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?” *American Economic Review* 104 (10): 3186–3221.

Olley, G. Steven, and Ariel Pakes. 1996. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica*, 1263–97.

Ottonello, Pablo, and Thomas Winberry. 2023. “Investment, Innovation, and Financial Frictions.”

Perla, Jesse, Christopher Tonetti, and Michael E. Waugh. 2021. “Equilibrium Technology Diffusion, Trade, and Growth.” *American Economic Review* 111 (1): 73–128.

Peters, Michael. 2020. “Heterogeneous Markups, Growth, and Endogenous Misallocation.” *Econometrica* 88 (5): 2037–73. https://doi.org/10.3982/ECTA15565.

Poschke, Markus. 2013. “‘Entrepreneurs Out of Necessity’: A Snapshot.” *Applied Economics Letters* 20 (7): 658–63.

———. 2018. “The Firm Size Distribution Across Countries and Skill-Biased Change in Entrepreneurial Technology.” *American Economic Journal: Macroeconomics* 10 (3): 1–41.

Restuccia, Diego, and Pedro Bento. 2015. “Misallocation, Establishment Size, and Productivity.” In *2015 Meeting Papers*. 1391. Society for Economic Dynamics.

Restuccia, Diego, and Richard Rogerson. 2017. “The Causes and Costs of Misallocation.” *Journal of Economic Perspectives* 31 (3): 151–74.

Romer, Paul M. 1990. “Endogenous Technological Change.” *Journal of Political Economy* 98 (5, Part 2): S71–102.

Ruggieri, Alessandro. 2022. “Trade and Labour Market Institutions: A Tale of Two Liberalizations.” University of Nottingham, GEP.

Ruiz-García, J. C. 2019. “Financial Frictions, Firm Dynamics and the Aggregate Economy: Insights from Richer Productivity Processes.” Working Paper.

Schoar, Antoinette. 2010. “The Divide Between Subsistence and Transformational Entrepreneurship.” *Innovation Policy and the Economy* 10 (1): 57–81.

Timmer, Marcel, Gaaitzen J. de Vries, and Klaas De Vries. 2015. “Patterns of Structural Change in Developing Countries.” In *Routledge Handbook of Industry and Development*, 79–97. Routledge.

Ulyssea, Gabriel. 2020. “Informality: Causes and Consequences for Development.” *Annual Review of Economics* 12 (1): 525–46.

Van Biesebroeck, Johannes. 2005. “Firm Size Matters: Growth and Productivity Growth in African Manufacturing.” *Economic Development and Cultural Change* 53 (3): 545–83.

World Bank. 2023. *Poverty and Inequality Platform (Version 20230328_2017) [Data Set]. World Bank Group. Www.pip.worldbank.org. Accessed the 15th September, 2023.*

A Data and Empirical Evidence

A.1 Data cleaning details

In the following subsection, we describe in detail the data cleaning steps we take to ensure that the data is consistent over time and that results are not driven by different forms of measurement error. The main data cleaning steps relate to cleaning plant-level labor (wage bill and workers) and output (value added) over time. We further discuss how we deal with "dynamic" outliers (e.g. unrealistic within-plant jumps in value added or the wage bill) and observations with extreme labor shares that pose numerous problems for the estimation and computation of the model. Besides these mentioned cleaning steps, we also drop a few clear outliers, such as when the plant ID is misreported or missing or when magnitudes of multiple reported variables are impossible. (Give details on final cleaning: How does raw data differ from cleaned data). At last, we report details on how we clean the capital series, industry codes and measure plant entry and exit.

A.1.1 Cleaning labor and the labor wage bill

The manufacturing census consistently reports a plant's total number of workers (including paid and unpaid) as well as separately the number of paid versus unpaid workers and the number of production and non-production (including managerial) workers. The main cleaning step we apply to ensure consistency over time is to drop all plants with less than 20 total workers (which is enforced by BPS starting in 1990, but not before), drop plants that report zero paid workers or that report more paid workers than total workers. This step drops slightly less than 2% of plant-year observations with dropped observations concentrated before 1990. We also identify a bunching at 99 workers in the years 2013-2015 (roughly 3-4% of plants), which we interpret as true bunching driven by actual policy changes and thus do not correct.[21]

For the structural model, we build on a plant's reported total wage bill. This variable is the sum of the total wage bill for other workers and for production workers. In principle, it includes all payments to labor, including in-kind transfers, overtime pay, bonuses and social contributions (e.g. pension and accident allowances). Since the survey asks about current workers and doesn't separately ask about severance pay, we treat the reported wage bill as excluding severance pay.

We take three main cleaning steps for the reported wage bill. First, we correct systematic

[21] For example, Indonesia introduced an occupational safety and health regulation in 2012 that was targeted at manufacturing firms and mandates any workplace with more than 100 workers to implement additional work safety measures.

misreporting for the year 2011. Looking at the evolution of the distribution of the total wage bill across years, we find that 2011 is the only clear outlier with the only bimodal wage distribution across all years. In 2011, the bottom 20% of observations show exactly the same per worker wage at an unrealistically low level that is below the bottom 1% of observations in 2010 and 2012 and roughly 50 times lower than the average minimum wage in 2011. Given that these are well-defined misentries, the remainder of the distribution is well-behaved and we observe most plants before and after 2011, we opt for imputing misentries using within-plant averages across 2010 to 2012, enforcing linear growth for 2011.

Second, we correct for misreporting in the wages for non-production workers. Non-production workers account for roughly 16.5% of overall employment in our data. However, 17% of plant-year observations reportedly employ zero non-production workers. This also means that plants may not always report all managerial workers, which is likely if the managerial staff partly owns the establishment (and is thus not formally employed). To the extent that all payments to managerial workers should be counted as labor costs, the Indonesian data may significantly underestimate labor costs. We cannot correct for this form of underreporting. However, we can correct for the following: About 10% more plant-year observations report zero wage payments to non-production workers than plant-year observations reporting zero non-production workers. That is, there are plants that report employing non-production workers, but paying them no wages. This could be in part due to some plants reporting managerial staff who receive remuneration other than wages or due to plants simply not reporting non-production worker wages. In any case, we think it is better to impute wages here, using plants' reported wages for production workers and the average year-specific pay gap across production and non-production workers for plants that report both wages. We find that wage premia for non-production workers were around 90% in 1975 and declined to around 20% in 2015. In the end, the overall importance of this correction is small, because the correction only applies to a small number of observations.

In the third and last cleaning step, we correct the total reported wage bill across periods where the exact questions and components of the total wage bill changed. While plants were asked consistently to report total payments including cash and in-kind wages, pensions and other social contributions between 1975-1995, survey questions changed most notably between 2001-2003, 2004-2010 and 2011-2015. Looking at changes in the distribution of reported total wage bills, the 2004-2010 period appears to be the most problematic period in which total wages are systematically underreported vis-a-vis the other periods.

We correct for changes in the measurement of the total wage bill over the period 2004-2010 by exploiting within-plant changes in reported wage bills across changes in the measurement period (from 2003 to 2004 and from 2010 to 2011) and further utilizing information across different types of workers and the reporting of the number of workers by type. Plants i report x_{it}^{j}, the total reported wage bill for worker type j in period t. Specifically, we assume that: $x^j_{itm} = w_t h^j_{it} \epsilon^j_{it} \tau^j_m$, where w_t is the wage (in line with our model), h^j_{it} are plant-worker-type-specific efficiency units of labor, ϵ^j_{it} captures idiosyncratic measurement error and τ^j_m is systematic underreporting that is constant within a worker-type and within the measurement period m where the same questions to elicit the worker-type-specific wage bill are asked. We assume that $\tau^j_m \in (0,1]$ for the period 2004-2010 and unity otherwise.

Our approach separately identifies wages from the measurement error τ^j_m across time and thus allows to correct for an important part of measurement error that leads to the underestimation of the total wage bill and a time inconsistent measure of the wage bill. For separate identification, we assume that over two consecutive periods, plants use the same average human capital within types of worker: $\frac{h^j_{it}}{l^j_{it}} = \frac{h^j_{it+1}}{l^j_{it+1}} = \alpha^j_i$. This allows some plants to specialize on high productive production workers or other plants to specialize on low productive managerial staff, but restricts changes in the average human capital within plant-worker-types. As long as the plant-specific number of workers by type are reported either without measurement error or with constant plant-worker-type measurement error, the assumption allows to identify:

$$\mathbb{E}_i \frac{\frac{x^j_{i,t+1,m}}{l^j_{i,t+1,m}}}{\frac{x^j_{i,t,m}}{l^j_{i,t,m}}} \equiv \mathbb{E}_i \frac{\tilde{x}^j_{i,t+1,m}}{\tilde{x}^j_{i,t,m}} = \frac{w_{t+1}}{w_t} \mathbb{E}_i \frac{\epsilon^j_{it+1}}{\epsilon^j_{it}}$$

Within measurement periods m, one can show that under realistic magnitudes for the measurement error, the following holds:[22]

$$\mathbb{E}_i \frac{\tilde{x}^j_{i,t+1,m}}{\tilde{x}^j_{i,t,m}} \approx \frac{w_{t+1}}{w_t}$$

Across measurement periods, separate identification of the change in measurement and the wage is impossible without further assumptions. To see this, write:

$$\mathbb{E}_i \frac{\tilde{x}^j_{i,t+1,m'}}{\tilde{x}^j_{i,t,m}} = \frac{w_{t+1}}{w_t} \mathbb{E}_i \frac{\epsilon^j_{it+1} \tau^j_{m'}}{\epsilon^j_{it} \tau^j_m} \approx \frac{w_{t+1}}{w_t} \frac{\mathbb{E}_i \epsilon^j_{it+1} \tau^j_{m'}}{\mathbb{E}_i \epsilon^j_{it} \tau^j_m} = \frac{w_{t+1}}{w_t} \frac{\mathbb{E}_i \tau^j_{m'}}{\mathbb{E}_i \tau^j_m}$$

where we have made use of the same approximation as above. Even then, changes in the measurement error across measurement periods cannot be separately identified from wage

[22]Specifically, a first-order Taylor series approximation around the mean of the measurement errors gives: $\frac{w_{t+1}}{w_t} \mathbb{E}_i \frac{\epsilon^j_{it+1}}{\epsilon^j_{it}} \approx \frac{w_{t+1}}{w_t} \frac{\mathbb{E}_i \epsilon^j_{it+1}}{\mathbb{E}_i \epsilon^j_{it}}$. With a second-order Taylor series approximation, we get: $\mathbb{E}_i \frac{\epsilon^j_{it+1}}{\epsilon^j_{it}} \approx \frac{\mathbb{E}_i \epsilon^j_{it+1}}{\mathbb{E}_i \epsilon^j_{it}} \left[1 - \frac{Cov_i(\epsilon^j_{it+1}, \epsilon^j_{it})}{\mathbb{E}_i \epsilon^j_{it+1} \mathbb{E}_i \epsilon^j_{it}} + \frac{Var_i(\epsilon^j_{it})}{\left[\mathbb{E}_i \epsilon^j_{it}\right]^2} \right]$. Plugging in realistic measurement error, the 2nd-order correction is very small. For example, if measurement error within plants is positively correlated (which is the likely case), then the two correction terms go in opposite directions. Furthermore, both the covariance term and the variance term are close to zero for reasonable magnitudes of measurement error.

changes. To solve this issue, we interpolate wages from wage growth in the previous period and the next period (for which measurement does not change), assuming that wages grow smoothly over time.

In our case, we set $\tau_m^j = 1$ for all measurement periods except the period 2004-2011. To identify $\tau_{m'}^j$ for 2004-2011, we are now actually over-identified, because we can identify the measurement error from variation between 2003-2004 or from 2010-2011. We choose to use 2003-2004 because 2010-2011 featured a change in the minimum wage in Indonesia, which partly explains a large increase in plants' total wage bills and we do not know how to separate this change from a change in the measurement. Following this approach, we find that $\tau_{m'}^j \approx 0.94$, similar when restricting to production workers only or when looking at all workers. We enforce the production worker correction across all plants for the measurement period 2004-2011 (which means that their wage bills get upward corrected by $1/\tau_{m'}^j$, a correction of roughly 6%).

A.1.2 Cleaning output / value-added

Throughout the paper, we use a consistent definition of value-added output. This definition coincides with how the Indonesian statistical agency (BPS) constructed value-added output for some, but not all years. Specifically,

$$\text{Value-added}_{it} \equiv \text{Gross income}_{it} - \text{Intermediates}_{it}$$

where Gross income$_{it}$ $\equiv$ Gross sales$_{it}$ + electricity sales$_{it}$ + revenue from industrial services$_{it}$ + other income$_{it}$ + Δvalue of semi-finished products$_{it}$ and Intermediates$_{it}$ $\equiv$ Raw materials$_{it}$ + Total fuel/electricity expenditures$_{it}$ + Other expenses$_{it}$. All inputs in the accounting identities are reported in their current values of Rupiah, which we deflate to 2010 constant Rupiah based on the aggregate CPI. We start by dropping observations with missing or negative gross income, which are less than 0.5% of observations.

Next, we construct a time-consistent measure of intermediates. The main issue is that intermediates are likely underreported since the survey only asks for specific categories of expenditures and intermediate expenditures have likely become more complex over time, leading plants to underreport parts of their expenditures. This leads to an overestimate of value-added output and an underestimate of capital and labor cost shares at the plant-level. We correct intermediate expenditures in two steps.

In the first step, we look at one main expenditure category of intermediate inputs for which we know that time inconsistency is an issue. Specifically, Other expenses$_{it}$ are reported inconsistently over time because not all components of other expenses are enumerated in every year. In the following, we describe the components of Other expenses$_{it}$ and how we impute them consistently across plants over time. In the years with the most detailed survey questions, Other expenses$_{it}$ (X_{it} in short) are the sum of three components (indexed by c): (1) expenses for other goods (consisting of packaging, spare parts and stationary), (2) manufacturing services, repair and maintenance, and (3) remaining other expenses (with detailed subcomponents for some years). We improve the measure of intermediate expenditures by imputing these three subcomponents in cases where they are missing. Similar to the components of the labor wage bill, we deal with underreporting of other expenses by exploiting within-plant differences in reporting around years with changes in survey questions. We separately impute missing fractions of each of the three components of other expenses, using further information on subcategories j within components c, bringing all series to the most complete level of reporting in the years 2006 and 1996/1997.

Specifically, we assume that $\forall c, j : \ X_{icjt} = \alpha_{icj} Y_{it}$. That is, any other expense category (or subcategory) is a plant-subcategory-specific fraction of gross income Y_{it}. Since expenditures for specific subcategories are systematically missing in some years, but gross income Y_{it} is reported for all years, we impute complete missing subcategory expenditures as follows: For plants that we observe across different measurement periods, we impute their expenditure shares from average within-plant expenditure shares around the time of missing. For example, expenses for other goods are missing between 1998 and 2005, which we impute using the plant average of plant-specific expenditure shares in 1997 and 2006 together with plant-year-specific gross income Y_{it}. This ensures within-plant consistency and allows for plant-category-year-specific variation in expenditures. For plants for which we do not observe expenditures in other years, we use the aggregate category-specific expenditure share around the time of missing.

On top of this, we correct reported expenditures for the remaining other expenditures for the period 1975-1984 in which the reported series is clearly underreported in comparison to post 1985. For this correction, we again proceed separately for plants that we observe across measurement periods and for plants that we do not, using either the average within-plant difference in reported ratios or the ratio of aggregate expenditure shares across the two measurement periods as correction factors.

Overall, this first step of cleaning intermediate inputs ensures more time consistency, but does not have a sizable effect on overall intermediate expenditures. In the second step of correcting intermediate expenditures, we deal with the sizable remaining decline in the intermediate expenditure share across plants over time. For example, the aggregate intermediate expenditure share declines by more than 10 percentage points from roughly 0.65 in 1980 to 0.525 in 2015 (mostly driven by a decline in the raw material input share). We expect that a major part of this decline is in fact measurement error. One simple reason could be that plants use more processed intermediate inputs, which they do not fully report as "raw materials." To distinguish this driver from industrial composition effects (e.g. industries relying on intermediate inputs declining in relative importance over time), we construct the following correction: We regress $log(\phi_{ijt}/(1-\phi_{ijt})) = \alpha_j + \alpha_t + \epsilon_{ijt}$ where ϕ_{ijt} is the intermediate expenditure share of plant i in 5-digit industry j at time t. We use the log odds ratio to ensure that any correction we implement gives expenditure shares that are bounded between zero and one. α_j and α_t capture industry and time fixed effects. We interpret α_t as our time-varying bias term, using α_{1975} as the normalization factor (for which the bias is zero). Controlling for industry fixed effects ensures that the bias terms do not capture variation in intermediate expenditure shares from changes in the industrial composition. Corrected intermediate input shares are then given by $\tilde{\phi}_{ijt} = exp(\alpha_j + \alpha_{1975} + \epsilon_{ijt})/(1 + exp(\alpha_j + \alpha_{1975} + \epsilon_{ijt}))$. The correction maintains plant-level variation in intermediate expenditure shares and delivers both within-plant and aggregate time consistency. To ensure that the regression is well-specified, we initially drop observations with non-positive intermediate inputs or value added as well as observations with missing value added. This drops less than 2% of observations. After the correction, we recompute intermediate expenditures and value added.

A.1.3 Identifying problematic outliers: jumps and extreme labor shares

The last important cleaning step we take, is to identify problematic outliers that are likely misentries and would have an outsized role on the model estimation and inference.

We start with "dynamic" outliers by which we refer observations that are outliers within the time series of an individual plant. We treat a plant-year observation as a dynamic outlier if the total wage bill or value added output series (X_{it} in the following) of an individual plant sees a sizable one-time jump after which it reverts directly back. We consider two different measures: the year-to-year within-plant change $\frac{X_{it}}{X_{it-1}}$ and the year-to-year within-plant quantile difference $q_t(X_{it}) - q_{t-1}(X_{it-1})$. For both measures, we first identify a potential outlier if any of the two measures is below the 10th or above the 90th within-year percentile for the respective measure. For example, we classify the following observation a potential outlier: Between 1993 and 1994, a plant's quantile of value added output changed by more than the 90th percentile of this year's distribution of value added output quantile changes. We classify any potential outlier as an actual outlier only if the following year is also identified as a potential outlier whose change goes in the opposite direction.[23] This ensures to identify jumps, while the initial and final level do not need to coincide, allowing for plant-, time- and variable-specific drifts. Big one time changes are explicitly not counted as outliers, treating them as true shocks. We drop any dynamic outlier that we detect through this procedure. Note that this procedure also identifies observations that change back and forth multiple times in a row as outliers as long as their changes are very large. In total, this procedure drops almost 10% of plant-year observations and roughly 10% of total reported value added.

