Journal of Development Economics 97 (2012) 339–351

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

Creative accounting or creative destruction? Firm-level productivity growth in Chinese manufacturing ☆

Loren Brandt [a], Johannes Van Biesebroeck [b,*], Yifan Zhang [c]

[a] University of Toronto, Canada
[b] K.U. Leuven and CEPR, Centre for Economic Studies, Naamsestraat 69, 3000 Leuven, Belgium
[c] Lingnan University, Hong Kong

ARTICLE INFO

Article history:
Received 9 February 2010
Received in revised form 11 February 2011
Accepted 23 February 2011

JEL classification:
D24
O14

Keywords:
TFP
Industrial development
Reallocation
Entry
Exit
China

ABSTRACT

We present the first comprehensive set of firm-level total factor productivity (TFP) estimates for China's manufacturing sector that spans China's entry into the WTO. For our preferred estimate, which adjusts for a number of potential sources of measurement error and bias, the weighted average annual productivity growth for incumbents is 2.85% for a gross output production function and 7.96% for a value added production function over the period 1998–2007. This is among the highest compared to other countries. Productivity growth at the industry level is even higher, reflecting the dynamic force of creative destruction. Over the entire period, net entry accounts for over two thirds of total TFP growth. In contrast to earlier studies looking at total non-agriculture including services, we find that TFP growth dominates input accumulation as a source of output growth.

1. Introduction

China has enjoyed impressive labor productivity growth averaging nearly 8% for a period now spanning three decades. Considerable debate persists over the sources of this growth and the relative contributions of improvements in total factor productivity (TFP) versus the mobilization of additional resources, notably physical and human capital. Studies using aggregate data and combining agriculture and non-agriculture typically find TFP contributing approximately half of labor productivity growth (Bosworth and Collins, 2008; Perkins and Rawski, 2008).

In a widely cited study focusing solely on the non-agriculture sector covering the period up through 1998, Young (2003) paints a much less impressive picture of China's growth story. Correcting for potential biases in official deflators and the measurement of human capita, but otherwise using official data, Young reduces the estimate of productivity growth for the sector between 1978 and 1998 from a very respectable 3% to a more pedestrian 1.4%. Over this period, non-agriculture was the source of nearly 80% of GDP.[1]

The aggregate results hide important heterogeneity. TFP growth in industry, which represents forty percent of GDP and is the source of 90% of exports, is likely to be much higher than in the service sector, to which reform and market liberalization have only come with a long lag (Bosworth and Collins, 2008). Earlier empirical studies also identify a significant gap in productivity in industry between the rapidly expanding non-state sector and state-owned firms (Groves, et al., 1994; Jefferson and Rawski, 1994). Qualitatively, rising firm capabilities and productivity in industry have been linked to the expanding role of market forces, massive entry of new firms, and intense competition (Brandt et al., 2008).

An analysis of Chinese manufacturing on par with that carried out for other countries has been handicapped by a lack of firm-level data sets. This constraint is gradually being relaxed, allowing more in-depth analysis at the micro level of key aspects of behavior in manufacturing that are missed at the macro level—see, for example, Bai et al. (2006), Dougherty et al. (2007), Hsieh and Klenow (2009), and Park et al. (2010). This paper builds on that work.

☆ We wish to thank three anonymous referees and the editor for excellent comments that have greatly improved the paper. Funding by Canada's SSHRC and the EU's ERC is gratefully acknowledged.

* Corresponding author.
E-mail address: jo.vanbiesebroeck@econ.kuleuven.be (J. Van Biesebroeck).

[1] Brandt and Zhu (2010) revise Young's original estimates upwards, reflecting revisions to official GDP figures, and biases in Young's deflator for services.

0304-3878/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jdeveco.2011.02.002

Drawing on an unbalanced panel of firms between 1998 and 2007 that represents approximately 90% of gross output in manufacturing, we present the first comprehensive set of firm-level productivity estimates for Chinese manufacturing that spans China's entry into the World Trade Organization (WTO). The absolute size of China's manufacturing sector and its exports make this important in its own right. Over the period we examine, we find firm-level TFP growth of manufacturing firms averaging 2.85% for a gross output production function and 7.96% for a value added production function.

Total TFP growth for the manufacturing sector was even higher due to massive entry of new firms with above average productivity levels and growth rates and the exodus of inefficient incumbents. When new firms replace exiting firms, the reallocation of input factors tends to enhance efficiency. Over the full sample period, our results identify net entry as the source of more than two thirds of total productivity growth, exceeding its contribution in U.S. manufacturing (Haltiwanger, 1997).[2]

In all, we find that TFP growth coming from improvements in continuing firms (the intensive margin of TFP growth) and through net entry (the extensive margin of TFP growth) was the source of over half of value added growth in manufacturing over the 1998–2007 period. TFP's contribution to labor productivity growth is even higher at two-thirds. The rest of the growth in value-added was the result of increases in total capital and labor use in manufacturing, much of which was associated with the entry of new firms. Our findings for the manufacturing sector are sharply at odds with the view of Young (2003) and others (Zheng et al., 2006) that productivity growth outside of agriculture has been mundane or ordinary. However, our results reveal that aggregate TFP growth in Chinese manufacturing remains constrained by limited efficiency-enhancing input reallocations between active firms, confirming results in Hsieh and Klenow (2009).

These findings have important implications for government policy. First, the high firm-level TFP growth estimates imply that Chinese manufacturing output growth will not disappear any time soon as input accumulation diminishes. The labor force will peak in a few years (Perkins and Rawski, 2008), and rates of investment are expected to come down as China rebalances. TFP growth will also help firms in China weather rising labor and other input costs. Second, increasing competitive pressure and the adoption of new technology are often mentioned as drivers of TFP growth. Learning is not only important to the upgrading efforts and productivity growth among continuing firms, but is also equally important to the contribution of new entrants. For entrants, there are two dimensions to learning: first, identifying new opportunities making successful entry possible and second, improving productivity subsequent to entry. Policies that facilitate both kinds of learning are the key to sustained growth in the medium term. Third, as input growth slows and the technology gap with advanced countries narrows, further reforms to enhance efficient allocation of resources still provide important growth potential. A policy of liberalizing entry and facilitating exit has already played an important role in this regard. Removal of constraints that underpin productivity differences among existing firms, including those between the state and non-state sectors will have to be tackled next.

Working with firm-level data for China has its difficulties. One of the additional contributions of this paper is to carefully describe and document these data. We make publically available online the complementary data we have constructed, including deflators, industry concordances, adjustment to capital stock series, etc. that are required to make full use of the data. Furthermore, in light of important concerns of Young and others, we examine the robustness of our results to a host of measurement issues. We show how alternative treatment of key variables often reduces productivity growth, but does not alter the basic picture.[3]

A particularly important aspect of the data work was the construction of linkages over time in firm-level observations when firm ID codes changed. This often occurs when active firms are restructured and it is important not to classify such instances as exit and subsequent entry. We find that one-sixth of the Chinese firms in our sample have at least one ID change. The ability to track firms as they are being restructured is an important precondition to being able to conclude that net entry has been the dominant force in productivity growth in Chinese manufacturing.

The remainder of the paper is organized as follows. In the next section we describe our methodology for measuring productivity. Section 3 describes the data set and the construction of the key variables. An online Appendix provides more detailed documentation. In Section 4 we describe the Chinese results at the firm level, the performance of entrants and exiting firms, and the aggregate productivity growth experience. The latter allows us to “line up” our findings for industry with estimates from the literature for the entire economy. We also decompose the productivity residual to identify the types of heterogeneity most important to the aggregate evolution of productivity. Section 5 concludes.

2. Productivity measurement

The most widely used measure of productivity is labor productivity, the ratio of value added to the number of hours worked or the number of workers. In China's national accounts the share of labor earnings in GDP is only around one half for the full economy and it is even lower in manufacturing. Not accounting for capital intensity is likely to paint a misleading picture.

Multi-factor productivity is only defined relative to a particular production technology – input aggregator – which we can characterize by a production function:

$$Q_{it} = A_{it} F_{it}(X_{it}). \tag{1}$$

It is inherently a relative concept, and we can write it in general as

$$\ln\left(A_{it}/A_{j\tau}\right)_k = \ln\left(Q_{it}/Q_{j\tau}\right) - \ln\left(F_k(X_{it})/F_k\left(X_{j\tau}\right)\right). \tag{2}$$

For productivity growth comparisons, the same firm enters the numerator and denominator ($i = j$) and for productivity level comparisons we fix time instead ($t = \tau$). Even though the production function in Eq. (1) is allowed to differ between firms and over time–as denoted by the subscript on the input aggregator–we have to use a uniform technology (k) for both units to perform the productivity comparison in Eq. (2).

To accurately measure productivity, one needs to accurately measure inputs and outputs and to estimate the input substitution possibilities that the technology allows. The first task is described at length in the next section; we now turn to the second task. Van Biesebroeck (2007, 2008) compares alternative methodologies to estimate productivity and finds different estimates to exhibit very high correlations. The assumption of a uniform production technology for all firms in an industry stands out as one modeling choice that the results are sometimes sensitive to.[4] Therefore, we implement two estimation procedures.

[2] Recent qualitative work of Brandt et al. (2008) and work with cross-sectional data for 1998 and 2005 by Jefferson et al. (2008), already point to entry and exit as important drivers of the dynamism in the manufacturing sector. Here, we provide decomposition results for China's manufacturing sector that are directly comparable to other studies in the literature.

[3] In particular, using different price deflators does influence absolute growth estimates, but the relative contribution of TFP growth and input accumulation in output or labor productivity growth are only affected to the extent that price deflators are biased differently for wages, capital, and output.

[4] This mattered in particular for the evaluation of learning-by-exporting effects.

The benchmark productivity measure is a straightforward Törnqvist index number, which does not require the estimation of any parameters. Caves et al. (1982a) illustrate how this can be interpreted as an exact measure of the relative productivity of two observations. In particular, it is the geometric average of the ratio in Eq. (2) using either the technology of observation "*it*" or that of "*jτ*". The intuition is that a cost-minimizing firm will make sure the relative factor price ratio equals the local elasticity of substitution between inputs of the production technology. As a result, factor shares can be used to control for input substitutability. The main benefit is the ability to allow for technology heterogeneity in the input elasticities across observations.

Productivity growth is calculated in the usual way as

$$TFPG_{it}^{IN} = (q_{it} - q_{it-1}) - \overline{S_{it}}(l_{it} - l_{it-1}) - \left(1 - \overline{S_{it}}\right)(k_{it} - k_{it-1}), \tag{3}$$

where $\overline{S_{it}} = (S_{it} + S_{it-1})/2$ is the average wage bill in value added.[5] Small cap variables represent logarithms and the three variables q, l, and k indicate output (value added), labor, and capital. To compare the productivity level across firms within the same industry, Caves et al. (1982b) propose the following multilateral productivity measure:

$$\ln TFP_{it}^{IN} = (q_{it} - \overline{q}_t) - \widetilde{S_{it}}\left(l_{it} - \bar{l}_t\right) - \left(1 - \widetilde{S_{it}}\right)\left(k_{it} - \overline{k}_t\right). \tag{4}$$

Each firm is compared to the hypothetical average firm in the industry. The weight on the labor input difference is $\widetilde{S_{it}} = \left(S_{it} + \overline{S}_t\right)/2$ and one minus this value for capital. While this measure is not transitive – the input weights differ across observations – it does allow for a comparison with the same benchmark while still allowing for technology heterogeneity.

To verify robustness, productivity is also estimated using a particular functional form for the production function. The parameters are estimated using two methodologies, one pioneered by Olley and Pakes (1996) and a new one by Ackerberg et al. (2006). Olley and Pakes (1996) invert the investment equation non-parametrically to proxy for unobserved productivity. An intermediary estimation step controls for the non-random sample selection induced by the differing probability of exit for small and large low-productivity firms. In many applications, a large number of zero investment observations have to be omitted when the investment equation is inverted. In the Chinese high-growth context we only observe negative real investment for 1% of continuing firms. However, we do not observe investment directly, but construct it from the capital stock information, which will smooth the investment series.

The Ackerberg et al. (2006) approach avoids this problem by using the value of intermediate inputs in the proxy estimator, which is virtually never zero. All coefficients, both on variable and quasi-fixed inputs, are recovered in the second stage using a GMM estimator. One advantage of this approach is the solid identification results.

To obtain productivity level and growth estimates with the parametric approaches (P), one simply has to replace the input weights in Eqs. (3) and (4) with the estimated sector-specific input elasticity parameters. The productivity level for firm i at time t is calculated as

$$a_{it}^{P} = q_{it} - \hat{\alpha}_L^S l_{it} - \hat{\alpha}_K^S k_{it}, \tag{5}$$

and productivity growth boils down to $a_{it}^P - a_{it-1}^P$. The superscripts on the coefficient estimates indicate that we estimate the production function separately for each industry. To normalize the productivity level estimates, recall that productivity is only a relative concept, and one can simply subtract the average productivity across all firms at time t in the same industry or include industry dummies in the regressions.

3. Data

We utilize firm-level data for the period 1998–2007 that are the product of annual surveys conducted by the National Bureau of Statistics (NBS). The survey includes all industrial firms that are either state-owned, or are non-state firms with sales above 5 million RMB (hereafter referred to as the "above-scale" firms). Industry is defined here to include mining, manufacturing and public utilities.

An important contribution of this paper is the construction of complementary information that is needed to use these data. This includes industry concordances, deflators for all nominal variables, programs to match firms over time, and to construct a real capital stock series. We provide an elaborate online appendix to document all ancillary information, to illustrate patterns in the data, and to show robustness checks for key results.[6] Here we only provide a short introduction.

Table 1 provides descriptive statistics for the sample, including information on the number of firms, total value-added, sales, etc. With a few exceptions, these data aggregate almost perfectly to totals for the same set of variables reported in the Chinese Statistical Yearbook. Totals are also nearly identical to those for firms extracted from the 2004 Census that are either state-owned enterprises (SOEs) or non-SOEs with sales larger than 5 million. Comparison with the full census of firms reveals however that 80% of all industrial firms are excluded from our sample. Fortunately, they account for only a small fraction of economic activity: In 2004, below-scale firms employed 28.8% of the industrial workforce, but produced only 9.9% of output and generated 2.5% of exports.[7]

We use unique numerical IDs to link firms over time. Firms occasionally receive a new ID as a result of restructuring, merger, or acquisition. Where possible, we have aimed to track firms as their boundaries or ownership structure changes, using information on the firm's name, industry, address, etc., to link them. The fraction of firms in a year that can be linked to a firm in the previous year increases over time from 84.5% in the first two years (1998–1999) to 92.2% in the final two years (2006–2007). Overall, 95.9% of all year-to-year matches are constructed using firm IDs, and 4.1% using other information on the firm. These other matches are still important as one-sixth of all firms that are observed for more than one year experience a change in their official ID over the period of analysis.

For the analysis in this paper, we focus only on manufacturing firms. This provides an unbalanced panel of firms that increases in size from 148,685 firms in 1998 to 313,048 in 2007.[8] On average, the annual rate of attrition in our sample is slightly less than 14%. Out of our original sample in 1998 of 148,685 firms, 33,054 firms or just under a quarter, survive through 2007. Exit was more than offset by entry, which averaged nearly 20% per annum. Fig. 1 reports exit and entry rates by ownership type. Noteworthy is the sharp increase in the number of sample firms between 2003 and 2004. This can be attributed to the Industrial Census, and the identification of firms, largely private in ownership, that should have been in the sample in earlier years, but were left out because of a less than perfect business registry.

With the exception of the capital stock, construction of most of the other key variables in our analysis, e.g., gross output, value added, employment and wages, is fairly standard. Details are reported in the online

[5] These equations are for a value added production function. When we use a gross output production function, material input enters similarly as labor input, and the dependent variable is changed accordingly.

[6] http://www.econ.kuleuven.be/public/n07057/China/.

[7] These comparisons are detailed in the Appendix. A comparison using the 1995 Census produces very similar results, which increases our confidence that the NBS decision rule on which firms to include in their annual sample is not introducing any systematic bias in our estimates.

[8] The unit of analysis is the firm, and not the plant, but other information in the survey suggests that more than 95% of all observations in our sample are single-plant firms.

Table 1
Summary statistics on the underlying firm-level data set.

Year	Number of firms	Value added	Sales	Output	Employment	Export	Net value of fixed assets
1998	165,118	1.94	6.41	6.77	56.44	1.08	4.41
1999	162,033	2.16	6.99	7.27	58.05	1.16	4.73
2000	162,883	2.54	8.42	8.57	53.68	1.46	5.18
2001	169,030	2.79	9.24	9.41	52.97	1.61	5.45
2002	181,557	3.30	10.95	11.08	55.21	2.01	5.95
2003	196,222	4.20	14.32	14.23	57.49	2.69	6.61
2004	279,092	6.62	20.43	20.16	66.27	4.05	7.97
2005	271,835	7.22	24.69	25.16	68.96	4.77	8.95
2006	301,961	9.11	31.36	31.66	73.58	6.05	10.58
2007	336,768	11.70	39.97	40.52	78.75	7.34	12.34

Notes: all values are denoted in trillion RMB and employment in millions of workers. All industrial firms are included while the analysis in the paper is limited to firms in the manufacturing sector. A comparison with corresponding values in the China Statistical Yearbook, the China Statistical Abstract, and the 1995 and 2004 Census is in the online Appendix.