[23] We do this separately by measure so the identification of an outlier is based on within-measure changes only.

The second and last cleaning step we take is to drop observations with extreme labor shares. These are observations with reported labor shares below 5% and above 500% (roughly differing from the median labor share of 50% by a factor of 10). Extreme labor shares are likely a combination of overreported value added and underreported wage bills and these observations have a sizable impact on aggregates. They make up roughly 3% of observations, but account for 41% of total reported value added. While we think that many of these plant-year observations with large value added are correctly classified as being "granular" in their importance for output (e.g. many of these plants consistently report being large over time), their exact value added output and wage bills are likely mismeasured.

A.1.4 Cleaning capital

For cleaning plants' self-reported capital stock, we draw on the cleaning steps in Cali, Le Moglie, and Presidente (2021), which is the most thorough attempt at cleaning the Indonesian manufacturing plant capital series that we are aware of. The cleaning steps draw in part on the perpetual inventory method (PIM). Details can be found in Cali, Le Moglie, and Presidente (2021).

A.1.5 Cleaning industry codes

Industry classifications changed over time, starting with ISIC 2 in 1975 and moving to ISIC 3, ISIC 3.1 and ISIC 4 by the end of our data period. For harmonization, we start by fixing a plant's first reported 5-digit industry (the most disaggregated level reported). While plants may reasonably change industries over time, we opt for fixing industries to have a time-consistent plant-level measure of industry. We then build backward correspondences at the 2-, 3-, 4- and 5-digit industries respectively using within-plant changes in industry classifications across changes in classification systems (that is, correspondences that map from later year classifications to earlier ones). For plants that enter later, we enforce these within-plant correspondences. In the case of one-to-many mappings (e.g. the same industry code in ISIC 3.1 maps to different codes in ISIC 3), we enforce the most common one. Note that this only matters for plants that are not observed previously. In the case of no linking (e.g. a plant enters in ISIC 4 with a code that has no observed backward linkage), we check codes manually and use official crosswalks. In the cases where we think that industries are truly "new," we simply create a new industry code. In the end, we are left with 140 unique 5-digit industry codes, 120 of which are initial ISIC 2 codes and 20 codes are new industries. At the 2-digit level, we have 9 different industries.

A.1.6 Plant entry and exit

In principle, classifying entry and exit should be straightforward: whenever a plant with a new panel identifier enters our panel, we would record this as plant entry and whenever that plant reports in t but does not report in $t + 1$, we would record the plant as having exited at the end of t. In practice, this classification would inflate plant entry and exit because occasional non-reporting is common. This is because of actual non-reporting and – as described above – because we explicitly drop plant-year observations with misreported entries. We thus only classify a plant as having exited if we do not observe reporting by the plant at any future time period. Similarly for entry, we only count the plant as entering if it is the first time the plant identifier has entered the panel. This difference is quantitatively important: the unconditional exit rate drops almost by half from around 14% to 7.9% if we follow our classification. As we discuss in the main text, 7.9% is close to other exit rates reported for India, Mexico and the US (e.g. Hsieh & Klenow 2014).

Where and how does this matter and could this new classification bias results differentially over time? Throughout the analyses, we mostly draw on within-plant changes that are robust to compositional changes due to entry and exit. In the structural model, we explicitly continue modeling plants that are non-reporting but non-exiting, correcting for (some forms of) differential non-reporting over the state space. Bias does arise from non-reporting or misreporting that is correlated with the state space for certain estimation steps. For example, estimating conditional exit probabilities clearly suffers from bias if measured exit probabilities are biased over the state space. Two potential issues may be particularly important in our case: non-reporting due to the cutoff of 20 workers and that our measure of exit may inflate exit towards the end of our data because we cannot distinguish permanent exit from temporary non-reporting. We think that both issues likely introduce biases that are small in magnitude.

As for the cutoff of 20 workers, the issue would be particularly problematic if plants regularly moved back and forth over the threshold or if plants with more than 20 workers moved permanently below the threshold, which we would wrongly classify as plant exit. We do not think that these are important issues in the Indonesian data. For example, few plants shrink and as we show in the main text, plants with 20 workers become relatively less important vis-a-vis larger plants over time. Also, given that pre-1990, plants often continue reporting even if they move below 20 workers, we find that movements around the threshold are rare. As for the classification of exit towards the end of the sample, we note that non-reporting actually seems to decline over time and 2015 (the last year of the data) is a census year in which enumeration is most complete.

A.2 Further main descriptives

Figure A.1 reports the evolution of the share of employment and value-added output that is captured by the Indonesian manufacturing plant census (1975-2015) in comparison to aggregate manufacturing value-added output and employment as reported in the GGDC 10-sector database (1975-2012) and the Economic Transformation Database (1990-2018).[24] The series tend to increase over time, which captures the fact that medium- to large-size manufacturing plants become more important over time. However, there is important variation in the shares over time that is not well-captured by a simple secular increase in the importance of medium- to large-sized manufacturing plants. In the model, we allow for the time-varying importance of the rest-of-the-economy, and differentially so for output and employment, including the part of manufacturing that our data misses. Note that for output, the increase is much stronger if we do not clean the value-added series, because a few plant entries can have an outsized effect on total value-added (e.g. for only a minimally cleaned output series, we found that the output share of our manufacturing data can increase to up to 80% by 2015, entirely driven by a few plant entries).

In the main text, we also report that our manufacturing panel misses 99% of manufacturing plants in Indonesia. This is based on information on a random five percent sample of all manufacturing establishments from the Indonesian Economic Census in 2006 reported in Hsieh and Olken (2014). We also verify that our micro-data is consistent with capturing all manufacturing plants with more than 20 workers. For example, based on the 2006 census sample (as reported in Hsieh and Olken (2014)), manufacturing plants with more than 50 workers should capture 34% of total manufacturing employment, while this figure is 32% based on employment in our micro-data (29.5% after cleaning) and taking the aggregate sectoral employment from the GGDC 10-sector database as denominator. Given that the manufacturing plant panel includes new plants based on the Economic Census, coverage is more complete after Economic Census years.

Next, Figure A.2 reports the evolution of aggregates in the Indonesian economy, showing the series in logs to better visualize how growth rates changed over time.

For completeness, Figure A.3 reports the full year-to-year evolution of employment shares for different plant sizes.

We further report estimated Pareto tail coefficients for the manufacturing data in Figure A.4. We follow Chen (2022) in constructing two simple, but alternative measures of the Pareto tail.

[24] We merge the latter two series consistently over time by enforcing the more recent vintage and aligning all series before 1990 to be consistent with the evolution after 1990. Specifically, for each variable, we take the ratio of the GGDC10 and ETD series in 1990 to be the amount that the 1990 GGDC series needs to be adjusted by. We then similarly correct each year from 1975 to 1990 by a correction factor that equals unity in 1975 (no correction in 1975), is equal to the full correction in 1990 and is taken from an equal-spaced, smoothed series in the years between.

Figure A.1: Representativeness of manufacturing panel over time

Notes: Evolution of the employment and value-added output share captured by the Indonesian manufacturing plant census (1975-2015) in comparison to aggregate manufacturing value-added output and employment as reported in the GGDC 10-sector database (1975-2012) and the Economic Transformation Database (1990-2018). Straight lines report (unweighted) time averages. The black dotted line highlights the Asian Financial Crisis.

Let $F(x,t)$ be the CDF of the underlying distribution of plant employment, and f the density function. Then $\tilde{F}(x,t) \equiv 1 - F(x,t)$ denotes the fraction of plants with size greater than x, and $\tilde{F}^{emp}(x,t) \equiv \int_x^\infty y dF(y,t)$ the total employment in plants with size greater than x. In addition, let T_L be the employment size threshold for large plants and T_S for small plants. Assuming that $F(x,t)$ follows a Pareto distribution with shape parameter k_t, we have:

$$k_t = 1 - log\frac{\tilde{F}^{emp}(T_L)}{\tilde{F}^{emp}(T_S)}/log\frac{T_L}{T_S}$$

Alternatively,

$$k_t = -log\frac{\tilde{F}(T_L)}{\tilde{F}(T_S)}/log\frac{T_L}{T_S}$$

In both cases, Pareto tails can be computed in the absence of knowing the employment share or fraction of plants below 20 workers, because these shares cancel out. Panel A reports estimated Pareto coefficients for different thresholds T_L for the employment share measure, while Panel B reports the estimated Pareto coefficients for the same thresholds T_L but for the number of plants instead.

In accordance with the main text, we estimate Pareto tails by decade in 1975, 1985, 1995, 2005 and 2015. The overall trend in Pareto tails is consistent for different measures and different thresholds T_L: the tail of the employment distribution grows markedly thicker over

Figure A.2: Evolution of aggregate and sectoral employment and output (in logs)

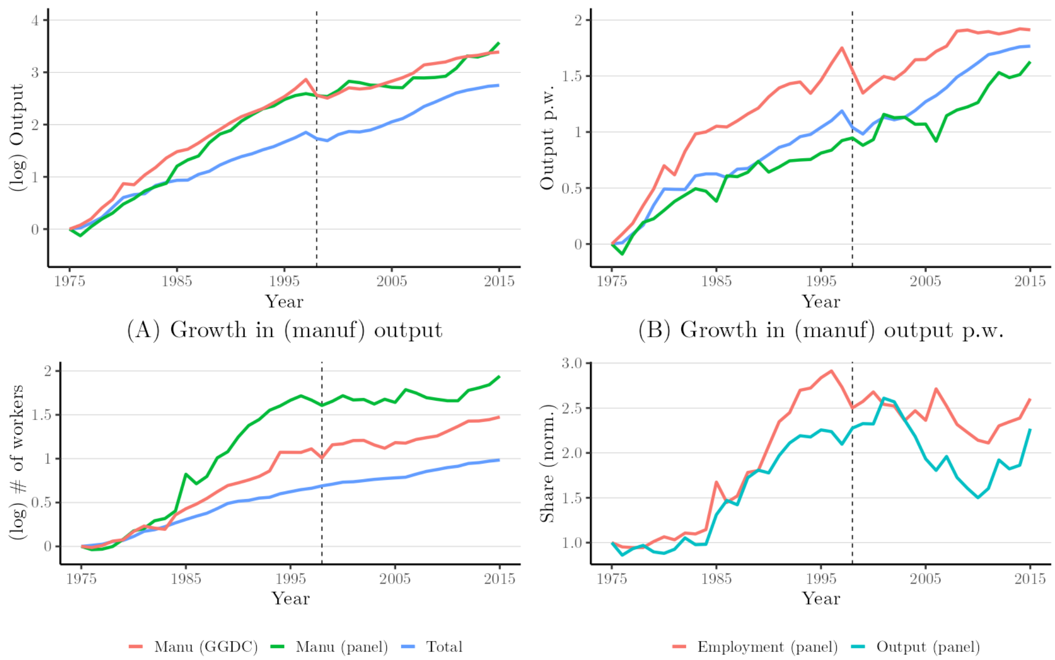

(A) Growth in (manuf) output

(B) Growth in (manuf) output p.w.

(C) Growth in number of workers

(D) Manuf employment + output shares

Notes: (Economy-wide) Total and GGDC are based on joining the GGDC 10-sector Database (1975-2012) and the Economic Transformation Database (1990-2018). Panel refers to the Indonesian manufacturing plant census (1975-2015, 20+ workers). All series are normalized by their respective value in the first year. (A) and (B) use value-added output.

time. However, there are important differences both across measures and across different thresholds, which is not in line with a common Pareto distribution in the cross-section. The quantitative implications are also very different for the two different measures, because Pareto tails below 1 imply that not even the mean of the distribution is defined.

Figure A.5 reports changes in the plant age distribution over time. Average plant age increased by roughly 40% between 1975 and 2006. While the 1975 plant distribution does feature very old plants, by far most plants are very young. In contrast, 30 years later, the plant age distribution is far more equally distributed, featuring more medium-old plants and relatively fewer very young plants.

A.3 Additional results for iterating on initial distribution

In this subsection, we provide further results on the reduced-form exercise of iterating on the discretized initial plant distribution. We start out by showing that 1975 and 1976 are good starting years, and if anything, give conservative estimates. We then show that results are very similar when accounting for entry and exit.

Figure A.3: Evolution of employment shares in large Indonesian manufacturing plants

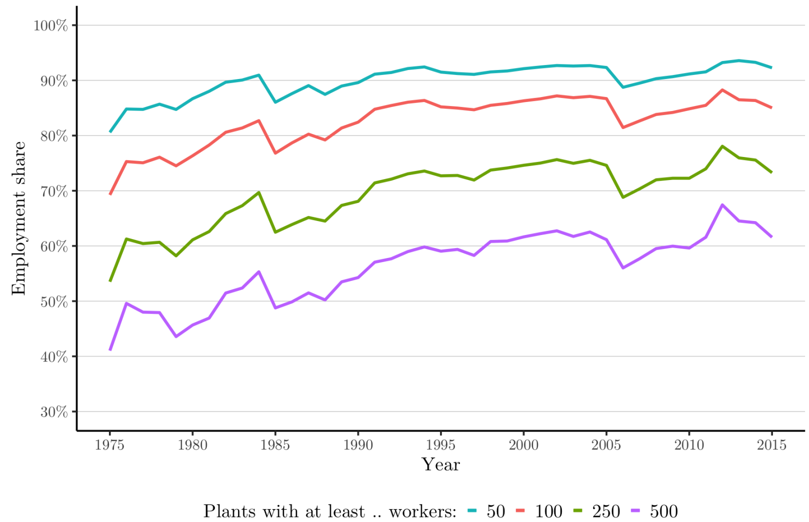

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers.

Figure A.4: Evolution of Pareto tail coefficients

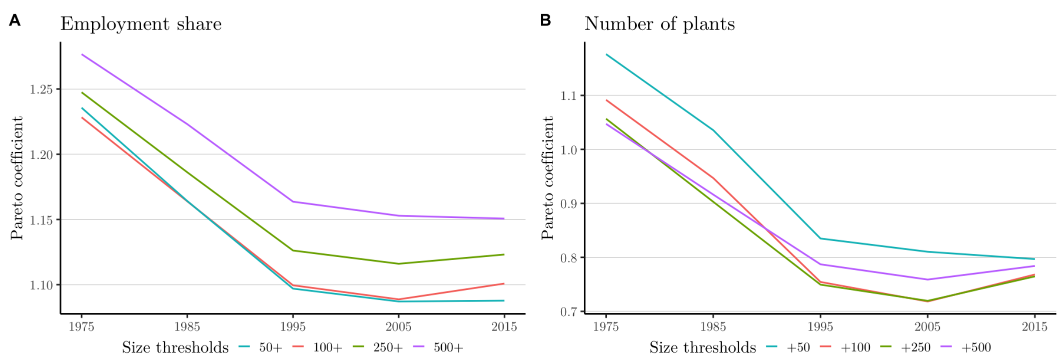

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers.

Results are similar when varying transition matrices and averaging transition matrices over multiple years. From an historical point of view, taking 1975 as the starting year and 1975 and 1976 as the initial years from which we construct the reduced-form transition matrix, is conservative, because (1) the two years fall in between the two periods of growth accelerations identified by Hausmann, Pritchett, and Rodrik (2005) for Indonesia (which are dated to 1967-1974 and 1987-1994 respectively), and (2) they are not affected by any notable labor or financial market reforms and predate the major tax reform of 1976 (see: Hill 2000). This is not to say that the 1970s were economically without important events. Oil prices rose dramatically in 1973 and inflation became a major macroeconomic issue that was followed by interest rate hikes and ceilings on commercial bank credit in 1974 (Hill 2000, 294). There were also important export-promoting trade policy reforms throughout the 1970s, but during a time in which Indonesia was still a very closed economy. Based on World Bank national

Figure A.5: Evolution of the age distribution across plants

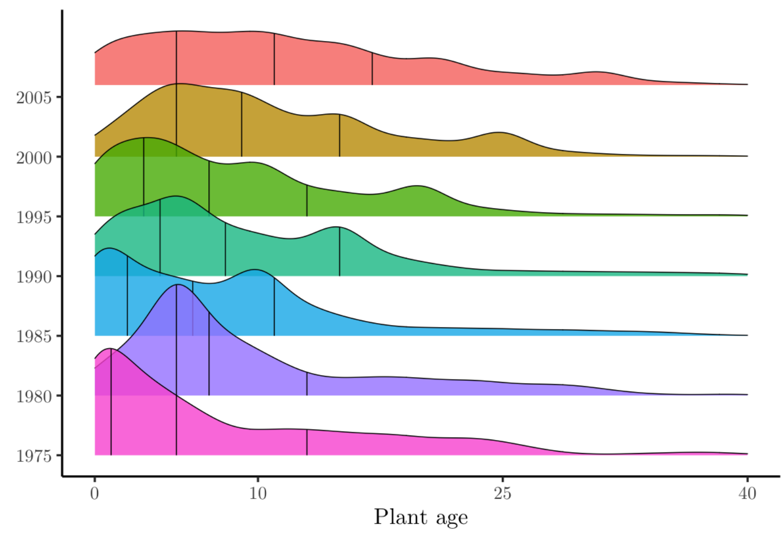

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. Only showing data for years 1975, 1985, 1995,2006. The last year is 2006, because 2006 is the last year where plant age is separately asked in the survey. After 2006, we only observe plant age for surviving plants, biasing estimates of the cross-sectional age distribution.

accounts data, exports made up around 22% of GDP in 1975 whose share actually slightly decreased from 1975 to 1976, alleviating the concern that the growth between 1975 and 1976 is purely driven by trade reforms.

Figure A.6 shows that taking transition matrices for any other starting year (e.g. 1985 as starting year for 1985-1986 transitions) gives, if anything, stronger results than the ones reported for 1975 and 1976. Most years see much more growth in the average plant size and the employment share of large plants. Importantly, all years show an eventual increase in the employment share of large plants, giving credence to the idea of a tail that slowly fills up. Furthermore, any other starting year in the 1970s would have given much stronger results. E.g. taking transitions between 1976 and 1977 would have explained 67% of the average size increase and 96% of the employment share increase over time. We also considered averaging transition matrices across multiple years and obtained very similar results.