Appendix, as well as the construction of deflators (and the complete series) for gross output, material inputs, and capital investment.

Employee compensation includes wages, employee supplementary benefits, unemployment insurance, retirement benefits, health insurance and housing benefits. Reported compensation however appears to underestimate total payments to labor. Labor's share of value added is only 34.2% in our sample, compared to around 55% in the national income accounts. In some of our productivity estimation, we follow Hsieh and Klenow (2009) and inflate wage payments by a constant factor for all firms to work with a wage share consistent with the national average.

Construction of a real capital stock series is difficult for two reasons: first, firms do not report fixed investment; and second, firms report information on the value of their fixed capital stock at original purchase prices. Use of nominal values runs the risk of introducing systematic biases related to a firm's age. We develop a procedure explained in the Appendix that converts estimates at original purchase prices into real values that are comparable across time and firms. Our preferred benchmark estimates imply an increase in (real) capital per worker of 35.9% between 1998 and 2007, or nearly twice the growth rate implied by official statistics. This adjustment lowers estimated productivity growth rates substantially, as the capital share in output in China is extremely high. Investment, which is used in the Olley–Pakes proxy estimator, is then obtained simply from the usual equation of motion for the capital stock.

Finally, some firms have missing observations for variables needed to calculate productivity. This arises either because the information was not originally reported, or because of negative values for variables such as the real capital stock or value added. We further drop all firms with less than 8 employees as they fell under a different legal regime. As a result, 17% of the original firms are dropped from the sample in 1998, but this fraction declines to 6% in each year after 2001.

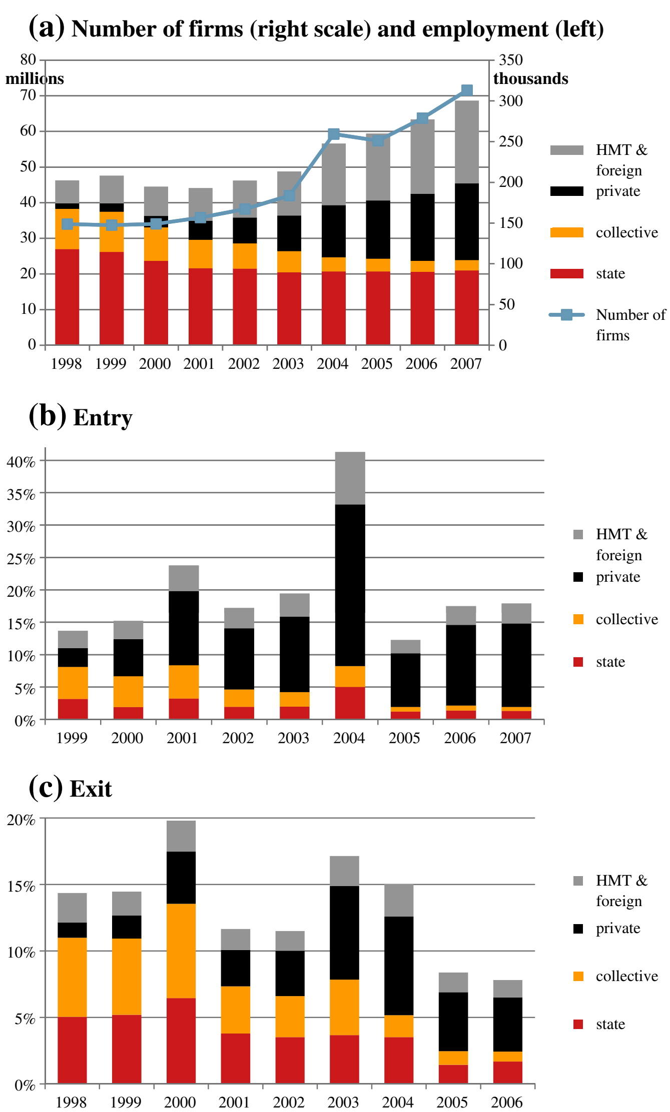

Fig. 1. Employment, entry and exit by ownership type.

4. Results

4.1. Setting the stage

There have been several growth accounting exercises carried out at the macro level that use the standard Solow method to estimate the contributions of resource accumulation and productivity to growth in China since 1978. These studies differ in terms of adjustments they make to the official data, as well as the periods they cover. Several of the more widely cited studies for the entire economy include Bosworth and Collins (2008), Brandt and Zhu (2010). Chow and Li (2000), Holz (2006) and Perkins and Rawski (2008). In general they suggest TFP growth of between 3.5 and 4.0% per annum, or more than half of the growth in output per worker over the last three decades.

The role ascribed to TFP by these studies is significantly larger than Young (2003) finds in his widely cited article examining productivity growth in non-agriculture for the period between 1978 and 1998. Brandt and Zhu (2010) show how revisions by the NBS to earlier GDP figures, and the use of a superior deflator for services raises Young's estimate of TFP by nearly three-quarters. They also illustrate that TFP growth for non-agriculture was a lot higher in the 1998–2007 period than for the two preceding decades.

Non-agriculture includes industry, itself made up of manufacturing and construction, and the service sector. Estimates of service sector productivity are limited, but the general view is that TFP growth has lagged that in industry as a result of less rapid market reform and liberalization. This suggests even higher rates of TFP growth in industry than estimated for non-agriculture alone.

Aggregate TFP growth for the manufacturing sector can occur in three important ways: (i) through firm-level productivity growth, (ii) exit of below-average productivity firms or entry of above-average firms, and (iii) the reallocation of production factors from less to more productive continuing firms. Each factor will be investigated in the following three subsections.

Chinese industry has experienced a significant amount of restructuring. Entry and exit is one dimension, but we might expect transitions between ownership types or between different subsectors of manufacturing to be equally important. In Section 4.5, we carry out decompositions along these lines. We also reconcile our findings for manufacturing with those for the entire non-agricultural sector.

4.2. Firm-level productivity growth

4.2.1. Benchmark estimates

The benchmark estimates for average productivity growth rates by year are depicted in Fig. 2. These measures use the index number formula from Eq. (3) for a value-added production function and the full unbalanced panel of firms. Firm-level growth rates are aggregated using value added weights, averaged over the initial and end year.

Output growth has been extremely rapid in the Chinese manufacturing sector, outstripping the well-documented economy-wide growth rate of 10%, by almost a factor of two. Here we find that firm-level total factor productivity growth has been extremely high as well, rising from an average of 2.9% in 1999 to 14% in 2005, before declining to 11.5% in 2007. The average over the full period stands at 9.6%.[9]

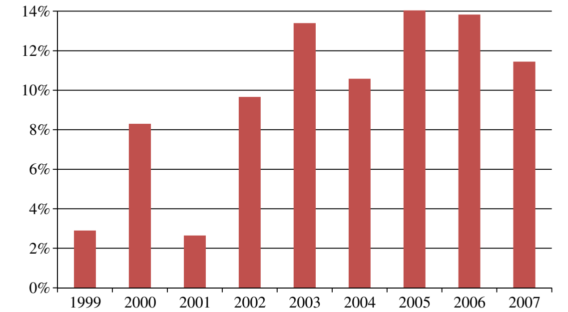

Note:Value-added weighted average of firm-level year-on-year TFP growth estimates obtained using the index number methodology on the full unbalanced sample of Chinese manufacturin firms.

Fig. 2. Benchmark firm-level TFP growth estimates.

In Fig. 3 we present several alternative TFP growth estimates. These measure different objects, but the overarching message of extremely rapid productivity growth prevails. In each case, productivity growth is significantly lower in the first four years of the sample, 1998–2001, than in the period following China's accession to the WTO.

Results in the top-left quadrant of Fig. 3 show that productivity growth rates are almost 4% higher if a parametric production function is assumed and the Olley–Pakes (OP) estimation methodology used. The Ackerberg–Caves–Frazer (ACF) methodology yields even higher estimates. The difference with the index number results is consistent with firms expanding input use more rapidly for those inputs that they use relatively more intensively, which is not unexpected if they are exploiting their comparative advantage.

However, a probably more important difference is the assumption of constant returns to scale for the index number results, while returns to scale are estimated to be decreasing in almost all industries using the parametric methods. The average sum of the labor and capital input elasticities is only 0.80 for OP and barely 0.70 for ACF. In a rapidly growing economy, this will automatically translate into higher productivity growth estimates. These averages seem implausibly low. Measurement error leading to downwardly-biased coefficient estimates is one possible explanation. De Loecker (2007) advances an alternative: if firms have price-setting power, positive markups lead to a downward adjustment of the input elasticities for a production function in value terms. Lacking firm-level information on price levels, we cannot correct for this along the lines of Foster et al. (2008). Jointly estimating the demand function with productivity, as in De Loecker (2007), is also beyond the scope of this paper, and thus we focus attention on the index number estimates.[10]

By construction, productivity growth is much lower for a gross output production function, and the magnitude of the difference is illustrated in the top-right quadrant. Abstracting from weighting issues and assuming that intermediate inputs are proportional to output, the ratio between TFP growths calculated using a gross output and a net output production function should equal the share of value-added in gross output. In the sample, the former averages 0.274 and the latter 0.290. TFP growth on a gross output basis of 2.89% annually is still extremely high—explicit comparisons with other countries follow.

In the bottom-left quadrant, the sample is limited to the balanced panel of firms active throughout the entire sample period. These firms enjoy higher productivity growth prior to WTO entry, but the differences vanish post entry. This is the result of two opposing forces. On the one hand, many of the firms active in 1998 are performing poorly and will exit the industry in subsequent years. Firms in the balanced panel outperform this group by a factor of almost two to one. On the other hand, many of the new entrants are very productive and will enjoy especially rapid productivity growth.

Finally, the unweighted averages of the firm-level productivity growth rates, reported in the bottom-right quadrant, are more than 2% per year lower than the value-added weighted averages.[11] From this we can conclude that, somewhat unusually, large Chinese firms are increasing productivity at a higher than average rate. The positive correlation holds for all three productivity measures (index, OP, and ACF), for all ownership categories, and using value added as well as employment weights.[12] The restructuring of large state-owned firms is one driver for this pattern. In addition, the inclusion rule in the sample based on annual sales implies that some small firms are only included in the sample by virtue of an extremely high productivity level, from which further improvement might be difficult.

To put the extraordinary productivity growth performance documented in Figs. 2 and 3 into perspective, it is important to keep in mind that at least four beneficial factors were jointly at play. First, the sample period covers the cyclical upswing following the Asian financial crisis. Second, China's entry into the WTO and its integration into the world economy lead to exports rising by 25% annually over the sample period, especially benefitting the manufacturing sector. Third, restructuring of SOEs and collectively-owned firms started in earnest in the mid-1990s, and accelerated through the early part of the period we analyze. And fourth, liberalization and competitive pressures in the manufacturing sector exceeded those in most other sectors, like services or utilities.

4.2.2. Robustness checks

An important message from Young (2003) is that measurement issues matter. Using aggregate statistics for the entire non-agriculture sector, he shows that with alternative price deflators and adjustments for input quality the productivity growth estimate is reduced from 3% to 1.4% per year over the 1978 to 1998 period.

While we believe our benchmark estimate uses the most appropriate assumptions, we have explored the sensitivity of our estimates to alternative assumptions.[13] Estimates in Fig. 4 consider four reasons why TFP growth could be biased upward. These follow directly from the definition of productivity growth as output growth minus weighted input growth: (1) price inflation is underestimated.

[9] To construct averages in Figs. 2 and 3, the top and bottom percentiles are dropped, but the effect is minimal.

[10] The online Appendix contains robustness checks for some key results using parametric productivity estimates.

[11] The weight used is the average in the initial and the end year of the period.

[12] There is an abundance of international evidence that firm growth is negatively correlated with size, but the evidence for productivity growth is a lot weaker. Griliches and Mairesse (1983) find a negative correlation between size and productivity growth for firms in the United States and France, but the plant-level evidence for the U.S. in Baily et al. (1992) points to a weak or insignificant relation once they control for firm-level growth.

[13] The benchmark estimate in Fig. 4 is slightly below the 9.6% average reported earlier as the sample has been made consistent across all rows in Fig. 4 and the top and bottom percentiles are not dropped.

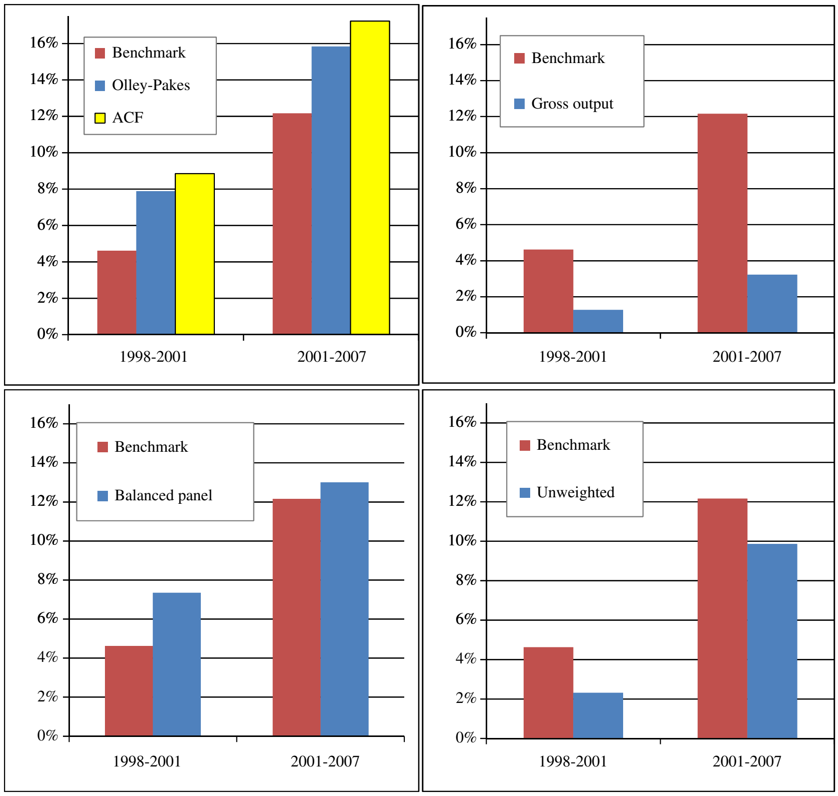

Note: The benchmark estimates always use the index numbers to calculate productivity (instead of the Olley-Pakes parametric methodology); use value added weights (instead of unweighted averages); are calculated on the full sample of firms (instead of the balanced panel); and use a value added production function (instead of gross output). The alternative calculations in the four panels use the alternative assumptions indicated in brackets here.

Fig. 3. Alternative TFP growth estimates.

(2) Labor input growth is underestimated. (3) The weight on capital is too low. (4) Initial capital stocks are overestimated and hence capital growth underestimated. Each of these concerns is addressed in some detail in the working paper version of the paper, see Brandt et al. (2009), but we discuss the bottom line here.

Young (2003) argued that out of convenience firms increasingly tended to report the same value of output at current and constant price levels. As we directly observe the reported output at the two price levels, we can verify this claim directly for the period of our analysis. The fraction of firms in the manufacturing sector that report the same output value using current and constant prices ranges from 16.9% to 21.4%, without any distinctive time trend. Many firms even report higher numbers at constant prices than at current prices, indicating price declines. To verify the sensitivity of our productivity estimates, we also used an aggregate 14-sector deflator, both for output and a corresponding input deflator. The results in Fig. 4 suggest an annual productivity growth that is 0.48% lower per year. This change is smaller than in Young (2003) for the period 1978–1998. The problems he highlights may be more severe for the rapidly expanding service industry.

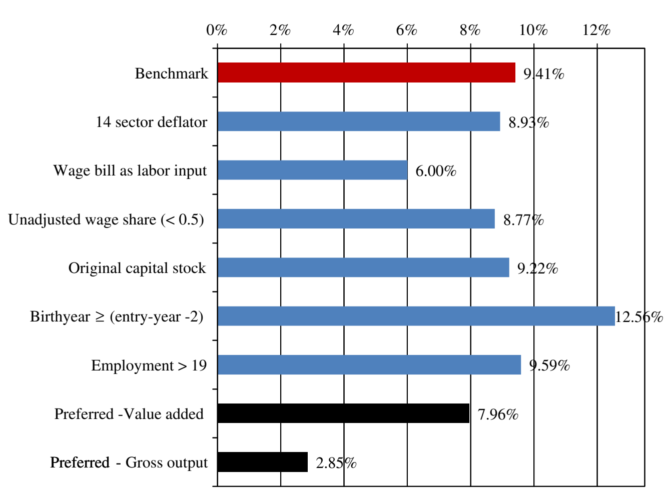

Fig. 4. Robustness of TFP growth estimates to a variety of assumptions.

The second important adjustment that Young (2003) proposes is to control for increases in human capital. Over the twenty-year period he studies, 1978–1998, educational attainment of the labor force increased rapidly, and contributed to output growth. This is less likely to be important in our shorter time period. In the absence of annual information for human capital at the firm-level, we utilize growth in the wage bill as an imperfect proxy for a broader labor input measure encompassing all human capital improvements. This measure grows a lot more rapidly than the absolute number of employees. The results in Fig. 4 illustrate the large impact of this adjustment, which lowers average productivity growth by 3.41% per year.