Next, we considered two variations on the exercise to account for entry and exit. To begin with, note that entry and exit is potentially very important, especially if entering plants differ from exiting plants. Of the roughly 6,800 plants with more than 20 workers operating in 1975, less than 12% were still operating in 2015. On the other hand, as shown in Figure 2, the number of active plants increased by a factor of 4 between 1975 and 2015. This means that the vast majority of active plants in 2015 did either not exist or was not captured in the 1975 census. To capture the role of entry and exit, we amend the previous exercise by including a state-0 which captures inactive plants or potential entrants. This means that both the

Figure A.6: Reduced-form transition dynamics from initial conditions in 1975 and all year-to-year transition matrices

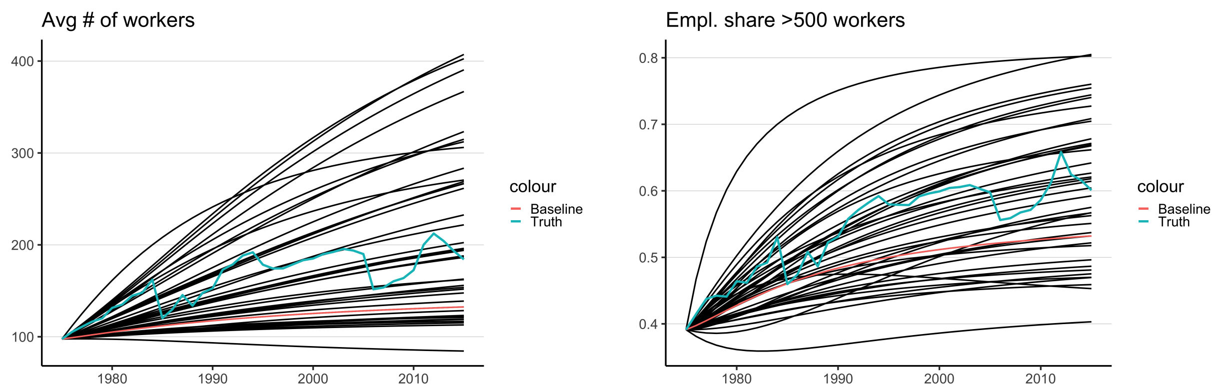

Notes: Reduced-form transition dynamics implied by initial plant size distribution in 1975, but taking transition matrices from each year-to-year pair in the data.

initial distribution is defined over an additional state-0 and the transition matrix will feature transitions into (exit) and out of state-0 (entry). To construct the new transition matrix, we can use observed entry and exit flows. Since transition matrix entries are computed as the share of flows from bin x in period t into any other bin in period $t+1$, we can readily compute transitions from an active state to an exit state. However, we cannot directly compute entries from inactivity, because the baseline is fundamentally undetermined. We do not know how many inactive or potential plants there are. This means we can also not directly compute the new initial distribution that includes the measure of plants in state-0. Since both the transition matrix and the initial distribution depend on the number of inactive plants, this number cannot be identified from observables in the first two periods alone. In theory, we can pin down the initial number of inactive plants by enforcing that the transition matrix stays constant over time and by feeding in another moment, the change in the number of plants between 1976 and 1977. However, the initial periods saw an initial decrease in the number of plants between 1975-1976 and a subsequent increase between 1976-1977. To match this pattern, we would have to enforce a negative transition matrix entry for staying inactive.

To avoid this, while giving almost indistinguishable results, we instead assume that the share of inactive plants that stay inactive is 0. This identifies the transition matrix and we then consider two additional exercises where we keep this transition matrix fixed. In the first version of the exercise with entry and exit, we simply iterate on the initial distribution and the transition matrix. This keeps the total number of plants (inactive + active) constant, while introducing interesting entry and exit dynamics that directly affect the evolution of the plant size distribution over time. Results for this exercise are given by the lines "Entry + Exit" in Figure A.7. While the long-run results are almost unchanged to the previous results, introducing entry and exit does speed up transition dynamics considerably, providing a much

Figure A.7: Reduced-form transition exercise with entry and exit

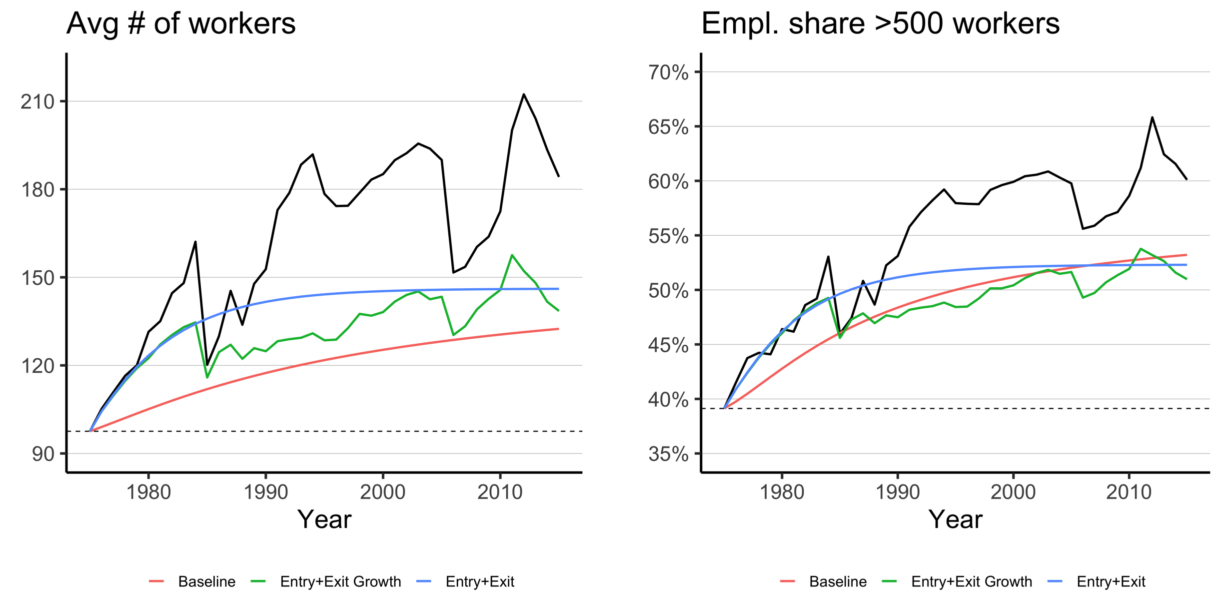

Notes:

better out-of-sample fit for the early transition period. This is driven by observed exiting plants being smaller and less productive than observed new entrants in 1976. With a positive share of inactive plants staying inactive each period would slow these predicted transition dynamics down.

In the second version of the exercise with entry and exit, we additionally vary the number of plants that enter each period. Specifically, we exactly match the increase in the number of active plants over time as shown in Figure 2, while taking information on new entrants and exits only from 1975 and 1976. In contrast to the previous exercise with entry and exit, here we do take limited information on future plant entry and thus it does not lend as well to predicting future changes in the plant size distribution. However, this exercise gives a more complete picture of the importance of entry and exit observed in the data. Results are given by the lines denoted "Entry + Exit Growth" in Figure A.7. The series again behave very similarly as before, but we can more clearly see that important year-to-year fluctuations in the real data series are driven by entry shocks. For example, the inclusion of many more plants in 1985 had important medium- to longer-run effects on the evolution of the size distribution.

A.4 Further details on exit behavior

Here, we provide evidence that exit behavior only varies little with plant productivity and does not clearly respond to aggregate shocks. Figure A.8 Panel A shows that exit rates vary

Figure A.8: Main plant exit patterns

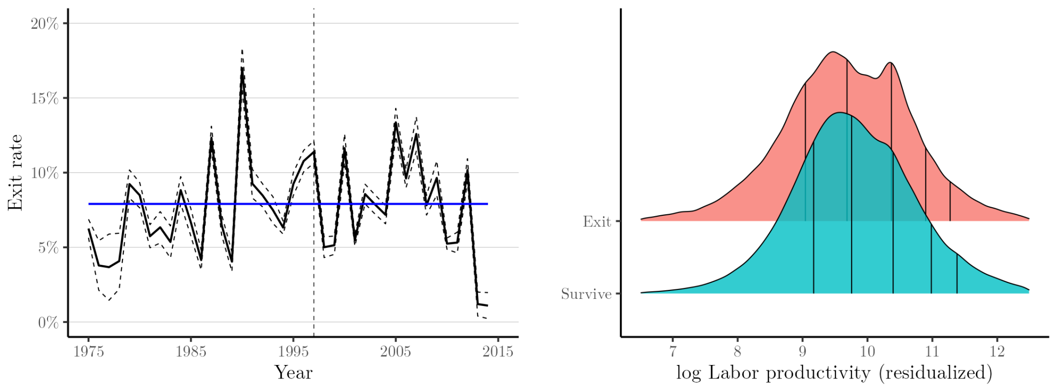

(A) Avg exit rates over time

(B) Distribution: Survive vs Exit

Notes: Panel A: Residualized by 5-digit industry fixed effects. Standard errors are two-way clustered by industry and year. Results show 95% confidence bands and blue line gives unconditional average. Panel B: Labor productivity is measured as value added per worker, residualized by 5-digit industry-year fixed effects. Vertical lines report 25th, 50th, 75th, 90th and 95th percentiles respectively.

quite strongly over time, but are not straightforwardly affected by measurable aggregate economic shocks. For example, exit rates actually decreased during the Asian Financial Crisis in 1998 & 1999. To focus only on within-industry variation, we residualize exit rates by 5-digit industry fixed effects here. Figure 5 Panel B shows productivity distributions of exiting and surviving plants using value added per worker as a simple measure of (labor) productivity and only using within-industry-time variation by residualizing the measure by detailed 5-digit industry-year fixed effects. Surviving plants are more productive on average than exiting plants, but given strong overlap in the two distributions, most plants do not exit because of their productivity. A dynamic implication of this difference is that much of plant exit is not driven by productivity so that it takes time for unproductive plants to leave the economy and productivity improvements from selective exit take time to materialize.

A.5 Details and robustness for production function estimation

We start out by proving formal identification of the production function for the different cases (static vs. dynamic capital, time variation in production functions, industry heterogeneity). We then discuss our estimation strategy and provide detailed estimation results for the case with full flexibility on the time variation in production functions but without industry heterogeneity. In the last part, we then consider the case of industry heterogeneity.

A.5.1 Identification of production function

We have the following setup. Log output by firm i at time t is given by

$$y_{it} = x_{it} + f(h_{it}, k_{it})$$

where x_{it} is productivity, h_{it} is labor input, k_{it} is capital and the price of the homogeneous production good is normalized to unity throughout. We leave $f()$ unspecified here to make clear that the identification proof is non-parametric. In the estimation and model, we assume that $f()$ is Cobb-Douglas. We also suppress industry variation here, but identification extends naturally to the case with industry variation in production functions. We start with the more general case where both capital and labor are chosen dynamically and then discuss the simpler case when capital is statically chosen.

Following the literature, we assume that in the case of dynamic capital input choices, capital is pre-determined. The input choices can then be written as non-parametric functions of the relevant state-space:

$$\begin{aligned} h_{it} &= f_h(h_{it-1}, k_{it}, x_{it}, \Omega_t) \\ k_{it} &= f_k(k_{it-1}, h_{it-1}, x_{it-1}, \Omega_{t-1}) \end{aligned}$$

where we have specified in which sense capital is pre-determined. For notational simplicity, we drop the dependence on Ω_t throughout, because identification arguments are cross-sectional. At last, productivity follows a general first-order Markov process with

$$x_{it} = f_x(x_{it-1}, u_{it}) \qquad with \quad u_{it}|x_{it-1} \sim U(0,1)$$

where u_{it} is an innovation. This representation follows the Skorohod representation of random variables and is without loss of generality (see Demirer 2022). Output, labor and capital are strictly monotonic in productivity, which imposes weak regularity conditions on the productivity process x_{it}, such that each can be inverted for productivity:

$$\frac{\partial y}{\partial x} > 0 \implies x_{it} = f_y^{-1}(h, k, y, \Omega_t) \tag{16}$$

$$\frac{\partial h}{\partial x} > 0 \implies x_{it} = f_h^{-1}(h_{-1}, h, k, \Omega_t) \tag{17}$$

$$\frac{\partial k}{\partial x} > 0 \implies x_{it-1} = f_k^{-1}(k, k_{-1}, h_{-1}, \Omega_{t-1}) \tag{18}$$

We now adapt the identification proof by Demirer (2020):

$$
\begin{aligned}
&h = f_h(h_{-1}, k, f_x(x_{-1}, u)) = f_h(h_{-1}, k, f_x(f_y^{-1}(h_{-1}, k_{-1}, y_{-1}), u)) = \tilde{f}_h(h_{-1}, k, k_{-1}, y_{-1}, u) \\
&u = F_{h|h_{-1},k,k_{-1},y_{-1}}(h|h_{-1}, k, k_{-1}, y_{-1})
\end{aligned}
$$

Intuitively, two firms with the same current capital, previous labor, previous capital and previous output, but different today's labor differ only in innovation to productivity. Using the identified u, we can then identify the production function using the control function $f_x(x_{-1}, u)$ for unobserved productivity x:

$$
y = f(h, k) + f_x(x_{-1}, u) = f(h, k) + f_x(f_y^{-1}(h_{-1}, k_{-1}, y_{-1}), u)
$$

A semi-parametric regression of y on the known function $f(h, k)$ of observables and a non-parametric term in observables/identified terms $(h_{-1}, k_{-1}, y_{-1}, u)$ identifies the output elasticities of interest.

In the case where capital is chosen statically (e.g. via a frictionless rental market), the identification approach simplifies. Specifically, dependence on k drops out in the sense that input choices are now given by:

$$
\begin{aligned}
h_{it} &= f_h(h_{it-1}, x_{it}, \Omega_t) \\
k_{it} &= f_k(h_{it-1}, x_{it}, \Omega_t) = f_k(h_{it}, x_{it}, \Omega_t)
\end{aligned}
$$

Output and labor are strictly monotonic in productivity such that:

$$
\frac{\partial y}{\partial x} > 0 \implies x_{it} = f_y^{-1}(h, y, \Omega_t) \tag{19}
$$

$$
\frac{\partial h}{\partial x} > 0 \implies x_{it} = f_h^{-1}(h_{-1}, h, \Omega_t) \tag{20}
$$

Identification is then given by:

$$
\begin{aligned}
&h = f_h(h_{-1}, f_x(x_{-1}, u)) = f_h(h_{-1}, f_x(f_y^{-1}(h_{-1}, y_{-1}), u)) = \tilde{f}_h(h_{-1}, y_{-1}, u) \\
&u = F_{h|h_{-1},y_{-1}}(h|h_{-1}, y_{-1})
\end{aligned}
$$

Now, without dependence on capital, two firms with the same previous labor and previous output, but different today's labor differ only in innovation to productivity. Using the identified u, we can then identify the production function using the control function $f_x(x_{-1}, u)$ for unobserved productivity x:

$$
y = f(h, k) + f_x(x_{-1}, u) = f(h, k) + f_x(f_y^{-1}(h_{-1}, y_{-1}), u)
$$

Figure A.9: Estimated capital and labor output elasticities for each year between 1990-2015

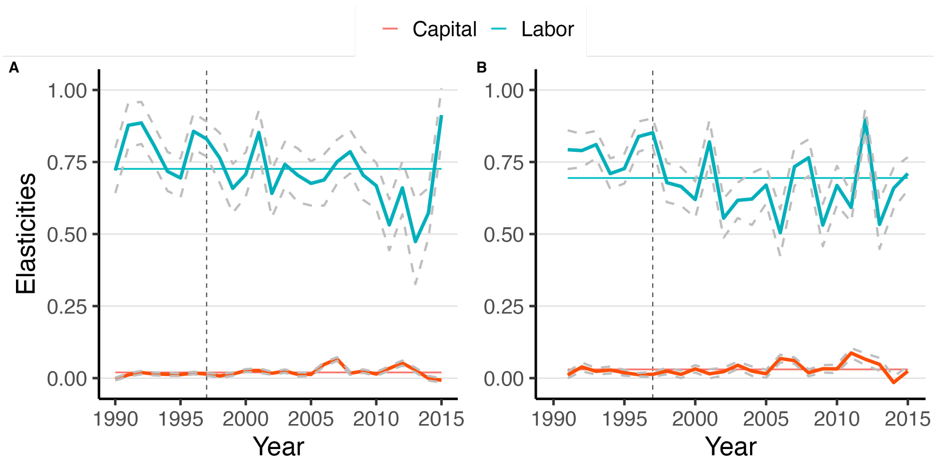

Notes: Panel A gives estimates based on a static choice of capital, Panel B gives estimates assuming capital is dynamically chosen. The latter requires to drop year 1990 in the estimation, because the estimator requires previous capital choices. Horizontal lines give average estimates over time. Grey dotted lines give 95 percent confidence bands (note that standard errors are not yet corrected for the two-stage estimation).

A semi-parametric regression of y on the known function $f(h,k)$ of observables and a non-parametric term in observables/identified terms (h_{-1}, y_{-1}, u) identifies the output elasticities of interest.

A.5.2 Production function estimation with time-variation but no industry variation

In the following, we report our production function estimation results. We start by showing results based on the sample from 1990 until 2015, which includes data on plant-level capital. We then discuss estimates from 1975-1990 and estimates with further industry heterogeneity. Importantly, (1) the estimated elasticities for labor are not biased by excluding capital, and (2) the estimated capital elasticities post-1990 are very low, meaning that any choice on how to model capital has only very small effects on productivity estimates.

Figure A.9 reports estimated capital and labor elasticities. Panel A gives estimates based on assuming that capital is statically chosen, while Panel B allows capital to be dynamically chosen. The estimates assume a common production function across manufacturing industries, but allow for fully flexible elasticities over time. Allowing for time-series variation is important, because policy functions are generally time-varying if the economic environment changes so that pooling estimates across years without allowing input choices to vary over time is model inconsistent. Elasticity estimates are remarkably stable and do not show a clear trend over time. Estimates are also very similar whether one assumes static or dynamic capital input choices. For example, the average estimated labor elasticity varies by less than 5% across the two different estimators (from 0.726 to 0.694). Furthermore, the estimated labor elasticity is close to 2/3, a common value in the literature. Note, however, that this is in a context where both the aggregate labor share in manufacturing (around 0.25) and the median labor share (around 0.54) are substantially below the estimated elasticities. Our model accounts for this systematic difference. Estimated capital elasticities are much lower than commonly estimated/used values in the literature and we discuss this point further below.