To the extent that wage increases capture increased hours worked per employee or the higher human capital content in labor input, they should indeed be subtracted from output growth in calculating productivity. To the extent that wages rise because workers are working with more capital now or are able to appropriate rents, the 6.0% now obtained underestimates productivity growth. There is some evidence for this.[14] In addition, more competitive labor markets

[14] For one, the difference between wage growth and employment growth increases markedly as the business cycle became more expansionary and labor markets were getting tighter. The average differential in the two growth rates rises from 2.5% in 1999 to 10.3% in 2006. Newly entering firms also experience higher wage growth than incumbents, presumably in part because they have to compete harder to attract workers away from other firms.

Table 2
Average TFP growth in different countries.

Country	Study	Period	Sector	TFP growth
Firm-level: value added production function				
China		1998–2007	Manufacturing	0.080
Slovenia	De Loecker and Konings (2006)	1994–2000	Manufacturing	0.085
Vietnam	World Bank (2007)	2001–2003	Manufacturing	0.073
U.S.	Baily et al. (1992)	1982–1987	Selected manufacturing sectors	0.031
Chile	Pavcnik (2002)	1979–1986	Manufacturing	0.028
Firm-level: gross output production function				
China		1998–2007	Manufacturing	0.028
Korea	Ahn et al. (2004)	1990–1998	Manufacturing	0.035
Taiwan	Aw et al. (2001)	1986–1991	9 manufacturing sectors	0.021
Mexico	Tybout and Westbrook (1995)	1984–1990	Manufacturing	0.019
U.S.	Haltiwanger (1997)	1982–1987	Manufacturing	0.017
Japan	Ahn et al. (2004)	1994–2001	Manufacturing	0.003

Note: where published productivity growth estimates for other countries are available for several time periods, we have taken those from cyclical expansions to be comparable with the Chinese macroeconomic situation. The Chinese averages are the “preferred” estimates incorporating three adjustments on the benchmark estimates, as at the bottom of Fig. 4.

and a more liberal stand towards migration increasingly allow the best firms to attract the best workers. This permits a more efficient employment of the available labor, but because it is reflected in wages it will not show up anymore as a productivity contribution in the alternative measure.

A third ingredient in the productivity growth calculations is the relative weight on employment and capital input growth. We followed Hsieh and Klenow (2009) and inflated the average wage share to half of the value added to approximate the fraction in the national accounts. However, if labor is actually claiming a much smaller share of value added in manufacturing than in the rest of the economy, one should use the unadjusted wage share, which averages only slightly above one third. The result in Fig. 4 shows that using the unadjusted wage shares lowers TFP by 0.64% per annum, which is as expected given that capital input has grown more rapidly than labor input.

The fourth and final factor with an impact on the productivity growth calculations is the construction of the real capital stock. The measure reported by the NBS – original value at purchase prices – adds undeflated investment flows to undepreciated past investments. In the capital series used in the benchmark estimates, we controlled as accurately as possible for depreciation and capital price inflation in the initial capital stocks and then used a perpetual inventory measure to roll the stock forward. If instead, we directly deflate the original capital stock measure, average productivity growth is estimated to be 0.19% lower. While the average does not change a lot, the relative productivity of firms is more broadly affected. Especially for older or capital-intensive firms, many of whom are state-owned, the choice of capital series matters.

A more radical solution would be to drop all firms with a birth year of more than 2 years prior to the first year they are observed. On this sample, productivity growth is estimated a lot higher. While the problem of estimating the initial capital stock is much reduced, these firms are also a lot younger and higher productivity growth is not unexpected. We explore productivity dynamics over a firm's life-cycle further in the next section. One final robustness check is to eliminate firms with fewer than 20 employees. As the inclusion rule for the sample is size-based, some small firms are likely to be included only because of an extraordinary productivity performance from which further improvement is difficult. The average TFP growth changes in the expected direction, but the magnitude of this effect is small.

Going forward, we adjust the benchmark productivity estimates for the three issues discussed first. As the truth is likely in between our benchmark assumptions and the extreme assumptions used in Fig. 4, we go half way on each adjustment. We use the geometric mean of the detailed and aggregate deflators; the geometric average of labor input growth based on employment and on wage costs; and we only adjust the wage share half of the way to one-half of value-added. Each of these three corrections has a downward effect on productivity growth and our ‘preferred’ productivity growth estimates, at the bottom of Fig. 4, now average 7.96% for a value-added production function and 2.85% on a gross output basis.

We gain further perspective by comparing these numbers with similar statistics for other countries taken from the literature. Averages for both value-added and gross output total factor productivity growth for several countries are reported in Table 2. Where information was available for several sub-periods, we chose cyclical expansions to match the Chinese situation.

Compared to these other countries China's performance remains exemplary, but not off the charts. Its record is comparable to Slovenia and Vietnam, two other transition economies. After the recession in the early 1980s, the record of U.S. plants is also better than commonly assumed, with a total factor productivity growth of 3.1% per year. Using a gross output production function, the Chinese average of 2.85% over the full sample period and 3.2% over the last six years is comparable to the Korean performance over the 1990–1998 period. Results for other countries indicate that a sustained TFP growth in excess of 2% is a rare event, even in cyclical booms.

4.3. Entering and exiting firms

The results in the preceding section average across the growth rates of continuing firms, new entrants, and firms about to exit the industry. For the aggregate productivity record it is also important (i) where in the distribution the new firms are entering and where the old ones are disappearing, and (ii) how productivity growth evolves over a firm's life cycle. These factors are particularly important for China given the high observed rates of firm turnover reported in Fig. 1.

Regression results with firm-level productivity as dependent variable on a set of five post-entry dummies are reported in Table 3. As controls these regressions include a full set of industry-year fixed effects and dummies for ownership type and coastal provinces. Only firms that report a birth year at most two years prior to their appearance in the sample are labeled as entrants.[15] In the first two regressions, columns (1) and (2), the control group includes all firms, including firms that we observe throughout and those that entered earlier or that reported an earlier birth year.

The results are similar whether we use the productivity growth rate or the level as the dependent variable. Firms are found to enter approximately at the average productivity level of incumbents, but in

[15] Recall that the sample only includes firms selling at least 5 million RMB per year. Some new entrants have been producing for some time before they achieve this sales-threshold or before they are ‘discovered’. Improvements in the business registry over time make an ever larger share of newly appearing firms true entrants.

Table 3
Productivity evolution.

(a) Entering firms				(b) Exiting firms			
Relative to	Full sample		Balanced sample		Full sample		Balanced sample
	TFP growth (1)	TFP level (2)	TFP level (3)		TFP growth (4)	TFP level (5)	TFP level (6)
t_0 (entry)		0.031	−0.047	t_0 (exit)	−0.145	−0.141	−0.258
		(0.004)	(0.005)		(0.002)	(0.003)	(0.004)
$t+1$	0.139	0.147	0.073	$t-1$	−0.005	−0.079	−0.169
	(0.003)	(0.004)	(0.005)		(0.003)	(0.003)	(0.004)
$t+2$	0.016	0.130	0.058	$t-2$	−0.042	−0.065	−0.154
	(0.003)	(0.004)	(0.005)		(0.003)	(0.003)	(0.004)
$t+3$	−0.014	0.086	0.017	$t-3$	0.009	−0.019	−0.108
	(0.003)	(0.004)	(0.006)		(0.003)	(0.003)	(0.004)
$t+4$	−0.001	0.053	−0.022	$t-4$	0.024	−0.007	−0.091
	(0.004)	(0.005)	(0.006)		(0.004)	(0.004)	(0.004)

Note: OLS regressions of TFP level and growth rates on dummies for the entry and post-entry years, in columns (1)–(3), and exit and pre-exit years in columns (4)–(6). Province and sector-year fixed effects are included. These results use the index number estimates of productivity; comparable results using the Olley–Pakes and Ackerberg–Caves–Frazer productivity measures are in the Appendix.

their first full year they improve dramatically. One year post-entry they already show a distinct productivity advantage. This difference gradually narrows again in the following years. The growth regression still shows a growth advantage of recent entrants in their second year even though their productivity level advantage declines because surviving entrants appear to be selected more on productivity growth than productivity level.

In the third column, we compare entrants only to the balanced panel of continuing firms, which provides a more stable comparison group. It shows firms entering at a slightly lower productivity level, but rapidly improving productivity makes them overshoot the incumbents' productivity level after their first complete year in the sample. Subsequently they converge quickly to the balanced panel firms and by the third year they are virtually indistinguishable. The relative standing of entrants is less favorable in this selected sample, which omits firms that will exit in future years. Firms that manage to survive over the entire turbulent period are clearly not a random group.

The results in Table 3 use the index number productivity estimates. In the online Appendix, we report comparable results using the two parametric productivity measures. The patterns are extremely similar, only the productivity level differences are even more pronounced. [16] While entrants have below average productivity in their first year, compared to all other firms, their strong productivity record shows them well ahead after four years. Compared to the balanced panel firms, they are estimated to enter at a 19–31% lower productivity level, but after four years they have closed the gap as well. Four years after entry, all productivity growth differences are statistically insignificant.

The overshooting and subsequent decline in productivity levels over the entrants' first few years is largely caused by the pooling of all entry cohorts. On the one hand, new entrants are compared not only to incumbents, but also to earlier entry cohorts, which is a sizeable group. On the other hand, the entry process has been changing over the sample period. In Fig. 5, we plot coefficients like those in column (3), but now estimated year by year.[17] For clarity, we lumped the entry year and the first full year post entry.

The most noteworthy feature is the decline in the initial productivity level of new firms. As the Chinese market liberalized, one could expect the entry process to change. In a dynamic entry model such as Hopenhayn (1992), a lower productivity draw would still lead a firm to enter the industry if fixed entry costs had been reduced. The greater market opportunities, especially after WTO accession in 2001, could also lead to more experimentation and more opportunistic entry.

A second feature is the relatively stable productivity premium over incumbent firms that entrants converge to by their third or fourth year in the sample. The coefficients in Fig. 5 only capture the evolution between surviving firms from each annual entry cohort and the stable group of balanced firms.

If market selection weeds out the worst performing entrants, we expect the average productivity for survivors in subsequent years to improve upon their entry-year productivity. This is indeed the case in every year, but this process runs its course particularly rapidly from 2002 onwards. This is all the more remarkable as the gap between the initial and eventual productivity levels increases over time. Only for the cohorts that entered in 2000-2001 – more uncertain recession years following the Asian financial crisis and preceding China's WTO accession – did the convergence to the stable productivity premium take more than two years.

Productivity level and growth results for exiting firms are reported in the right panel of Table 3. The patterns confirm well with findings for other countries and suggest that firm-exit contributes positively to the aggregate productivity record. In the final year of operation, firms that subsequently exit are a lot less productive than the average

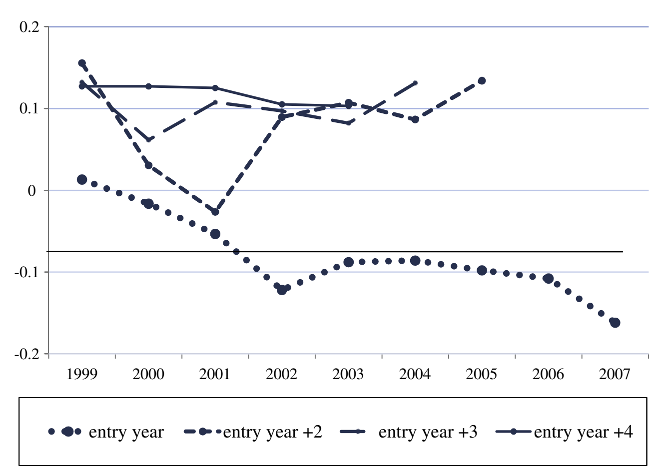

Note: The markers represent coefficient estimates for productivity levels in post-entry years relative to balanced panel firms, comparable to those in column (3) of Table 3. Estimates are performed separately by entry cohort and the two-year moving average of coefficients is plotted. Productivity is estimated using the index number method; only recently-founded firms (at most 2 years ago) are counted as entrants; the year of entry and the first complete post-entry years are lumped together.

Fig. 5. Evolution of relative productivity levels for new entrants.

[16] The lower initial productivity level for entrants using the parametric estimates is again the result of scale economies estimated to be decreasing. In relative terms, it penalizes the smaller average size of entrants. As they rapidly converge in size to the average firm, the productivity premiums also converge to the index number results.

[17] To smooth out the annual variations, we plot the two-year moving average of each coefficient.

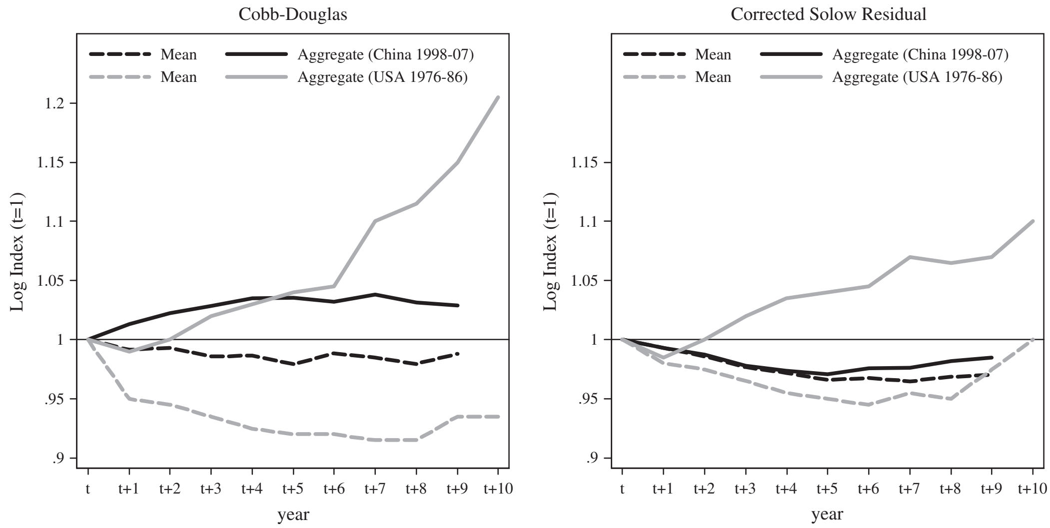

Source: U.S. results are from Figure 1 in Bartelsman and Dhrymes (1998); for consistency, we apply the same methodology for Chinese firms.

Notes: Cobb-Douglas productivity measures are calculated from a production function estimated by least squares. Solow Residuals refer to a Tornqvist index number using the wage share as weight on labor input and the material share for materials; in the United States a "capital share" is observed, while in China constant returns to scale are assumed. Both productivity measures are purged from year and 4-digit industry effects, and residuals are used. The dashed lines are unweighted averages, the solid lines weigh plants or firms by their aggregate input ($L^{\alpha_L}M^{\alpha_M}K^{1-\alpha_L-\alpha_M}$) share.

Fig. 6. Productivity contribution of resource reallocation between active firms.

surviving firm. The orders of magnitude vary from negative 13.5% for the index number results to negative 35–39% for the parametric methods (in the Appendix).

These differences are very large and they open up almost entirely in the firms' last years of operation. Three and four years prior to exit, productivity growth differences for firms about to exit are mostly small and positive, and the productivity level differences do not display any pattern. From then onwards, however, their productivity starts to deteriorate very rapidly. Differences with the balanced panel of firms are even larger and a negative productivity gap is already apparent four years prior to exit.

4.4. Resource reallocation between active firms

The previous results illustrate that resource reallocation associated with firm turnover contributes positively to aggregate productivity. The replacement of exiting firms with new entrants has an immediate positive effect and the higher growth rate of entrants boosts aggregate growth yet further in subsequent years. We now look at the movement of production factors between active firms to see whether such reallocation has similar positive effects for aggregate productivity.

We do this by applying the methodology that Bartelsman and Dhrymes (1998) used for U.S. plants to the Chinese firms in our data set. In an efficient factor market, we expect that resources are put to their most productive use. Aggregate growth can result from above-average productive plants gradually commanding a greater share of inputs over time, which is exactly the case in the United States. We reproduce their estimates, with comparable results for China in Fig. 6.

The graphs show the evolution of the unweighted productivity average (dashed lines) and the aggregate productivity level (solid lines), both normalized by their level in the initial year. Log productivity is purged from time and four-digit industry effects, before applying the exponential function and aggregating. The weight to construct the aggregate level is the input aggregate, $L_{it}^{\alpha_L}M_{it}^{\alpha_M}K_{it}^{1-\alpha_L-\alpha_M}$, but results are similar for output weights. The gap between the two lines is the extent to which the weight is tilted towards more productive plants.[18] Results in the left graph use the residuals from a Cobb–Douglas production function estimated with least squares. In the right graph, a Törnqvist index (Solow residual) for productivity is constructed using factor shares as input weights.[19]

The U.S. pattern (in gray) is remarkable. The unweighted average drops by 0.45% annually if the Cobb–Douglas productivity measures are used and is virtually unchanged for the corrected Solow Residual. A decline reflects less productivity dispersion as the average of the log-productivity measures are normalized to zero in each year and the exponential function is convex. The weighted average, on the other hand, corresponds to the evolution of aggregate productivity and increases annually by 1.9% in the case of the Cobb–Douglas based measures and by 1.0% for the Solow Residual. Note that these effects are solely the result of changing input weights and are in addition to 0.8% and 0.2% average annual productivity increases that are common to all plants. The different evolution between the unweighted and weighted averages can be the results of larger plants increasing productivity more rapidly or of plants with the highest productivity level increasing their weight over time. Other evidence pointing to large persistence in relative productivity levels leads Bartelsman and Dhrymes (1998) to prefer the latter explanation.