Elasticity estimates for each year are based on the following estimation steps: In the first step, we flexibly estimate the rank (taking the empirical cumulative distribution function) of labor conditional on previous labor and previous output. From this estimate, we then back out a monotonic transformation of the productivity innovation u as the difference between the observed and estimated rank. In the second step, we then estimate a log-log regression in capital and labor on output, but flexibly controlling for the different components of the control function (the estimated u and previous labor, output and capital). The dynamic and static capital estimators differ only in the variables that we condition on in each of the two estimation steps. For both estimators and both estimation steps, we draw on generalized additive models (GAMs) as a flexible and robust way to estimate semi-parametric models (Hastie 2017). We obtain very similar results when choosing flexible polynomial regressions.

Given our estimation approach, why are capital elasticity estimates so low? We think the main reason here is a standard attenuation bias in the capital elasticity estimates given substantial measurement error in observed plant-level capital. Apart from the control function term, the second step estimation is a standard linear regression in capital, so that any classical measurement error in capital will attenuate the estimated capital elasticity. Why do we suspect measurement error in the capital series? As reported in Cali, Le Moglie, and Presidente (2021), one common issue in the reported plant-level capital series is misreporting in the units, which exactly shows up as a log-additive measurement error. Apart from such unit misreporting, capital – as is well known – is also more susceptible to misreporting because it is a stock that not all plants necessarily keep track of (in contrast to cost flows such as the labor bill). Inferring changes in the stock directly from reported investments and assumptions on capital-type-specific depreciation rates (as is done in perpetual inventory methods and the capital series based on Cali, Le Moglie, and Presidente (2021) that we draw on), can mitigate some of this measurement error but unlikely all.

Besides attenuation bias from noisy capital reporting, the capital series likely also suffers from more systematic biases that complicate their use. To show this, we consider the case where capital is chosen statically in which case we can directly make use of plants' first-order conditions instead of estimating the capital elasticity from the output regression. This first-order approach does not generally suffer from attenuation bias because one can estimate capital elasticities from average or median capital shares, which are robust to log-additive measurement error. Specifically, the static capital input choice conditional on the assumed

Figure A.10: The effects of alternative estimates of capital elasticities

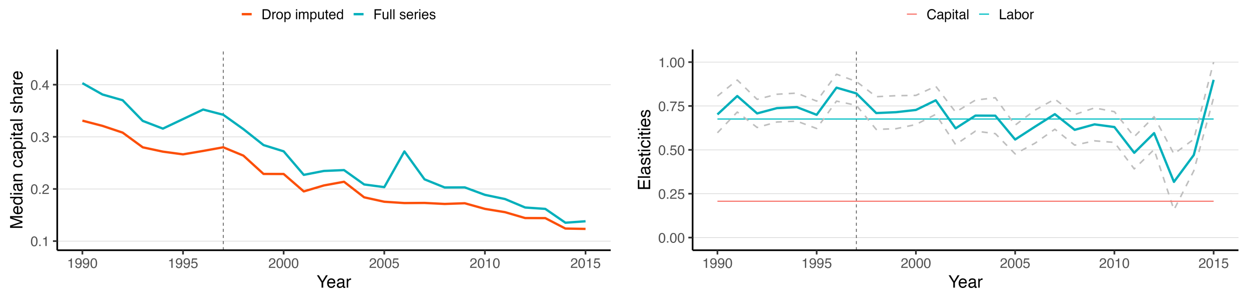

Notes: Left: Evolution of capital shares based on the capital series in Cali, Le Moglie, and Presidente (2021). For better comparability, using also their real value added series which deflates output by sector-specific prices. To construct capital shares, assume that rental rate is 14 percent (interest rate of 4 percent and depreciation of 10 percent) and value added tax is 10 percent. Right: Estimated labor elasticities when assuming that capital is statically chosen and its output elasticity is fixed at the median from Cali, Le Moglie, and Presidente (2021), which is 0.207.

production structure implies the standard condition:

$$\alpha_{jt} = \frac{(r_t + \delta_t)k_{it}}{(1 - \tau_y)p_t y_{it}}$$

where α_{jt} gives the potentially industry-time-specific capital elasticity, $(r_t + \delta)$ gives the competitive rental rate of capital, τ_y gives a value-added tax and $p_t y_{it}$ gives plant revenue.[25] Figure A.10 (left) plots changes in median capital shares over time. Capital shares more than halved between 1990 and 2015, which would also imply a halving of capital elasticities in the case of static capital choices. This seems unlikely, among others because of global capital-biased technological change, the advent of industrial robots and a running out of labor intensive industrialization in Indonesia. Also, in the case of adjustment frictions, one would have expected an increase in the capital share as the economy is catching up, not a decline. More likely, we think plants systematically underreport new capital investments and we might overestimate depreciation of existing capital.[26] Taking noisy capital estimates together with systematic misreporting, it is hard to trust the Indonesian plant-level capital series. We thus choose to stick with the low estimated capital elasticities for the baseline model and results, which implies that observed capital variation has very small effects on plant output and labor decisions. Still, to gauge how sensitive our estimates are to low capital elasticities, Figure A.10 (right) also reports estimated labor elasticities in case where we enforce the much higher capital elasticity based on the median capital share. We find that estimated labor elasticities

[25] If observed capital features log-additive measurement error ($\tilde{k}_{it} = k\varepsilon$), then, $\mathbb{E}_i \frac{(r_t+\delta_t)\tilde{k}_{it}}{(1-\tau_y)p_t y_{it}} = \alpha_{jt}$.

[26] Of course, one may also explain a strong decline in estimated revenue elasticities by changes in markups. However, the required magnitude of such markup changes also seems unrealistic. Through the lens of a standard monopolistic competition model, the difference in estimated revenue elasticities would translate to a roughly 100 percentage point increase in the markup. That is, if a product is sold at 50% above marginal cost, it would now be sold at 150% above marginal cost.

Figure A.11: Estimated output elasticities of labor for each year between 1975-2015

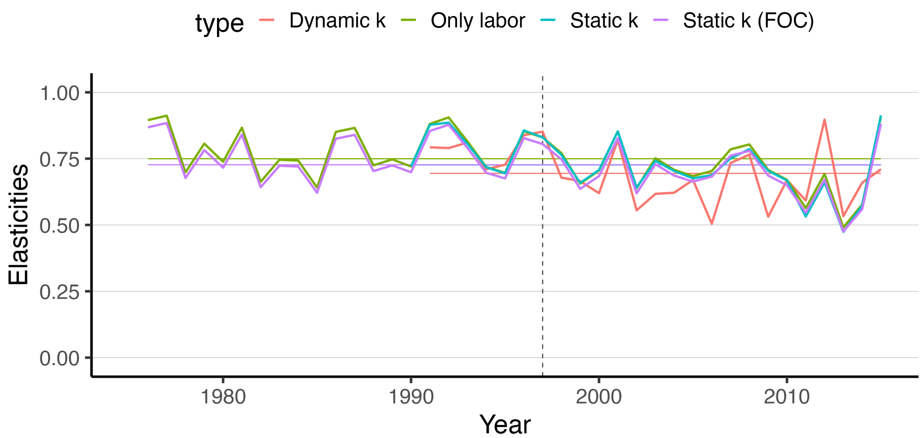

Notes: Dynamic and static k report previously estimated labor elasticities. Only labor shows estimates based on the assumption that the production function only features labor. Static k (FOC) assumes that the capital choice is static and enforces the average estimated capital elasticity from the Dynamic k estimator (which nests the static k estimator) from 1990-2015.

are almost unchanged when assuming such a higher capital elasticity.

Next, we consider production function estimation for the period 1975-1990 for which we lack data on capital. We follow two different approaches to understand whether there have been important changes in production functions over time. In the first approach, we simply assume that the production function does not feature capital and estimate only labor elasticities from 1975-2015, comparing this to the estimated labor elasticities estimated with capital data from 1990-2015. In the second approach, we assume a static choice for capital and enforce the estimated capital elasticities from before. As shown in Figure A.10, we find that for both approaches estimated labor elasticities are very similar to estimates after 1990 and that they are remarkably stable and do not show a clear trend over time. We interpret this as strong evidence that production functions did not systematically change over time.

A.5.3 Production function estimation with industry variation

We now consider industry-level variation in production functions. (Show 2 results: 1. Test for equality across industries. 2. Check how correlated productivity estimates are)

A.6 Details and robustness for aggregate technology estimates

In this section, we give a formal identification proof, estimation details for separating aggregate technology from idiosyncratic productivity and a discussion of the drivers of technology growth. To simplify the exposition, in the following we will denote the logarithm of a vector by lower letter cases. Hence, in slight deviation from the exposition in the main paper, we assume that productivity for plant i at time t is given by: $Y_{it} = Z_t exp(s_{it})$. So that the log-additive form is: $y_{it} = z_t + s_{it}(s_{it-1})$. The average of within-plant changes in log productivity is then:

$$\frac{1}{N^S_{t,t-1}} \sum_{i \in \mathcal{N}^S_{t,t-1}} \Delta y_{it} = \underbrace{z_t - z_{t-1}}_{\Delta z} + \underbrace{\frac{1}{N^S_{t,t-1}} \sum_{i \in \mathcal{N}^S_{t,t-1}} \Delta s_{it}}_{\text{Avg mean reversion of survivors}}$$

A.6.1 Identification

Proposition A.1 (Main identification result)**.** *Under the following four assumptions:*

1. *(**Common first-order stationary Markov process**) s follows the same general first-order, stationary ergodic Markov process for all i & t.*
2. *(**Selective exit**). The decision to exit after period t can flexibly depend on observables and unobservables X_{it} as well as productivity s_{it}, but may not depend on future productivity s_{it+1}. Specifically,*

$$\mathbb{P}(exit) = f(X_{it}, s_{it}, z_t) \quad with \quad \mathbb{P}_t(exit) \perp\!\!\!\perp s_{i,t+1} | s_{i,t}$$

3. *(**No complete exit over s**) $\mathbb{P}_t(exit|s_{it}) < 1 \, \forall s \in Supp(s)$*
4. *(**Connected support in s**) For each period t, there exists at least a subset of the support of s in that period which is fully contained in the support of all s in all future periods. Formally: $\forall t, \exists S_t \subset Supp(s_{it})$ for which $S_t \subset \cup_{\tau > t} Supp(s_{i\tau})$.*

the path $z_t \, \forall t$ is identified given some normalization z_τ for some $\tau \in [0, T]$ and $\max t \equiv T \rightarrow \infty$.

Proof. To already convey the idea of a suitable estimator for the time path of z_t, let us proof Proposition 1 constructively. Identification proceeds sequentially in two fundamental steps. In the first step, I show identification of the density of the stationary distribution of s, which is identified for $t \rightarrow \infty$. In the second step, the density of the stationary distribution is used to identify the path of z_t backwards by starting at some final time T. The density of the stationary distribution is key because it can be used to construct weights under which a weighted difference Δy_{it} exactly identifies Δz_t. Specifically, there exist weights ω_s such that $\sum_{i \in \mathcal{N}^S_{T+1,T}} \omega(s_{iT})(s_{iT+1} - s_{iT}) = 0$ (where $\sum_i \omega_s(s_i) = 1$). These weights recover the stationary distribution of s. Denote by $f^{SS}(s)$ the density of the stationary distribution at s and by $f_t(s)$ the density of the distribution of s at time t. Assuming that this distribution shares the support of the stationary distribution, we have:

$$\lim_{N\to\infty} \sum_{i\in\mathcal{N}_{t+1,t}^S}^{N} \frac{f^{SS}(s_{it})}{f_t(s_{it})}\Big(log(s_{it+1}) - log(s_{it})\Big) = 0$$

The weights are thus defined by $\omega_s(s_{it}) \equiv \frac{f^{SS}(s_{it})}{f_t(s_{it})}$ and are a function of the unknown density function of the stationary distribution of s. To identify the density $f^{SS}(s)$, start with the distribution of plants at t_0 over known y_{i0}. The idea is to follow survivors (as they follow the process for s), while replacing exiting plants with plants that stay in the panel that have similar y_{it}. More formally, denote the initial set of plants by $\mathcal{N}_0$ where each plant is given a uniform weight $\tilde{\omega}_{i0} = \frac{1}{N_0}$. We are interested in updating $\mathcal{N}$. For this, pass on the weight of each surviving plant and redistribute the weight of each plant that exits to close plants around them.[27] This gives $\mathcal{N}_1$. Updating in this way allows to eventually pass on weight to plants that have entered the economy, even if they have entered in an arbitrarily selective way. As $t \to \infty$, surviving plants will eventually populate the entire support of s and this procedure gives a synthetic sample $\mathcal{N}_\infty$ with weights $\tilde{\omega}_{i\infty}(s_{i\infty})$ that directly identify the density $f^{SS}(s)$.

The second step of the proof takes the identified density $f^{SS}(s)$ and works backwards from time T. Normalizing the final value z_T, one can show that z_{T-1} solves a fixed point problem. Specifically:

$$\sum_{i\in\mathcal{N}_{T,T-1}^S} \omega_{\hat{s}_{T-1}(z_{T-1})}(y_{iT}-y_{iT-1}) = z_T - z_{T-1} + \sum_{i\in\mathcal{N}_{T,T-1}^S} \omega_{\hat{s}_{T-1}(z_{T-1})}(s_{iT} - \hat{s}_{iT-1}(z_{T-1})) = z_T - z_{T-1}$$

where the last equality holds only if the guess z_{T-1} is correct. It thus gives a nonlinear equation in z_{T-1} (since the weights and the right-hand side depend on z_{T-1}). One can iterate on this procedure to identify the path of z_t backwards. At any point in time $t < T-1$, one can also alternatively guess z_{T-1} and instead of using weights at all, estimate the bias term $\sum_{i\in\mathcal{N}_{T,T-1}^S}(s_{iT} - \hat{s}_{iT-1}(z_{T-1}))$ directly using future survivors with similar s. This alternative relaxes the assumption of a common support with the stationary distribution and instead only requires that we can build a sample with similar survivors – requiring a much weaker connected support.

A.6.2 Estimation

Estimation proceeds along the lines of the constructive identification proof. In the first step, one sequentially builds the synthetic panel with weights $\omega_s(s_{it})$ (which sum to 1 in each

[27] As $N \to \infty$ and the assumption that exiting probabilities are always strictly lower than one, there always exists a plant that is arbitrarily close to an exiting plant.

year). In principle, one can use any standard matching estimator for passing on the weight for exiting plants. We find that a Kernel matching estimator works well, because matching is only based on one variable and the Kernel estimator distributes the weight widely across multiple observations, reducing variance.[28]

One can then estimate $f^{SS}(s)$ using observed s in the last period T and constructed weights $\hat{\omega}_s(s_{iT})$. Any standard density estimator such as a Kernel density estimator works here. To reduce variance, one can also estimate $f^{SS}(s)$ on the last x periods (where x is at the discretion of the researcher). In general, for any fixed T, the bias on the estimated weights is increasing in the persistence of the process as well as in the distance of the initial distribution from the stationary distribution. That is, for large T and low persistence, one can use more periods in the end to estimate $f^{SS}(s)$.[29] Once the density is estimated, one can then proceed by sequentially estimating the path z_t. For each period t and for each guess of z_t, this means one has to estimate $f_t(s_{it}(\hat{z}_t))$. Again, any standard density estimator works here. One can then construct the weights according to: $\omega_{st}(s_{it}) \equiv \frac{f^{SS}(s_{it})}{f_t(s_{it})}$. Alternatively, one can choose not to use weights and instead directly estimate the bias from mean reversion. In that case, one can again use any kind of matching estimator to match plants in t with productivity $s(\hat{z}_t)$ to future survivors with similar s. The variance in the bias estimate reduces with the number of matched plants such that one to many matches are recommended. As before, a Kernel-based matching estimator is a natural choice here. In either the approach with weights or with an estimated bias term, one then finds z_t that solves the fixed point problem, requiring a standard root finder. We have not formally proven uniqueness of the root, but in practice, we found no issue of a multiplicity of roots. In principle, any (weighted) moment of within-plant changes in productivity that preserves scalar multiplicity can be used for the estimation. In practice, we use (weighted) median changes in productivity as the median is less susceptible to outliers. Results are similar when taking the weighted average.

A.6.3 The drivers of aggregate technology

We now discuss the potential drivers of the estimated aggregate technology growth path. First, through the lens of standard endogeneous growth models such as Romer (1990), the increase in technology growth after the year 2000 could be driven by overall human capital improvements and a better integration into global markets. However, there seems to be no sharp change in human capital improvements nor in the integration into global markets. For example, Figure A.12 shows that the share of manufacturing output that is exported stays very stable around 20% and the fraction of plants that are exporting also remained flat since the mid 1980s.

[28]Note that one can readily match based on further variables such as detailed industries to minimize the risk of model misspecification.

[29]A formal treatment of optimally solving this trade-off is beyond the scope of this paper.

Figure A.12: Evolution of key exporter statistics

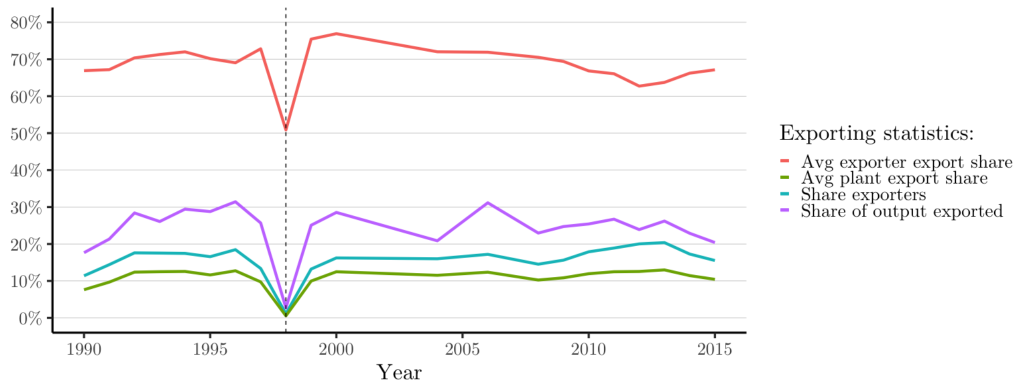

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers and reported export shares (out of total value-added).

Alternatively, the patterns may be in line with models of learning and imitation (e.g. Perla, Tonetti, and Waugh 2021), whereby the initial entry of many new and relatively unproductive plants lowered productivity growth and the subsequent better selection of plants increased the productivity growth from learning and imitation. While this may be an underlying driver of technology growth, we only find a very weak correlation between contemporaneous changes in aggregate technology and the evolution of average plant productivity (or other moments of the productivity distribution).

The increase in technology growth may also be in line with the recent theory in Ottonello and Winberry (2023) whereby constrained firms initially invest in factor accumulation and only later in activities that increase productivity. Given that we find little empirical evidence for capital deepening, only little plant-level labor deepening in the sense of rising labor shares and strong increases in financial access in the run-up to the Asian Financial Crisis, we are rather skeptical that this mechanism can explain large changes in aggregate technology. Still, given the limitations on the capital series and limited evidence on plant-level investments in technology, we cannot rule out that this mechanism is an important driving force of technology growth.

At last, since we cannot distinguish between productivity and demand drivers, changes in demand may also be an important driver of the patterns we observe. For example, the Asian Financial Crisis shows up as a more than 20% drop in technology, which is likely to be at least partly demand-driven. In line with this interpretation, Figure A.12 shows that plant-level exports almost completely plummet in 1998. Decreases in technology may also be partly explained by decreases in demand as the economy grows richer and consumers switch their demand towards services (e.g. Alder et al 2019, Comin et al 2021).

Figure A.13: Evolution of cross-sectional correlation of plant productivity and input share

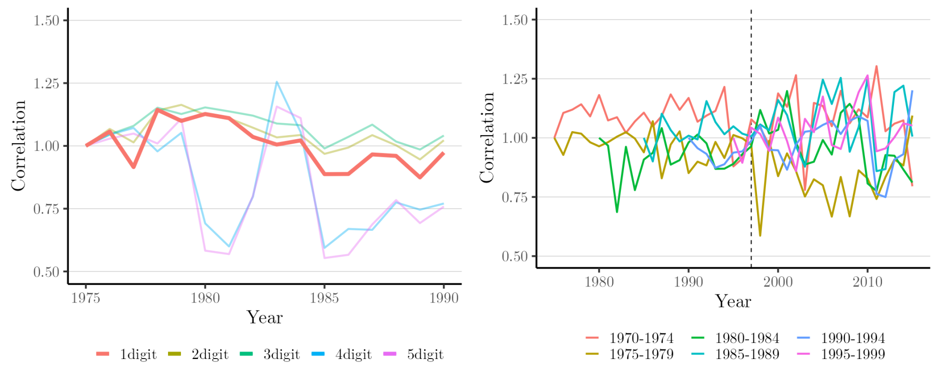

(A) Balanced Sample (B) Survivors (by cohort)

Notes: Input shares are computed based on a Cobb-Douglas aggregator as explained in the text. For within-industry results, we first estimate the correlation across plants in a given industry and year and then construct the weighted average correlation across industries using the industry's average share in manufacturing value added as a time-invariant weight. All series are normalized by the first year.

A.7 Additional results on changes in misallocation

In this section, we report two sets of additional results. First, we show additional evidence on the evolution of the covariance and correlation of plant-level productivity and input shares. Figure A.13 shows additional evidence on the correlation for a balanced panel of plants and for each cohort of plants between 1970 and 1999. We construct the balanced panel by selecting all plant-year observations between 1975 and 1990 for which the plant operated in 1975 and in 1990 and for which we observe more than 10 observations (to avoid dropping all plants for which individual years are missing or had to be dropped). Extending the time frame would drop too many plants.

For completeness, Figure A.14 also plots the covariance instead of the correlation. For the full sample, this covariance also does not increase over time due to the entry of small plants. For surviving plants, the covariance increases strongly. This is mechanical, because the sample shrinks over time, which naturally leads input shares to increase. Also, average productivity strongly increases among surviving plants, adding an additional trend. The correlation is robust to such common trends.

The second set of results is on an alternative measure of the allocation of resources. Figure A.15 reports changes in the dispersion of marginal revenue products of labor and capital; a sign of misallocation in the static model of Hsieh and Klenow (2009). Most importantly, we do not find any evidence for a decreasing dispersion in marginal revenue products over time, which could be linked to an "undoing of misallocation" that drives economic growth. Instead, we find evidence for an increase in the dispersion over time. Most of these increases happen after the Asian Financial Crisis in 1997. We think that the measured increases in the dispersion of marginal revenue products are at least in part driven by changes in measurement.

Figure A.14: Evolution of cross-sectional covariance of plant productivity and input share

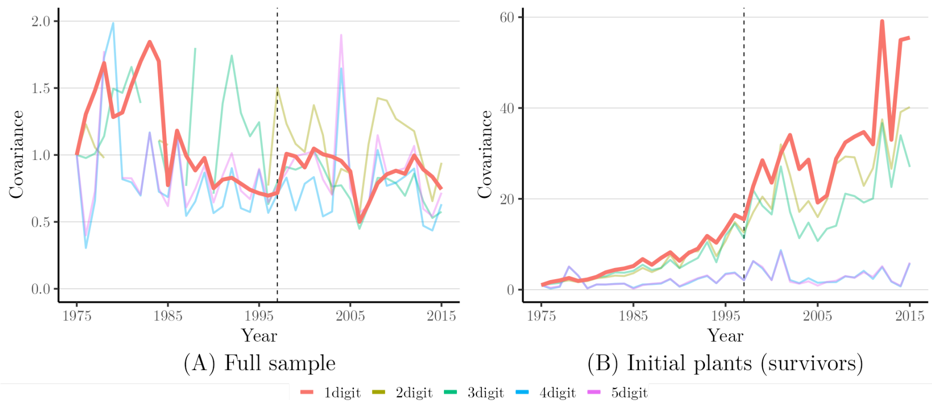

(A) Full sample

(B) Initial plants (survivors)

Notes: Input shares are computed based on a Cobb-Douglas aggregator as explained in the text. For within-industry results, we first estimate the covariance across plants in a given industry and year and then construct the weighted average covariance across industries using the industry's average share in manufacturing value added as a time-invariant weight. All series are normalized by the first year.

We refer the interested reader to the discussion of measurement changes further above.

A.8 Further event study evidence

In this section, we report event study results for hiring responses (in contrast to the labor share responses). We stick to the same treatment definition as previously. To study the dynamic hiring responses of plants to a permanent positive productivity shock, we look at the *worker share* (# worker / value added) instead of the *labor share* (wage bill / value added). The concern of looking at the labor share is that a positive productivity or demand shock at the plant level may lead workers to bargain for higher wages to share in the profit gains. If these bargaining gains slowly accumulate, then we misattribute slow increases in the labor share to labor adjustment frictions. Figure A.16 shows that this concern is unwarranted. Plants actually slowly increase hiring, analogously to the labor share.

\beginappendixB

B Model and Estimation

B.1 Adjustment costs as costs of managerial time

In the following, we show that adjustment costs can be microfounded as costs of scarce managerial time. Our goal is to make explicit how adjustment costs can capture the time

Figure A.15: Evolution of cross-sectional variation in marginal revenue products

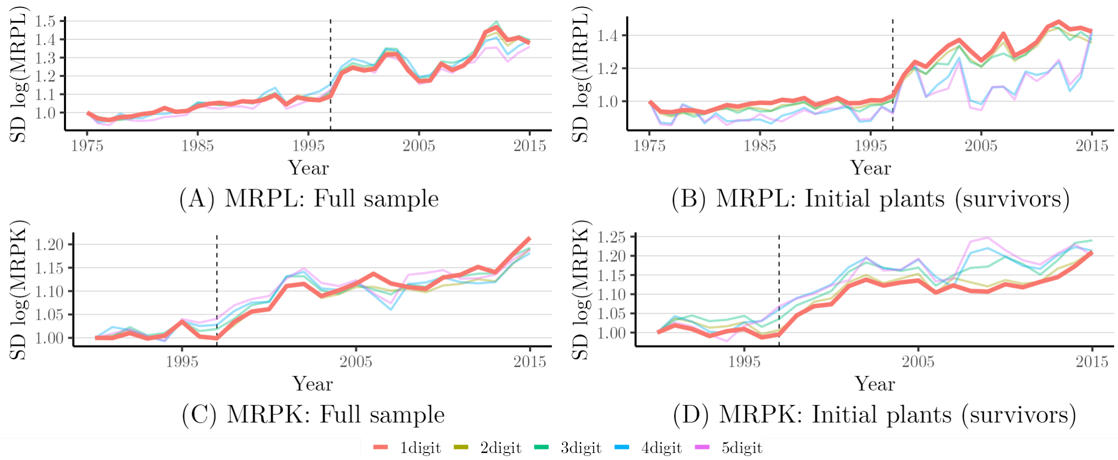

(A) MRPL: Full sample

(B) MRPL: Initial plants (survivors)

(C) MRPK: Full sample

(D) MRPK: Initial plants (survivors)

Notes: Evolution of cross-sectional standard deviation in marginal revenue products of labor and capital following Hsieh and Klenow (2009) and Gopinath et al (2017).

Figure A.16: Further event study results for worker shares

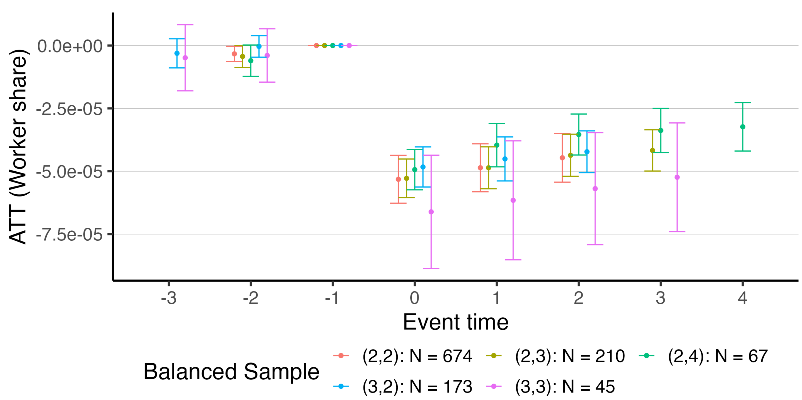

Notes: Worker share measured as ratio of number of workers over value added. Treatment definition as in main event study.

constraints of a manager working at a plant and to show why it makes sense to write adjustment costs in terms of the costs of labor w_t.

Suppose a plant owner solves the following problem:

$$V(s_{i,t}, h_{i,t-1}, \Omega_t) = \max_{h_{i,t}} \Bigg\{ y_{it}(s_{it}, h_{i,t}; z_t) - w_t h_t - w_t T(h_{i,t}, h_{i,t-1}) \\ + \lambda(s_{i,t}, h_{i,t}, \Omega_t) \left\{ -\mathbb{E}_c[c_F|\text{stay}] + \beta \mathbb{E}\Big[V(s_{i,t+1}, h_{i,t}; \Omega_{t+1}) | s_{i,t}, h_{i,t}; \Omega_t \Big] \right\} \Bigg\}$$

where $T(h_{i,t}, h_{i,t-1})$ gives the managerial time needed to implement changes in the workforce. We denote $T(h_{i,t}, h_{i,t-1})$ in terms of the efficiency units of a worker such that we can express the manager's time cost in terms of the wage w_t. One can think of $w_t T(h_{i,t}, h_{i,t-1})$ as the actual compensation that managers receive or as a combination of compensation and opportunity costs of managerial time.

Time costs are due to two main managerial tasks: (1) the task of hiring and firing, and (2) the task of reorganizing production. Conditional on the task, we think of a single unit of the task as requiring always the same amount of time (e.g. signing one contract always takes a fixed amount of time), but the total units needed depends on the organization and the amount of hiring. Hiring and firing requires (c_F^+, c_F^-) units of time for each unit of labor hired or fired $\Delta h \equiv |h_{it} - h_{it-1}|$. This time comes from filling out paperwork, signing the contracts and adding the worker to the books. Policies that affect the paperwork that plants need to fill out, will change these costs.

Next, for any workers the plant hires or fires, managers need to assign and explain changes in worker tasks. Both for hiring and firing, we assume that managerial time to assign and explain new worker tasks is proportional to the percentage change in the workforce. We assume this is for different reasons in the case of hiring and firing and thus the unit cost of changes in the workforce for hiring and firing can differ. For hiring, the plant hires $h_{it} - h_{it-1}$ workers who they need to explain their new task. The proportional time cost in the case of hiring comes from the possibility that each new worker can also learn from their coworkers. However, in the case of relatively many new hires, each new hire can learn from relatively fewer coworkers, thus increasing the time that the manager needs to add. In the case of firing, jobs may potentially not simply disappear, but need to be reorganized. In this case, the managerial time costs of reorganizing scale with the number of lost jobs $h_{it-1} - h_{it}$ (for which replacements need to be found) and are proportional to the percentage change in the workforce because this is the amount of time the manager needs to reexplain jobs to all existing workers.

In the end, the entrepreneurial time costs T_h are then given by the functional form reported in the main text:

$$T_t(h_{i,t-1}, h_{i,t}) = \begin{cases} c_{0,t}^+(h_{i,t} - h_{i,t-1}) + \frac{c_{1,t}^+}{2}\left(\frac{h_{i,t}-h_{i,t-1}}{h_{i,t-1}}\right)^2 h_{i,t-1} & \text{if } h_t > h_{t-1} \\ 0 & \text{if } h_t = h_{t-1} \\ c_{0,t}^-(h_{i,t-1} - h_{i,t}) + \frac{c_{1,t}^-}{2}\left(\frac{h_{i,t}-h_{i,t-1}}{h_{i,t-1}}\right)^2 h_{i,t-1} & \text{if } h_t < h_{t-1} \end{cases} \tag{21}$$

Note that the main text also features fixed costs that can easily be rationalized by fixed time costs of managers whenever there is a change in the organization.

B.2 Microfoundation of working capital constraint

The microfoundation of the working capital constraint can be derived from a standard limited enforcement problem (e.g. as in Buera and Shin 2013). Assume that plant managers need to first pay their workers before being able to produce and they do so by borrowing the entire wage bill $w_t h_{it}$ with a financial intermediary. For simplicity, suppose further that the time between production and paying the wage bill is $\varepsilon \to 0$ such that the costs of borrowing go to zero. Suppose further that the plant manager – after paying their workers and producing – could decide to run away with a fraction $\frac{1}{\tilde{\kappa_t}}$ of the borrowed resources $w_t h_{it}$. Isomorphically, the plant manager runs away with all of the resources, but is caught with probability $\frac{1}{\tilde{\kappa_t}}$. We assume that the only punishment in case of successful evasion is that the financial intermediary can now sue the plant manager and claim (part of) the output of the plant in period t. We assume that the claim is proportional to plant output net of value-added tax. Importantly, the plant manager never loses access to the plant and is not excluded from any future economic activity, ensuring that the constraint remains a static problem. In equilibrium, the financial intermediary will lend $w_t h_{it}$ only to the extent that no plant manager will renege on the contract, implying the financing constraint:

$$w_t h_{it} \geq \kappa_t y_{it}$$

B.3 Stationarized value function and balanced growth path after 2015

After 2015, we assume that plants expect wages, all costs and aggregate productivity to rise at the same growth rate $(1 + g)$ over time. This allows to capture realistic future growth in a parsimonious way and is in line with the entire economy being on a balanced growth path after 2015. An alternative would be to only enforce constant growth in costs and productivity and then solve for the actual endogenous wage path after 2015 that clears labor markets after 2015. This would require further assumptions on how other fundamentals in the economy evolve (e.g. wedges and technology in the rest of the economy and aggregate labor supply) and feature a continued transition towards an eventual balanced growth path (as long as assumptions on the changes in future fundamentals allow for a balanced growth path). Given that the growth path after 2015 is not identified, we think that our approach strikes a good balance between realism and parsimony.