The black lines represent the corresponding patterns for China. The decline in the unweighted average is similar to the U.S. New firms entering with a productivity level closer to the mean than exiting firms contribute to the decline. The big difference between the U.S. and

[18] The productivity decomposition in Olley and Pakes (1996) for the U.S. telecommunications equipment illustrates the same phenomenon. The gap between the two lines in Fig. 6 corresponds to the importance of the covariance term in the popular Olley–Pakes decomposition.

[19] As we cannot modify the U.S. results which use confidential census data, we had to used the same two productivity measures for China for comparability. Bartelsman and Dhrymes (1998) pool plants from three related two-digit industries and their Cobb–Douglas results use the same input coefficient estimates for all. We only use a subset from their longer time series, and start in 1976 after the U.S. recession. For China, we re-estimate the production function for each two-digit industry.

Chinese patterns is the vastly smaller contribution to the aggregate productivity growth of inputs shifting to more productive firms. For the Cobb–Douglas results, the reallocation of resources provides some positive effect, but the cumulative effect over nine years is barely 4%—the difference between the dashed and solid black lines on the left graph. For the Corrected Solow Residual, shifts in input weights are virtually unrelated to productivity differences; the solid line tracks the dashed line almost perfectly. This discrepancy is particularly revealing given the much higher firm-level growth rates observed in China, and the important positive productivity impact of reallocation at the extensive margin.

This is an important finding in its own right, but of even greater relevance in the Chinese context. Hsieh and Klenow (2009) report evidence of a very large dispersion in TFP for China. They also estimate much larger differences in marginal products of labor and capital than in the United States. More efficient allocation of resources could potentially lead to very large aggregate TFP gains: in counterfactual simulations they show a potential TFP-boost of 30 to 50 % if the differences in China were to be reduced to U.S. levels. The evidence in Fig. 6 shows that despite very large *potential* gains from reallocation of resources, the *realized* gains in China are much smaller than those in the United States.

In sum, we find that aggregate growth in China is constrained by extremely limited efficiency-enhancing input reallocations. This confirms a wealth of anecdotal evidence that describes the continued support of less efficient manufacturing firms with significant ties to the state.

4.5. From the micro to the macro level

4.5.1. Growth decomposition

We wish to link our firm-level, micro results with the literature on Chinese productivity that uses macro-level data. We also want to identify the types of heterogeneity most important to the aggregate evolution of productivity. We begin by aggregating inputs and output over our sample of firms to obtain manufacturing totals, and then decompose value-added growth into the contributions of capital and labor input growth, and a productivity residual.[20] This residual will then be further decomposed to investigate the contribution of a number of factors on aggregate TFP.

The first bar in Fig. 7 contains the Solow growth decomposition using all firms and the full sample period. Between 1998 and 2007, value-added in manufacturing grew at more than 22% per annum. Capital accumulation and labor (quality-adjusted) input growth accounted for 5.1% and 4.5% of annual output growth, respectively, or 43% of the total. The remainder, or 57%, can be attributed to productivity growth. The contribution of TFP is slightly higher than that found in the aggregate growth accounting exercises cited in Section 4.1.

Our estimate of aggregate productivity growth for manufacturing of 13.4% per annum exceeds the firm-level growth rates reported earlier. This is to be expected as (limited) resource reallocation towards more productive firms and (especially) the large productivity gap between entering and exiting firms make positive contributions. However, our residual growth rate is significantly higher than the 4.4% per annum estimated by Brandt and Zhu (2010) for the same period for the entire non-agricultural sector using macro data. The difference can be understood as follows.

First, the lower TFP for all of non-agriculture is linked in large part to the much slower TFP growth in services and the construction sector, which combined are larger than manufacturing. Restructuring in many services such as finance, retail, distribution, has been much slower than in manufacturing. Furthermore, the service sector had to

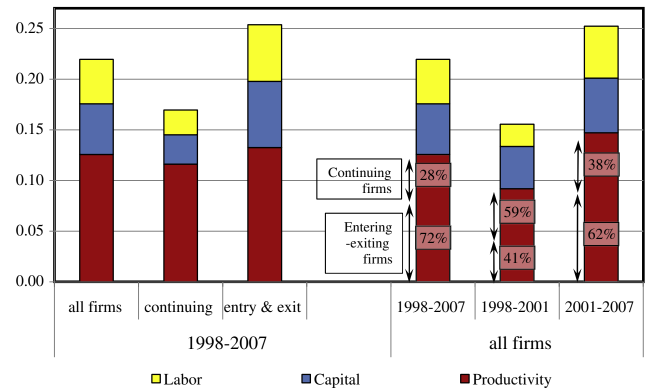

Note: The bars indicate the different components in a standard Solow growth decomposition showing average annual changes in log points. The weight on labor input growth is 0.42, the average of the firm-level sample average and that in the national accounts. The percentages on the right decompose productivity growth into changes at intensive and extensive margins, as discussed in the text.

Fig. 7. Output and productivity growth decompositions.

absorb a significant portion of the nearly 50 million workers laid off from the state sector during this period. Estimates of TFP growth in services in Brandt et al. (2010) for the period 1990–2007 are one-third to one-fourth of those for the secondary sector. Second, during our sample period, value-added in industry in the national income accounts grows less rapidly than estimates obtained from aggregating up firm-level data; both series for industry are reported in China's Statistical Yearbook.[21] This shows up as lower productivity growth at the macro level. And third, the manufacturing sector has been a huge beneficiary of the enormous infrastructure investments, which has helped to raise productivity. This addition to the capital stock can be subtracted in aggregate TFP estimates, but not in the firm-level estimates.

In the next two bars of Fig. 7, we depict the growth decomposition separately for the balanced panel of firms, and the newly entered firms combined with exiters. Output growth for the balanced panel of firms is significantly lower – 0.170 (18.5%) versus 0.254 (28.9%) annually – but this is almost entirely due to the much greater input factor mobilization by new entrants. The contributions to output growth of both capital and labor additions are more than twice as large for net entrants as they are for the balanced panel of firms.

The two far right columns in Fig. 7 show the growth decompositions separately for the years 1998–2001, and 2001–2007. Comparing the pre and post-WTO periods, annual output growth increased tremendously and this was accompanied by a nearly proportional increase in productivity growth. We do not show further decompositions, but the productivity growth increase in the post-WTO period is more pronounced for firms surviving between 2001 and 2007 than at the extensive margin, consistent with our observation in Section 4.3 that later entry cohorts entered the productivity distribution at lower points. At the same time, the tendency for new factor inputs to flow towards new entrants also diminishes over time.[22]

Earlier results in Section 4.4 pointed to a very modest role for resource reallocation between active firms. At the same time, marginal additions to input factors have been absorbed predominantly by

[20] Aggregating labor and capital overall firms requires that we use the same input weights for all observations. Note that this only affects the breakdown between the capital and labor inputs, and not the size of the productivity residual. Other adjustments we make to calculate aggregate TFP are in line with those discussed in the context of our preferred productivity measures in Section 4.2.

[21] In the construction of the national income data, the National Bureau of Statistics makes several adjustments to value-added aggregated from the micro-data. Unfortunately, we do not know much about these adjustments. A higher coverage of total manufacturing output by the sample of above-scale firms, notably in 2003, is one potential explanation for the diverging series. Higher non-manufacturing output by firms classified in the manufacturing sector by their main line of business is a second.

[22] Continuing firms lowered total labor input and received only one third of new capital in the 1998–2001 period, even though they were responsible for 72% of value added. In contrast, over the period between 2001 and 2007 they added workers and received 40% of new capital, even though their value-added share fell to 59%.

new entrants. The fraction of productivity growth coming from continuing firms versus net entry is a combination of the within-firm productivity growth rates of both groups, the relative productivity levels of entrants to incumbents, and the evolving output shares of both groups.

We implement the decomposition pioneered by Baily et al. (1992) who defined the aggregate productivity level as $\ln TFP_t = \sum_i \theta_{it}^Q \ln TFP_{it}$ and aggregate growth rate as the time difference of this object. This can readily be decomposed into the contribution of continuing and other firms by splitting the sums. As the aggregate share of both groups might change over time, Haltiwanger (1997) illustrated that it is preferable to normalize all terms by a constant and it is intuitive to use $\ln TFP_{t-1}$. For comparability with U.S. results, we use gross output shares as weights and the same methodology to construct the firm-level productivity estimates.