The value function in 2015 (denoted by T and suppressing dependence on Ω_T for expositional clarity) writes:

$$V_T^M(s_{i,T}, h_{i,T-1}) = \max_{h_{i,T} \in [\underline{h}, \bar{h}]} \Bigg\{ z_T s_{i,T} h_{i,T}^{\theta} k_{i,T}^{\alpha} - w_T h_{i,T} - (r+\delta) k_{i,T} - w_T AC(h_{i,T}, h_{i,T-1}) +$$

$$\lambda(s_{i,T}, h_{i,T}) \left\{ -\mathbb{E}_c[c_F | \text{stay}_{i,T}] + \beta \mathbb{E}\Big[V^M(s_{i,T+1}, h_{i,T}) | s_{i,T}, h_{i,T} \Big] \right\} \Bigg\} \tag{22}$$

In the following, we show that under the right normalization and assuming constant growth for all costs and productivity, we can write: $V_T^M(s_{i,T}, h_{i,T-1}) = \widetilde{V}^M(s_{i,T}, h_{i,T-1}) \tilde{z}_T$, which implies that we can solve for $V_T^M(s_{i,T}, h_{i,T-1})$ by first solving for the stationary $\widetilde{V}^M(s_{i,T}, h_{i,T-1})$ and then renormalizing by $\tilde{z}_T$. To show this, we proceed in two steps. First, we find the normalizing factor $\tilde{z}_T$, which needs to grow at a constant rate $(1+g)$. We do so by deriving the optimal static capital choice: $k_{i,T}^* = \left(\frac{\alpha}{r+\delta} z_T s_{i,T} h_{i,T}^{\theta} \right)^{\frac{1}{1-\alpha}}$. Plugging the capital choice into the value function gives:

$$V_T^M(s_{i,T}, h_{i,T-1}) = \max_{h_{i,T} \in [\underline{h}, \bar{h}]} \Bigg\{ z_T^{\frac{1}{1-\alpha}} s_{i,T} h_{i,T}^{\theta} \left(\frac{\alpha}{r+\delta} s_{i,T} h_{i,T}^{\theta} \right)^{\frac{\alpha}{1-\alpha}} - w_T h_{i,T} -$$

$$z_T^{\frac{1}{1-\alpha}} (r+\delta) \left(\frac{\alpha}{r+\delta} s_{i,T} h_{i,T}^{\theta} \right)^{\frac{1}{1-\alpha}} - w_T AC(h_{i,T}, h_{i,T-1}) +$$

$$\lambda(s_{i,T}, h_{i,T}; \Omega_T) \left\{ -\mathbb{E}_c[c_F | \text{stay}_{i,T}] + \beta \mathbb{E}\Big[V^M(s_{i,T+1}, h_{i,T}) | s_{i,T}, h_{i,T} \Big] \right\} \Bigg\} \tag{23}$$

from which we can see that for output and capital to grow at a constant rate, we need that $z_T^{\frac{1}{1-\alpha}} \equiv \tilde{z}_T$ grows at a constant rate. Dividing through by $\tilde{z}_T$ gives the deflated value function $\widetilde{V}^M(s_{i,T}, h_{i,T-1})$:

$$\widetilde{V}^M(s_{i,T}, h_{i,T-1}) = \max_{h_{i,T} \in [\underline{h}, \bar{h}]} \Bigg\{ s_{i,T} h_{i,T}^{\theta} \left(\frac{\alpha}{r+\delta} s_{i,T} h_{i,T}^{\theta} \right)^{\frac{\alpha}{1-\alpha}} - \frac{w_T}{\tilde{z}_T} h_{i,T} -$$

$$(r+\delta) \left(\frac{\alpha}{r+\delta} s_{i,T} h_{i,T}^{\theta} \right)^{\frac{1}{1-\alpha}} - \frac{w_T}{\tilde{z}_T} AC(h_{i,T}, h_{i,T-1}) +$$

$$\lambda(s_{i,T}, h_{i,T}; \Omega_T) \left\{ -\frac{\mathbb{E}_c[c_F | \text{stay}_{i,T}]}{\tilde{z}_T} + \beta \mathbb{E}\Big[(1+g) \widetilde{V}^M(s_{i,T+1}, h_{i,T}) | s_{i,T}, h_{i,T} \Big] \right\} \Bigg\} \tag{24}$$

where we have made use of the constant growth in $\tilde{z}$:

$$\frac{V^M_{T+1}(s_{i,T+1},h_{i,T})}{\tilde{z}_T} = \frac{V^M_{T+1}(s_{i,T+1},h_{i,T})}{\tilde{z}_{T+1}}\frac{\tilde{z}_{T+1}}{\tilde{z}_T} = (1+g)\widetilde{V}^M(s_{i,T+1},h_{i,T})$$

In the second step, we need to prove that all other aggregate time-varying components also grow at the same rate. In the deflated value function, output and capital do not depend on time-varying aggregates anymore. All terms featuring wages require that wages grow at the same rate $(1+g)$, such that $\frac{w_T}{\tilde{z}_T} = \tilde{w}$ is a constant. The trickier parts are $\lambda(s_{i,T}, h_{i,T}; \Omega_T)$ and $\mathbb{E}_c[c_F|\text{stay}_{i,T}]$. We prove that $\lambda(s_{i,T}, h_{i,T}; \Omega_T)$ does not depend on time Ω_T if all costs grow by the same rate and that expected fixed costs $\mathbb{E}_c[c_F|\text{stay}_{i,T}]$ grow at the same rate $(1+g)$.

Using the analytic formula for the survival rate, it is easy to see that the survival rate does not vary with time-varying aggregates as long as μ_{xT} and σ_{xT} grow at the same rate $(1+g)$:

$$\begin{aligned}
\lambda(s_{i,T}, h_{i,T}; \Omega_T) &= exp\left(-exp\left(-\frac{\beta\mathbb{E}\Big[\frac{V^M(s_{i,T+1},h_{i,T},\Omega_{T+1})}{\tilde{z}_T}|s_{i,T},h_{i,T},\Omega_{T+1}\Big] - \frac{\mu_{xT}}{\tilde{z}_T}}{\frac{\sigma_{xT}}{\tilde{z}_T}}\right)\right)\\
&= exp\left(-exp\left(-\frac{\beta\mathbb{E}\Big[(1+g)\widetilde{V}^M(s_{i,T+1},h_{i,T})|s_{i,T},h_{i,T},\Omega_{T+1}\Big] - \frac{\mu_{xT}}{\tilde{z}_T}}{\frac{\sigma_{xT}}{\tilde{z}_T}}\right)\right)
\end{aligned} \tag{25}$$

At last, we use the analytic formula for the expected fixed costs to show that they indeed grow at the same rate $(1+g)$:

$$\begin{aligned}
\frac{\mathbb{E}_c[c_F|\text{stay}_{i,T}]}{\tilde{z}_T} =& \beta\mathbb{E}\Big[(1+g)\widetilde{V}^M(s_{i,T+1},h_{i,T})|s_{i,T},h_{i,T},\Omega_{T+1}\Big]\lambda(s_{i,T},h_{i,T}) - \\
& \frac{\sigma_{xT}}{\tilde{z}_T}\Gamma\left(0, exp\left(-\frac{\beta\mathbb{E}\Big[(1+g)\widetilde{V}^M(s_{i,T+1},h_{i,T})|s_{i,T},h_{i,T},\Omega_{T+1}\Big] - \frac{\mu_{xT}}{\tilde{z}_T}}{\frac{\sigma_{xT}}{\tilde{z}_T}}\right)\right)
\end{aligned} \tag{26}$$

which does not depend on aggregate time-varying components.

B.4 Formal derivation of main accounting identity

We can start by giving a formal derivation of the main accounting identity that we use to validate our model.

$$
\begin{aligned}
Y_t &\equiv \sum_i y_{it} \\
&= \sum_i z_t s_{it} f(x_{it}) = \sum_i z_t s_{it} f(x_{it}) \frac{\sum_i f(x_{it})}{\sum_i f(x_{it})} \\
&= z_t * \sum_i f(x_{it}) * \sum_i s_{it} \frac{f(x_{it})}{\sum_i f(x_{it})} \\
&= z_t * \sum_i f(x_{it}) * \sum_i (s_{it} - \bar{s}_t + \bar{s}_t) \left(\frac{f(x_{it})}{\sum_i f(x_{it})} - \frac{1}{N_t} + \frac{1}{N_t} \right) \\
&= z_t * \sum_i f(x_{it}) * \left[\bar{s}_t + N_t cov \left(s_{it}, \frac{f(x_{it})}{\sum_i f(x_{it})} \right) \right]
\end{aligned}
$$

$$
ln(Y_t) = ln(z_t) + ln \left(\sum_i f(x_{it}) \right) + ln \left(\bar{s}_t + N_t cov \left(s_{it}, \frac{f(x_{it})}{\sum_i f(x_{it})} \right) \right)
$$

B.5 Details on estimation

In this section, we provide further details on the model estimation.

B.5.1 Taxes

In the following, we discuss how we map Indonesian corporate income and value added taxes (VAT) over the period 1975-2015 to our model economy. In both cases, we assume that tax rates are constant over time and uniform across firms/plants. This is a very accurate mapping for the VAT, but less accurate for the corporate income tax rate. Throughout, we abstract from the very important topic of tax evasion and enforcement, but we discuss the empirical evidence on this.

We start with the simpler VAT, which we fix in the model economy to a constant 10%, the rate which was officially introduced in 1985 and has remained unchanged in place until now (see: Gillis 1985; Hill 2000; Basri et al. 2021). Officially, the only exemptions are on exports, which we do not model and thus abstract from and there were higher luxury product rates in place that we also abstract from. The VAT replaced an older sales tax that was in place between 1951-1985, whose rates varied from 5 to 20%, but with most sales subject to a 10% rate (Lent and Ojha 1969, 537). Enforcement of the sales tax before 1985 was almost absent, leading to widespread evasion (e.g. see: Gillis 1985), but the introduction of the VAT greatly improved (self-)enforcement and reduced evasion (Hill 2000; Basri et al. 2021), so that our assumption of a flat 10% rate seems reasonable (at least since 1985).

Changes in the corporate income tax rates are slightly more complicated, with major reforms in 1985 and 2009. We simplify our analysis by assuming a fixed 20% corporate income tax rate across firms and over time. Before 1985, many different tax rates were in place, including top marginal rates at 45%, which were non-enforced (Gillis 1985). With the 1985 reforms, corporate income rates were reduced and homogenized, with the maximum marginal rate capped at 35% (Gillis 1985). Between 1985 and 2009, the corporate income tax rates followed a 3-tiered schedule of different marginal tax rates defined over taxable profits (see: Gillis 1985; Basri et al. 2021). The different cutoffs and marginal rates were varied slightly over time, adjusting in part to inflation. For example, in 1985, a tax rate of 15 percent applied to the first IDR 10 million, 25 percent to the next IDR 40 million, and 35 percent on any taxable profits in excess of IDR 50 million (Gillis 1985). By 2009, as documented in Basri et al. (2021), a rate of 10 percent applied for the first IDR 50 million in taxable income; a rate of 15 percent applied for the next IDR 50 million; and a rate of 30 percent applied on all taxable profits over IDR 100 million. After 2009, the corporate income tax system moved to a flat 25 percent rate, with a more complicated schedule of discounts based on gross income that led to effective tax rates below 25 percent (see Basri et al. (2021) for details). Tax evasion and enforcement for the corporate tax rate posed a larger problem than for the VAT, especially before 1985 but also after (Hill 2000, 51f.; Basri et al. 2021). In the end, our 20% flat tax rate assumption tries to parsimoniously capture average effective corporate income tax rates, while abstracting from important temporal and cross-sectional variation.

B.5.2 Estimation of borrowing constraint

To identify the borrowing constraint κ_t, note that the working capital constraint writes as:

$$\frac{w_t h_{it}}{(1-\tau_t^{VAT})y_{it}} \leq \kappa_t$$

The left-hand side of this constraint is a tax-adjusted labor share, which is directly observable. The constraint gives an inequality so that in the absence of measurement error in $w_t h_{it}$ or $(1-\tau_t^{VAT})y_{it}$:

$$\kappa_t = \max_i \left(\frac{w_t h_{it}}{(1-\tau_t^{VAT})y_{it}} \right) \quad \text{for } i \to \infty$$

as long as the constraint is strictly binding for any plant. It is easy to see that the constraint will always bind for some plants as long as there is a non-zero chance for large productivity losses and plants face some positive costs of adjusting labor. The bigger problem is that any estimator based on the maximum observed adjusted labor share will be strongly influenced by outliers and individual measurement error in the labor share. As we discuss in more detail in the data cleaning Appendix, changes in the survey questions over time is one reason why we expect systematic variation in the maximum reported labor share over time that is independent of actual changes in the borrowing constraint κ_t. We thus give up on identifying variation in κ_t over time. Instead, we opt for a robust estimator of κ, taking the 95th percentile of the observed adjusted labor share. Based on this estimator, we find that $\kappa = 1.7$.

B.5.3 Discretization details

We discretize productivity and labor. Specifically, we choose 30 grid points for idiosyncratic productivity, which we select based on quantiles of the productivity distribution. We choose more quantiles at the right tail of the distribution as these high productivity plants are key for the aggregate economy. Non-parametrically estimating the transition matrix of idiosyncratic productivity is quantitatively important as other oft-used processes such as an AR(1) log-normal process cannot replicate empirically observed productivity dynamics (e.g. see Ruiz-García (2019)). We discretize efficiency units of labor h_{t-1} on a grid of 1000 points that we choose based on equal spaced quantiles, ensuring that the entire labor distribution is well represented. Specifically, we choose the bottom 990 grid points based on quantiles (ensuring that all plants in the data can be mapped to the grid) and then use the last 10 grid points to extend the upper bound for labor to allow plants in the model to grow beyond what we observe in the data.

B.5.4 Fundamentals needed for model counterfactuals

Two key sets of model fundamentals are not needed for solving the baseline model, because they are linked to reduced-form statistics that are treated as fixed along the baseline equilibrium path: the path of potential entrant distributions and the fundamentals of the rest-of-the-economy including aggregate labor supply. For model counterfactuals, however, all fundamentals are needed, so we now discuss their identification.

The potential entrant distributions can be related to objects of the baseline equilibrium path: $PE_t(s_t, h_t; \mathbf{\Omega}) = E_t(s_t, h_t; \mathbf{\Omega})/\mathbb{P}_E(s_t, h_t; \mathbf{\Omega})$, where $E_t(s_t, h_t; \mathbf{\Omega})$ is the identified path of entrant distributions and $\mathbb{P}_E(s_t, h_t; \mathbf{\Omega}) = P\Big(V^M(s_{i,t}, h_{i,t}; \Omega_t)\Big)$ gives the path of entry probability distributions. The latter is a function of the incumbent's value function, which we directly obtain from the baseline model computation, and the entry cost distribution P. Since the potential entrant distributions and the entry cost distribution are not separately identified, we make the identifying assumption that the entry cost distribution is the same as the fixed cost distribution governing plant survival.[30]

[30] Given that plant entry and survival likely depend on similar economic forces (e.g. similar outside options for not running a plant), we think this gives a reasonable estimate. We also think this gives a conservative estimate of potential entry because most plants survive, implying that most potential entrants also enter. The assumption is a form of normalizing the distribution of potential entrants and is more general than

The time path of aggregate labor supply is given by the sum of aggregated labor supply in the two sectors of the economy: $H_t = H_t^R + H_t^M$. Total labor supply in manufacturing H_t^M is identified from aggregating up plant-level labor demand h_{it} over the computed equilibrium path. To obtain H_t^R, we use the total observed number of workers l_t^R in the Rest of the Economy and map this to the total efficiency units of labor in R accounting for differential worker selection across sectors in Indonesia.[31]

For the rest-of-the-economy, we can directly identify θ_R and the sequences of A_t and τ_t^R. For this, take plant first-order conditions to obtain: $\frac{\theta_R}{(1+\tau_t^R)} = \frac{w_t h_t^R}{y_t^R}$. We use observed y_t^R and can construct $w_t h_t^R$ to obtain the left-hand side. We assume that wedges behave such that the average of the right-hand side over time is exactly equal to θ_R. Labor wedges τ_t^R are backed out such that the previous equation holds exactly. Given θ_R and h_t^R, we can simply back out the sequence A_t using: $A_t = y_t^R / \left(H_t^R\right)^{\theta_R}$.

B.6 Details on Euler estimation

In this subsection, we provide more details on the Euler estimation procedure we use and derive all main results.

Derivations for Gumbel distribution

We start out by showing that the Gumbel distribution for fixed costs allows closed-form expressions for the survival probability and the conditional expectation of fixed costs. For expositional clarity, we suppress dependence on the aggregate state Ω_t, but note that all objects generally depend on the aggregate state.

$$\lambda(s_{i,t}, h_{i,t}) = exp\left(-exp\left(\frac{-(x(s_{it}, h_{it}) - \mu_t^x)}{\sigma_t^x}\right)\right) \tag{27}$$

$$\mathbb{E}_c[c_F|\text{stay}] \equiv \tilde{g}(s_{i,t}, h_{i,t}) = x(s_{it}, h_{it})\lambda(s_{i,t}, h_{i,t}) - \sigma_t^x \Gamma\left(0, exp\left(\frac{-(x(s_{it}, h_{it}) - \mu_t^x)}{\sigma_t^x}\right)\right) \tag{28}$$

normalizing the total number of potential entrants as often done in entry models (see Aguirregabiria 2021, Chp. 5).

[31] Specifically, we use the estimates of wage differences and worker selection across rural agriculture and urban non-agriculture from Hicks et al. (2017) for Indonesia. This leads us to estimate that average efficiency units of labor are roughly two times larger in M than in R. Hicks et al. (2017), using worker-level panel data from Indonesia, find that non-agricultural jobs earn about 2.5 times higher income than agricultural jobs, but that around 80% of this earnings gap is explained by selection as captured by individual-specific fixed effects. Through the lens of our model, this implies that manufacturing workers have on average more efficiency units of labor. We enforce the point estimates of Hicks et al. (2017) across all time periods.

where $x(s_{it}, h_{it}) \equiv \beta\mathbb{E}\Big[V(s_{i,t+1}, h_{i,t})|s_{it}, h_{it}\Big]$ and $\Gamma()$ gives the incomplete Gamma function. That is, in principle, $\mathbb{E}_c[c_F|\text{stay}]$ depends not only non-linearly on the parameters $\{\mu_t^x, \sigma_t^x\}$, but also depends directly on the unknown expected future value x. However, given that the continuation value x is simply an invertible function of (observable) $\lambda(s_{i,t}, h_{i,t})$, we can rewrite the term to substitute for x:

$$
\begin{aligned}
\tilde{g}(s_{i,t}, h_{i,t}) =& \mu_t^x \lambda(s_{i,t}, h_{i,t}) - \sigma_t^x \Big\{ ln\left(-ln\left(\lambda(s_{i,t}, h_{i,t})\right)\right)\lambda(s_{i,t}, h_{i,t}) + \Gamma\left(0, -ln\left(\lambda(s_{i,t}, h_{i,t})\right)\right) \Big\} \\
\equiv& \mu_t^x \tilde{g}_1(s_{i,t}, h_{i,t}) - \sigma_t^x \tilde{g}_2(s_{i,t}, h_{i,t})
\end{aligned}
$$

We will use these equations and the invertibility of exit rates for continuation values throughout.

Identification details

In the following, we derive the estimating Euler equation and then discuss identification. To derive the estimating Euler equation, we exploit the invertibility of exit rates as shown above and simplify terms to rewrite the Euler equation only in terms of observables and model parameters:

$$
\begin{aligned}
0 =& \frac{\partial y(s_{i,t}, k_{it}, h_{i,t}, z_t)}{\partial h_{i,t}} - w_t - w_t \frac{\partial C_h(h_{i,t}, h_{i,t-1}; w_t)}{\partial h_{i,t}} + \\
& \lambda(s_{i,t}, h_{i,t})\beta\mathbb{E}\left[-w_{t+1}\frac{\partial C_h(h_{i,t+1}, h_{i,t}; w_{t+1})}{\partial h_{i,t}}|s_{i,t}, h_{i,t}\right] \Bigg\{ \\
& 1 - \lambda(s_{i,t}, h_{i,t}) + ln\Big(\lambda(s_{i,t}, h_{i,t})\Big)\left[\frac{\mu_t^x}{\sigma_t^x}(2\lambda(s_{i,t}, h_{i,t}) - 1) - \Gamma\left(0, -ln\left(\lambda(s_{i,t}, h_{i,t})\right)\right) - \right. \\
& \left. ln\Big(-ln\Big(\lambda(s_{i,t}, h_{i,t})\Big)\Big)(2\lambda(s_{i,t}, h_{i,t}) - 1) - \lambda(s_{i,t}, h_{i,t})\frac{\partial\Gamma\left(0, -ln\left(\lambda(s_{i,t}, h_{i,t})\right)\right)}{\partial\lambda(s_{i,t}, h_{i,t})}\right]\Bigg\}
\end{aligned}
$$

Given the estimating Euler equation, we can now discuss the identification of the parameters. We discuss each of the three sets of parameters in turn.

Linear and convex adjustment cost parameters: Without giving a full identification proof, one can see that the Euler equation generally identifies marginal adjustment costs $\frac{\partial C_h(h_{i,t}, h_{i,t-1}; w_t)}{\partial h_{i,t}}$ non-parametrically. Given our functional form assumption on adjustment costs, linear costs c_0 & convex costs c_1 are identified as follows: c_0 adds as a fixed wedge between the marginal product and the marginal costs of labor, but any adjustments today save on adjustments tomorrow. Thus, linear costs are pinned down by the observed labor wedge across plants and the probability of switching between shrinking and growing as determined by the volatility of the productivity process. Asymmetric linear costs are identified from the differential behavior of growing and shrinking plants. The convex costs c_1 instead scale with labor growth and are thus identified from the variation in within-plant labor demand growth across periods, again conditioned by the observed volatility of the productivity process. Low labor demand growth despite a high labor wedge will point to strong convex adjustment costs. Again, asymmetry here is identified from differential growth and shrinking (conditional on the productivity process and the state).

Fixed adjustment costs: The Euler equation does not identify fixed costs F^+ & F^- since they do not enter marginal adjustment costs. However, we note that fixed costs are identified from the (time-varying) distribution of plants that are not adjusting and for whom the Euler equation does not hold. The idea is that the more plants choose to not change their labor inputs (as we condition on previous labor and vary productivity), the higher the implied fixed costs. In the data, driven by the choice of focussing on efficiency units of labor, we do not see any plant that remains strictly inactive. We can thus not rule out that fixed costs are zero and fix them to zero throughout. We also note that in a previous version of the paper, we estimated fixed costs indirectly by solving for the model equilibrium path and also found them to be noisily estimated around zero. In cases where one might be particularly interested in fixed costs of adjustment that induce inaction – such as the sluggish responses to aggregate shocks – we think it is best to either work directly with the number of workers or at least work with the nominal wage bill.

Cost parameters (exit): One can immediately see that the Euler equation only identifies the ratio $\frac{\mu_t^x}{\sigma_t^x}$. The reason is that the Euler equation captures the marginal effect of changes in labor demand on the survival probability, which only depends on the ratio of the level and dispersion of costs. What variation in the data identifies this cost ratio? While the dependence in the Euler Equation looks daunting, the cost ratio is jointly disciplined by the size of the labor wedge, the dispersion in survival probabilities and the size of marginal adjustment costs next period over current labor demand. Given empirically estimated survival probabilities, one can see that a higher cost ratio generally increases the labor wedge.[32] Thus, high observed labor wedges push towards lower cost ratios.

Estimation details

The estimation proceeds in two stages. In the first stage, we estimate reduced-form survival probabilities $\lambda(s_{i,t}, h_{i,t}, \Omega_t)$ and dynamic labor input choices $h(h_{it-1}, s_{it}, \Omega_t)$ conditional on the state space. In the second stage, we enforce these reduced-form objects to estimate the Euler equation for the structural parameters of interest.

[32]This is as long as marginal adjustment costs tomorrow are negative, since $\lambda(s_{i,t}, h_{i,t}) ln(\lambda(s_{i,t}, h_{i,t}))\,(2\lambda(s_{i,t}, h_{i,t}) - 1)$ is generally negative (since survival probabilities are generally higher than 0.5 and $ln(\lambda(s_{i,t}, h_{i,t})) < 0$).

We start with the estimation of conditional survival probabilities and labor input choices. To flexibly estimate both, we draw on generalized additive models for their combination of flexibility and robustness in estimating semi-parametric functional forms. However, subsequent parameter estimates are very similar when using flexible polynomial regressions in the first stage instead. We start with survival probabilities, which – through the lens of the model – are a nonlinear, time-varying function in current labor and productivity. They are time-varying because exit decisions depend on the aggregate state space through current wages and aggregate productivity as well as through perfect foresight over future wages and aggregate productivity. We estimate survival probabilities as the combination of year fixed effects and semi-parametric functions in labor, productivity and their interaction. In general, we find that the estimated survival probabilities make sense while improving on a simple linear model in labor, productivity and year fixed effects. Specifically, the GAM achieves an adjusted R^2 that is roughly 13% higher than the linear model and produces survival probabilities that are increasing in productivity.

For labor input choices, we instead estimate the GAM as the combination of year fixed effects and semi-parametric functions in previous labor, productivity and their interaction.[33] The adjusted R^2 of our GAM is around 96%, driven by the huge explanatory power that previous labor has for current labor. In fact, a simple linear regression of previous labor on current labor already reaches $R^2 = 0.95$ with a coefficient of auto-correlation close to unity. Hence, a policy function that simply says to stick with past labor already explains observed labor choices extremely well. Through the lens of the model and the Euler equation, in the presence of sufficient productivity variation, this already points to large adjustment costs. However, to identify adjustment cost parameters, we require that the estimated policy functions also show variation across productivity, since changes in productivity conditional on previous labor vary the returns to adjusting labor. Figure A.17 plots the policy functions implied by our estimated GAM. Panel A varies previous labor and fixes productivity at the median, while Panel B varies productivity but fixes previous labor at the median. We find that policy functions are monotonic in productivity (in line with the model), but nonlinear such that labor is declining more strongly for low productivities and increasing more strongly for high productivities.

In the second stage, we enforce these reduced-form objects to estimate the Euler equation for the structural parameters of interest. For this stage, we follow the CCP literature in

[33] We also looked at more flexibility in how policy functions could vary in labor and productivity over time (beyond a simple year fixed effect), however, we found this to give less robust results. The reason is that such higher flexibility is more likely to predict more erratic differences in input choices over consecutive years, which "violate" the smoothing motive implied by the Euler equation and which can then only be rationalized by extreme adjustment cost parameters. We believe that if the underlying data has less measurement issues (that are driven by common aggregate components) and is maybe taken from an economic environment where there are less changes over time, it is more feasible to allow for more flexibility in how policy functions can change over time.

Figure A.17: Step 1 of CCC estimation: flexibly estimated labor input choices

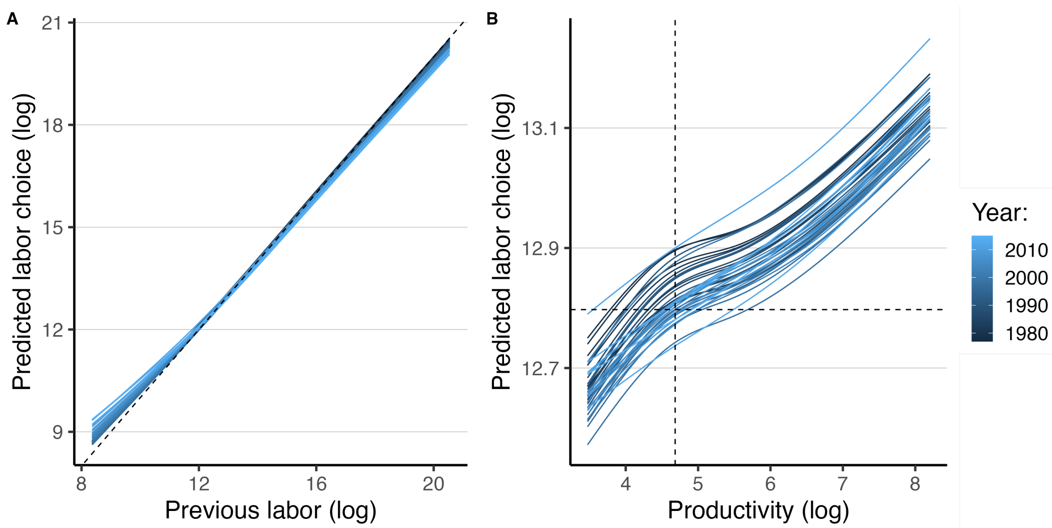

Notes: Panel A gives labor choices over previous labor (fixing productivity at the unconditional sample median). Panel B gives labor choices over productivity (fixing previous labor at the unconditional sample median). Dotted lines give unconditional sample medians.

imposing our model structure. This means that we impose the same discretized state space as in our model and the same process of idiosyncratic productivity to ensure that the Euler estimation gives parameter estimates that are fully consistent with our model. We also impose CCC-estimated input choices for plants at time t and for all states in $t+1$ conditional on the plant's state in t, greatly reducing the noise in the estimation: this means that for a plant with labor h_{it-1} and s_{it}, we do not use observed h_{it} and h_{it+1}, but instead the predicted values $\hat{h}_{it}(s_{it}, h_{it-1}, \Omega_t)$ and $\hat{h}_{it+1}(s_{it+1}, \hat{h}_{it}, \Omega_{t+1})$, where the latter is over all possible s to be able to compute the expectation term.[34]

We think it is important to briefly mention the computational gains here. An important step in the estimation is that we follow Bajari, Benkard, and Levin (2007) in exploiting the fact that adjustment cost parameters enter linearly in the problem, which is due to the linearity of the marginal adjustment cost specification and due to plants' risk neutrality. Practically, this means that we can compute all expectation terms outside the parameter loop, greatly speeding up the parameter estimation. It is important to highlight the computational gains from this step alone: We can estimate parameters from both steps in less than 2 minutes on a standard personal computer and moving to annually estimated parameters (increasing the number of parameters by a factor of 40) can sometimes even be faster – the reason is that the pre-computed terms all stay the same and instead of estimating X parameters on N data

[34]An alternative to our approach would be to directly use observed plant-level future labor adjustments as noisy realizations of expected labor adjustments, without enforcing model-based expectations (e.g. Hall 1979). Our approach is closer to our model, greatly reducing noise in the estimation, which is particularly problematic for the estimation of convex costs that disproportionately react to outliers. However, this makes our estimation approach – as other CCC/CCP estimators – more susceptible to model misspecification.

Table A.1: Main Euler estimation results

Parameters	Estimates	Std error	95% CI
c_0^+	0.735	0.010	[0.715,0.755]
c_1^+	36.656	0.059	[36.54,36.772]
c_0^-	0.000	0.011	[-0.022,0.022]
c_1^-	12.593	0.073	[12.45,12.736]
Cost ratio	-0.366	0.004	[-0.374,-0.358]

Details: Pooled across all consecutive plant-year observation pairs (N = 358,240). Adjustment cost parameters are restricted to be (weakly) positive and the cost ratio is bounded between -0.577 and -0.366 to ensure that median and mean costs are sufficiently far apart and rationalize dispersion in exit probabilities. Inference for corner solutions should be treated with care. Standard errors are not yet corrected for the multi-stage estimation.

points, we now separately estimate X parameters on N/T data points T times, which one can even parallelize.

We estimate structural parameters via nonlinear least squares (NLS). We do so by assuming that the Euler equation can be written as: $f(\Theta) + \eta_{it} = 0$, where η_{it} is model misspecification error or additive measurement error and Θ is the vector of parameters. Table A.1 reports estimated results. Table A.2 separately estimates parameters for the period before the Asian Financial Crisis and for the period after. One can see that estimated adjustment cost parameters are considerably higher post-1999 than before. At last, we also estimate adjustment costs at an annual level. Figure A.18 shows the yearly estimated convex adjustment costs and shows that they tend to increase over time.

\beginappendixC

B.7 Further model validation exercises

In this section, we show further model validation results. Specifically, Figure A.19 shows the distribution of labor shares. A key feature of the data, which the model captures, is that while average labor shares increase when holding productivity constant, the large observed shifts in productivity due to selection and productivity convergence imply that increasingly more production is concentrated in more productive plants. These productive plants, however, have substantially lower labor shares, in part because they were surprised by positive productivity shocks and adjust labor only slowly, and in part because they avoid large labor increases anticipating future mean reversion in productivity. Together, this implies that the aggregate

Table A.2: Euler estimation results: Pre vs. Post Crisis

Parameters	Estimates		Std error		95% CI	
	Pre-1997	Post-1999	Pre-1997	Post-1999	Pre-1997	Post-1999
c_0^+	0.635	0.731	0.008	0.013	[0.619,0.651]	[0.706,0.756]
c_1^+	28.590	38.495	0.098	0.069	[28.398,28.782]	[38.356,38.63]
c_0^-	0.000	0.000	0.006	0.013	[-0.012,0.012]	[-0.025,0.025]
c_1^-	7.206	15.605	0.167	0.088	[6.879,7.533]	[15.433,15.777]
Cost ratio	-0.366	-0.366	NaN	0.004	[NaN,NaN]	[-0.374,-0.358]

Details: Pooled across all consecutive plant-year observation pairs (N = 136,918 for pre, N = 209,075 for post). Adjustment cost parameters are restricted to be (weakly) positive and the cost ratio is bounded between -0.577 and -0.366 to ensure that median and mean costs are sufficiently far apart and rationalize dispersion in exit probabilities. Inference for corner solutions should be treated with care. Standard errors are not yet corrected for the multi-stage estimation.

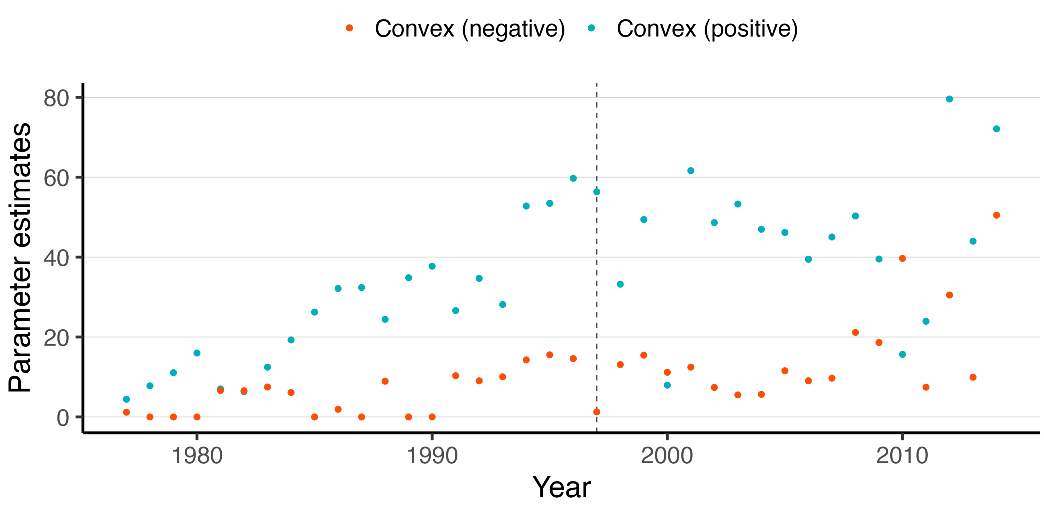

Figure A.18: Annually estimated convex adjustment costs. Details in text.

labor share is low and remains low over time, with the median labor share even declining in the data. Apart from the decline in the median labor share, the model correctly captures these distributional changes.

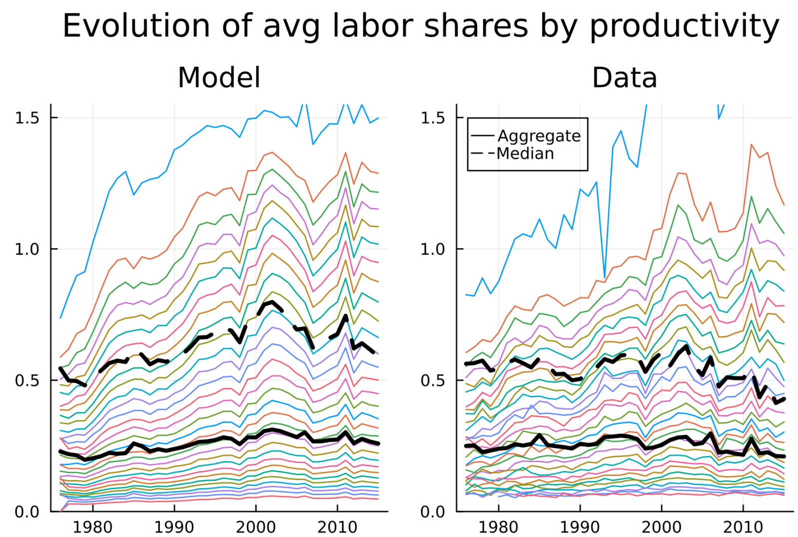

Figure A.19: Changes in the entire distribution: baseline model versus data.

C Counterfactuals and results

C.1 Further results on decomposing the drivers of growth

Figure A.20 reports the evolution of the employment distribution in the baseline (model) economy versus the counterfactual (model) economy where only initial conditions in 1976 play out over time (and all other fundamentals in 1976 are fixed).

C.2 Details on INPRES evaluation

This section provides further details on the model-based evaluation of the INPRES school construction program. We start out with more details on how we interpret the program through the lens of the model, how the program's effects map into changes in model fundamentals and how we identify counterfactual fundamentals of the economy had the INPRES school construction program not been implemented. In the second part, we then provide more detailed results on the model-based evaluation of the INPRES program.

We assume that the program's direct effect only goes through improving children's education as measured by human capital h in the model. We can be agnostic about how schools raised

Figure A.20: Evolution of employment distribution: baseline model versus initial conditions counterfactual

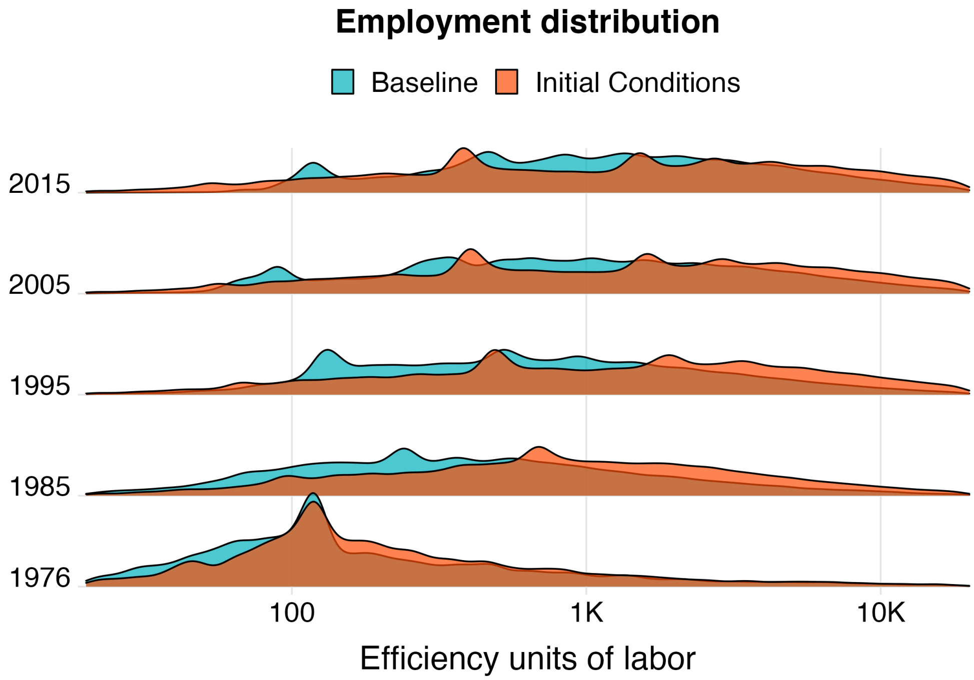

Notes: Details in the text.

human capital, capturing a combination of changes along the extensive margin (some children are induced into going to school in the first place), intensive margin (some children stay longer in school) and quality margin (more schools and teachers meant smaller classroom sizes and closer proximity that may provide additional time for learning). We assume that overall demographic changes are not affected. The direct effects of changes in human capital then induce a number of endogenous changes in the model. Specifically, the increase in human capital puts downward pressure on wages and drives up labor demand both in the rest of the economy and across all manufacturing firms. Increases in human capital also have an effect on the endogenous entry and exit of firms, but we assume that this only happens through input costs. Specifically, we rule out that increases in education may have a direct effect on entrepreneurial choices and the distribution of potential entrants. Given that the policy only had measurable effects on primary and secondary schooling outcomes and that entrepreneurs in larger manufacturing firms are more likely to have tertiary education, we think this is a reasonable assumption.

Formally, we model the INPRES program as changing individual-level human capital, which aggregates up to aggregate human capital H_t over time. While the model allows for individuals being differentially affected by the school construction program and also differentially select into different sectors, the general equilibrium results only depend on the change in aggregate human capital. Denoting by L_t the evolution of the number of workers (which we assume to be unaffected by the program), we need to know the average effect of the program on human capital per worker. Through the lens of the model, the average treatment effects estimated via differences-in-differences exactly capture the average differences in human capital h induced by the INPRES program for workers who were treated by the program, netting out aggregate changes in the wage. We can rewrite the effect of the overall program on wages as the combination of three separate effects that have been estimated in the literature: the effect of the program on school construction, school construction on years of education and years of education on wages. We further simplify the setup by assuming that we can treat the three terms as separate expectations (which is true under homogeneous treatment effects or when the shocks driving variation in the treatment effects are independent):

$$\begin{aligned}\mathbb{E}\left[\frac{\partial \text{wage}}{\partial \text{program}}\right] &= \mathbb{E}\left[\frac{\partial \text{wage}}{\partial \text{years of schooling}}\frac{\partial \text{years of schooling}}{\partial \text{no. of schools}}\frac{\partial \text{no. of schools}}{\partial \text{program}}\right] \\ &= \underbrace{\mathbb{E}\left[\frac{\partial \text{wage}}{\partial \text{years of schooling}}\right]}_{\approx 0.1 \text{ (Chaisemartin \& d'Haultfoeille '18)}} \underbrace{\mathbb{E}\left[\frac{\partial \text{years of schooling}}{\partial \text{no. of schools}}\right]}_{\approx 0.25 \text{ (Akresh et al '21)}} \underbrace{\mathbb{E}\left[\frac{\partial \text{no. of schools}}{\partial \text{program}}\right]}_{1.98 \text{ (per 1k children; direct measure)}}\end{aligned}$$

Following Akresh, Halim, and Kleemans (2023), we assume from this that the program on average increased years of schooling by half a year for individuals of any treated cohort. We further follow Akresh, Halim, and Kleemans (2023) by assuming that individuals join the workforce at age 18 and that differences in human capital induced by the program are constant over a person's life, in line with one-time educational gains. The primary schools built by the INPRES program between 1973-1979 are for children between the ages of 7-12 years, such that children fully treated by the INPRES program first joined the labor force by 1984. As in the existing literature, we assume that all cohorts born after 1968 benefit from the INPRES program. This assumes that the last cohort that we observe in 2015 still benefited from INPRES schools in 2009 (their last year of primary school).[35] To avoid having to deal with partial treatment, we further assume that cohorts before 1968 did not benefit from the program. Through the lens of our model, the INPRES program thus led to variation in aggregate human capital over the period 1986 to 2015.

We use the following steps to construct counterfactual paths of aggregate human capital in the absence of the INPRES program:

1. Start from aggregate human capital H_t given by model
2. For each year t between 1975 and 2015:
    - count share of working population affected by INPRES treatment (ϕ_t^T) & get

[35]The program initially planned for the INPRES-built schools to last for 20 years, however, Akresh, Halim, and Kleemans (2023) note that most even exist 40 years later. Our assumption implies that the maximum age for an INPRES school in our data is 36 years, well in line with the age range of INPRES schools.

aggregate human capital without INPRES: $\tilde{H}_t = H_t * (1 - \phi_t^T) + H_t * \phi_t^T * \frac{1}{\tilde{\beta}}$ where $\tilde{\beta}$ is the corresponding average treatment effect of the program (here: assume that this is 1.05 given as above)

We thus implicitly assume that all cohorts have the same average human capital. This is unlikely, but in the absence of better worker-level estimates of human capital, this is the best we can do. Still, we think that there are two biases that push in opposite directions so that we think the overall bias may not be too strong. First, younger cohorts likely have more human capital, which means we overestimate human capital in the absence of the INPRES program $\tilde{H}_t$. At the same time, individuals likely experience human capital increases over their life cycle such that young cohorts have less experience and less human capital, biasing our results in the opposite direction.

To construct the path of ϕ_t^T we draw on representative and harmonized population census data that we retrieve via IPUMS. For each available census wave $t \in \{1980, 1985, 1990, 1995, 2000, 2005, 2010\}$, we construct the share of working age individuals (between 18-65) who have been born in 1968 or later, which gives ϕ_t^T. For the years in between, we extrapolate from the previous census wave assuming there is no differential mortality risk. The treated share in the working population is zero before 1986 and then - due to baby-boom cohorts - increases rapidly to almost 20% by 1990, 50% by 2000 and 75% by 2015.

Applying this time path, we find that the INPRES program raised the annual economy-wide level of human capital by 3.6% by 2015. While important, this effect only accounts for less than 2% of the more than doubling of human capital per worker that we estimated over the entire time period from 1975-2015. These numbers also explain the model-based aggregate effects of the INPRES program that we find.

C.3 Details on FDI policy counterfactual

In this section, we provide further details on the model-based evaluation of attracting more foreign-owned plant entrants in Indonesian manufacturing. We start out with more details on the entry of foreign-owned plants, how we interpret this variation and map it to the model, how regulatory changes on FDI map to changes in model fundamentals and how we identify counterfactual fundamentals of the economy had Indonesia's FDI policy been different. In the second part, we then provide more detailed results on the model-based evaluation.

We start out by showing variation in ownership across Indonesian manufacturing plants over time. In Figure A.21 we plot the share of total manufacturing output that is owned either by the state (central + local govt), private domestic or foreign owners. We construct this by summing up all value-added output across plants but taking plants' reported ownership

Figure A.21: Evolution of ownership shares for Indonesian manufacturing plants

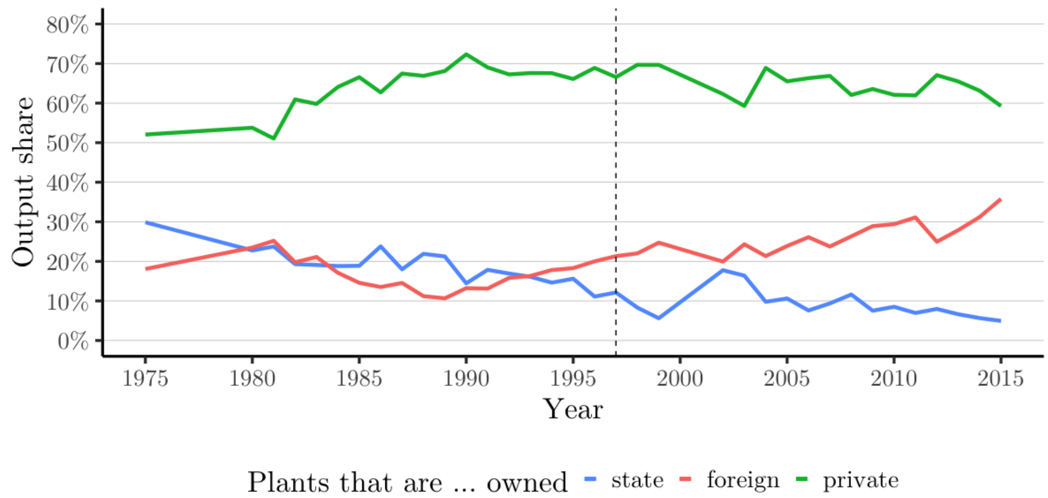

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. Figure reports the fraction of total manufacturing output that is owned by the state (either local or central government), domestic private owners or foreign owners measured by summing up all value-added output of plants and weighting plants' output by their respective reported ownership shares.

shares as weights. The main movement can be observed in the decline of state ownership from around 30% in 1975 to around 5% in 2015 and the rise of foreign ownership. The domestic private sector is by far the largest actor and owns between 60-70% of all manufacturing production. If we were instead to look at the share of plants, we find that more than 90% are fully domestically owned, which is stable over time. Again, we find that state ownership declines over time and foreign ownership increases, making up almost the entire remainder of 10% by 2015.

For the role of FDI policy, we are specifically interested in the effect on plant entry. Plant entry is particularly important, because most variation in foreign ownership shares is across and not within plants as plant-level ownership shares are relatively constant.[36] Figure A.22 thus reports evidence on the importance of foreign ownership among new entering plants.[37] Within a given year, foreign-owned plants (those that are majority foreign owned) make up around 3.8% of entering plants, but they account for 18-19% of total output among entrants. Entrants with some foreign ownership are almost always close to fully foreign-owned with average ownership shares around 80% and the median at 95%. Figure A.22 also documents

[36]For example, the variation in foreign ownership explained by plant fixed effects is 78% and the persistence in foreign ownership as measured by an AR(1) regression is $\rho \approx 0.9$. Restricting only to plants that were ever foreign owned gives slightly lower numbers with the $R^2 \approx 0.6$ and $\rho \approx 0.77$.

[37]We define new entering plants as plants that enter the panel for the first time. While foreign-owned plants are larger and unlikely to not make the cutoff of 20 workers, we want to make sure to compare foreign- and non-foreign-owned plants correctly. To this end, we further impose the restriction that the plant has to be younger than 10 years (which is the spacing of the censuses). The sample of entrants without the age restriction looks very similar.

Figure A.22: Evolution of the entry of plants with foreign ownership

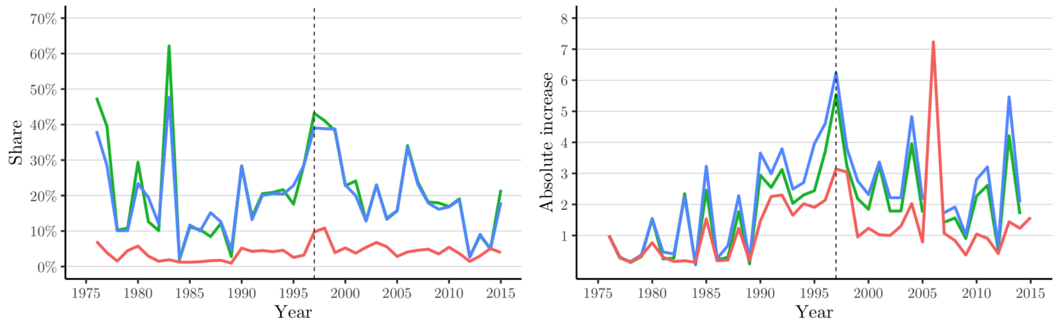

(A) Share of foreign entrants over time

(B) Number of foreign entrants (1976 = 1)

Based on: # of Plants Output (Majority) Output (Weighted)

Notes: Data based on Indonesian manufacturing plant census (1975-2015) with 20+ workers. In both Panels, the output measures aggregate up value-added across plants, with Majority summing up output for entering plants who are at least 50 percent foreign-owned, and Weighted summing up output based on the respective ownership share. Number of plants instead constructs the share and absolute increase based on the total number of majority-foreign-owned plants.

important variation across time - variation that we exploit for identifying the effect of FDI policy. Specifically, the Indonesian FDI regulatory regime turned increasingly restrictive throughout the 1970s, forbidding 100% foreign ownership and banning FDI entirely in some sectors of the economy (see: Hill 2000). This policy regime reverted only in the second half of the 1980s with simplifications and more transparency over existing restrictions introduced in 1987. In 1992, 100% foreign ownership was permitted again and the 1990s saw increasing attempts at luring foreign manufacturing plants. We can see some of these changes in observed entry, including a marked increase in the absolute and relative weight of foreign entrants between 1990 and the Asian Financial Crisis in 1997. We also see a correlation of these policy changes with the aggregate ownership series, with the importance of foreign ownership increasing steadily from 10% in the late 1980s up to 35% by 2015.

We exploit this policy variation to study the influence of changes in foreign direct investment policy in the Indonesian growth experience. We proceed as follows: We first identify the distribution of actual foreign entrants, which - through the lens of the model - is a reduced form object that masks the underlying distribution of potential foreign entrants. To take out the variation in foreign entry that is purely explained by changes in the economic conditions that make entry more or less attractive, we proceed as before and use the model-identified, time-varying entry probabilities to invert for potential foreign entry distributions. In the next step, we are interested in whether FDI policy changes can account for changes in these potential foreign entry distributions over time. To do so, we compare the period of "restrictive FDI" from 1975-1986 with the period of "FDI promotion" from 1987-1997. Given the known data limitations of annual variation in entry (e.g. the bunching of entry around census years), we aggregate the potential entry distributions within each of the two time periods. Luckily, the census waves do not fall in a year between, so that there is no ambiguity in how to attribute entrants. We then compare changes in these two aggregated potential entrant distributions. For simplicity, we measure the effect of changes in the FDI policy on changes in potential foreign entrants by comparing the weighted mass of potential entrants across the two periods, taking as weights either the plants' value added or employment at entry.[38]

C.4 Making a (more impressive) Growth Miracle

In this part of the Appendix, we move away from Indonesia's historical growth experience and ask whether and how Indonesia could have experienced a more impressive manufacturing growth miracle, closer in comparison to the experiences of countries such as China or Malaysia. For this, we study two important policy levers that both have sizable growth effects, but play out differently over time. Specifically, we look at reduced-form policy changes that either reduce (convex) labor adjustment frictions or increase the annual growth in aggregate technology in manufacturing, but deliver the same long-run growth in manufacturing output.[39]

Figure A.23 shows how the manufacturing miracle would have played out differently in the two alternative scenarios that both see a doubling of manufacturing output by 2015 compared to the baseline miracle economy. As expected, lower adjustment costs lead to faster hiring and thus faster transitions such that output growth is initially higher. With lower adjustment costs, far more large manufacturing plants emerge, driving up the average plant size. Growth in manufacturing technology, on the other hand, makes all plants more productive, leading to more entry and less exit of small plants, a stronger left tail and a much lower covariance of idiosyncratic productivity and resources.

[38]We thus only use the time series variation and not a differences-in-differences identification design. The model-based trend correction should ensure the validity of the approach and we do not think that it is credible to compare foreign entrants to domestic entrants in a differences-in-differences design, because changes in domestic entrants (e.g. due to changes in demographics) may likely show very different trends. An alternative with a DiD design would be to compare potential entrant distributions across different countries or across different industries that were differentially treated by the regulatory regime. This is an interesting approach that we leave for future work.

[39]To reduce labor adjustment frictions, we consider a policy package that reduces both linear and convex adjustment costs. For the linear hiring and firing costs, we consider a hiring subsidy for each new hire of around 25% of the annual wage bill. Changes in convex costs are harder to map directly to tangible policies. Given our microfoundation in terms of scarce managerial time and talent, we think of them as policies that improve managerial quality in the economy such as training programs. We consider a feasible policy mix that halves the estimated convex cost parameters (for hiring and firing) – in line with the lower end of annually estimated adjustment cost parameters that we find in the data. For aggregate technology in manufacturing, z_t, we consider a policy that raises its annual growth by a constant rate to the degree that manufacturing output in 2015 is the same as in the adjustment cost counterfactual.

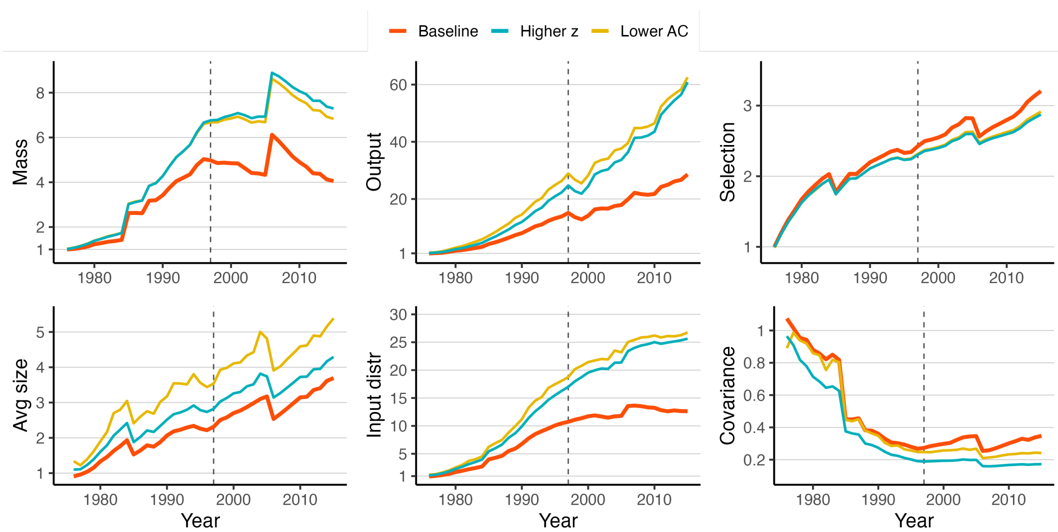

Figure A.23: Main miracle economy counterfactuals.