The percentage contribution of both groups in each period is indicated on the right bars in Fig. 7. Over the entire sample period, the contribution of net entrants to total productivity growth is 72%, even though their output share is only 59%. The comparable share for the U.S. between 1977 and 1987 was 26% of productivity growth. Foster et al. (2001) construct various decompositions for U.S. manufacturing and note that the relative contribution of net entry was one of the most robust patterns. Fernandes (2007) finds an even smaller role for the net entry contribution in Chile. As expected, the contribution of net entry is smaller over shorter time horizons as we illustrated that it takes a few years for new entrants to converge in productivity level to incumbents and to build up their output share. Still, even over the short three-year period prior to WTO-entry, net entry accounted for 41% of productivity growth, a lot higher than the 28% output share and a lot higher than in five-year periods in the United States.

There are a number of alternative decompositions. Petrin and Levinsohn (2006) in particular have argued that the above one is a good approximation of welfare growth only when the allocation of input factors is inefficient, i.e. if the marginal products of resources are not equalized across firms. Only in that case do resource reallocations and their corresponding output share changes contribute to aggregate welfare.

Hsieh and Klenow (2009) provide evidence that inefficient factor markets are an important phenomenon in China, nonetheless, it is instructive to take a look at the Petrin–Levinsohn decomposition that captures aggregate welfare changes assuming perfect factor markets. The two most important adjustments are the exclusion of share changes for continuing firms, only counting within-firm productivity growth, and the adjustment of the entrants' productivity level for the average productivity gap at entry, and vice versa for exiting firms. A complete comparison with the U.S. results for both decompositions is in the Appendix, where we have used published statistics from Foster et al. (2001) to construct an approximation to the Petrin–Levinsohn decomposition.

Even using this alternative decomposition the contribution of net entry is far more important in China than in the United States. On the one hand, the larger productivity gap between entrants and exiting firms in China is now adjusted for and partly taken out of the net entry contribution. On the other hand, the substantial output share increases for the most productive incumbents in the United States do not count anymore in the productivity contribution of continuing plants. These two opposing forces still lead to a larger contribution of net entry in China, of more than 50% of aggregate TFP growth, compared to 14% in the United States.

The absolute growth rates underscore the importance of net entry for China. Over the full sample period, the productivity contribution for continuing firms in China amounts to 1.1% per year, using either decomposition. For the 1977–1987 decade in the United States this ranges from 0.8% to 0.7% over the two decompositions. The absolute level of aggregate productivity growth generated through net entry, on the other hand, differs markedly. It is 2.9% per year for China versus 0.3% for the United States using the BHC decomposition and 1.2% versus 0.1% using the PL decomposition. In both cases, net entry contributes ten times as much in China. The same holds over the sub-periods, except in the PL decomposition for 1998–2001 where the ratio is only 3.5.

4.5.2. State versus non-state

In the Chinese context, the state versus non-state distinction figures prominently. At the macro level, Brandt and Zhu (2010) find significant productivity differences between these two components of the non-agriculture sector, with much of it related to capital allocation. Differences widened between 1978 and 1998 as productivity growth in the non-state sector averaged 4.6% per annum compared to only 1.5% in the state sector. As a result of restructuring in the state sector, which began in the mid 1990s, the gap in productivity growth has narrowed over time, with nearly identical rates of productivity growth in the two sectors between 1998 and 2007. Still, significant differences in the productivity levels remain, implying a potential role for productivity-enhancing reallocation of inputs.

Regressions of firm-level productivity or productivity growth on ownership dummies for the 1998–2007 period confirm that there are still significant differences between state and non-state firms in our sample. On average, state-owned firms were 27% less productive, while their annual productivity growth was 4.6% lower. These simple comparisons are potentially misleading however. Non-random ownership transitions, differential entry and exit rates, and differences in productivity gaps at entry or exit all play a role in determining total productivity growth for the two sectors. In addition, the extent to which new input factors are allocated to the most productive enterprises might also differ by ownership.

In fact, the aggregate growth decomposition in Fig. 8, which is similar to Fig. 7 but now performed separately by ownership category, reveals slightly higher productivity growth among continuing state-owned firms compared to either private or foreign firms. Between 1998 and 2007, annual productivity growth was 12.5% for SOEs, compared to 11.3% for private firms and 11.8% for foreign. The next two bars further indicate that productivity growth is even higher for continuing firms that changed ownership from state to private or foreign and it is lowest for firms making the reverse transition. These two categories are indicated on the graph as 'Privatized' or 'Nationalized' firms, and their respective productivity growth rates are 13.5% and 10.7% per year. Reallocation of inputs towards more productive firms within each category might have contributed to these growth rates. While this was not an important factor for the overall manufacturing sector, as demonstrated in Section 4.4, it can be important in some individual categories.

The same growth decomposition is also performed comparing entering and exiting firms within each ownership category separately, on the right in Fig. 8. Value-added growth differences for the three

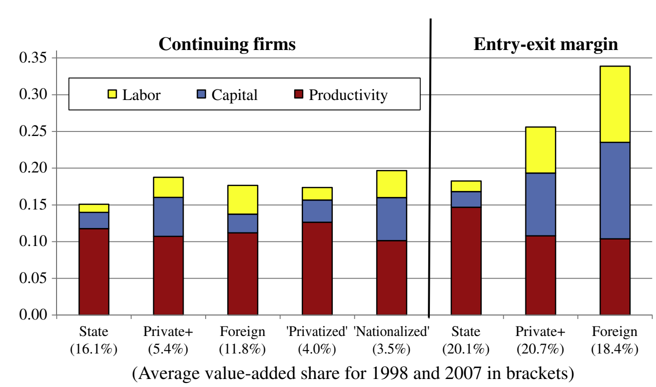

Note: Private+ combines collective-owned firms (and other hybrids) and purely private firms. The nationalized firms are those switching to majority state ownership, the privatized category are those switching from state to one of the other two or switches within the private cateogory.

Fig. 8. Growth decomposition by ownership categories (1998–2007).

groups – the sum of the three components – are clearly much larger than productivity differences. The relative success in attracting new input factors determined relative growth rates. New state firms that appeared between 1998 and 2007 were able to produce almost five times as much value-added as disappearing state firms, even though their real capital stock only grew marginally and employment was a quarter lower.[23] As a result, almost the entire output growth at the entry–exit margin for state firms is attributed to productivity growth. An important contributing factor is that exiting state firms are particularly unproductive.

Despite broadly similar rates of TFP growth across all eight categories, there are sizeable differences in value-added growth. For continuing firms this averages 16.3% for SOEs, compared to 20.6% and 19.3% for private and foreign firms, respectively. At the extensive margin, total value added of newly entered private firms grows 29.2% per year on average, which over the full period represents a ten-fold increase on the 1998 level. The output increase of new foreign-owned firms is even more rapid, but from a lower base. These differences in output growth are entirely driven by differences in input mobilization. In fact, the three categories that acquire the least additional resources – continuing, privatized, and new state firms – have been forced to rely most on productivity growth to boost output.

In the absence of much productivity-enhancing input reallocations between active firms, it matters greatly which entrants are able to attract marginal inputs. Even though productivity growth for state firms is relatively high, this does not mean that they should be attracting more resources. Fig. A.1 in the Appendix illustrates that the average productivity gap between state and private firms over our sample period was still large in most industries. As such, the potential for productivity growth through the reallocation of resources away from state firms is still there. In China, this potential tends to be realized through extensive margin changes. The least productive firms exit and are replaced by comparatively better firms. Input movements between continuing firms are less important, except for those that are accompanied by ownership changes from state to private or foreign, which are associated with high TFP growth.

4.5.3. Sectoral heterogeneity

Finally, we illustrate the distribution of sectoral contributions to aggregate TFP growth with a sunrise diagram pioneered in a famous paper by Harberger (1998). Mimicking an inverted Lorenz curve, all 2-digit manufacturing sectors are ranked left to right by the share of their contribution to TFP growth normalized by their value added share. Fig. 9 then traces cumulative productivity growth over all industries on the vertical axis against cumulative value added on the horizontal axis. By construction, this line lies above the straight line that would represent aggregate manufacturing growth of 13.4% per year over the entire sample period if all industries contributed in proportion to their value added share. A large gap between the two lines would indicate most productivity growth is concentrated in a few sectors.

Compared with corresponding U.S. graphs (over five year intervals) in Harberger (1998), several differences stand out. First, the Chinese curve is extremely flat. Even though the cumulative productivity line in the United States tends to lay closer to the straight line in periods of higher productivity growth, the smallest group of sectors that are able to produce the entire aggregate productivity growth never account for more than 55% of value added. The remaining sectors generate offsetting positive and negative contributions. In many 5-year intervals for the United States, this share is even below 25%. In China, sectors accounting for at least 85%

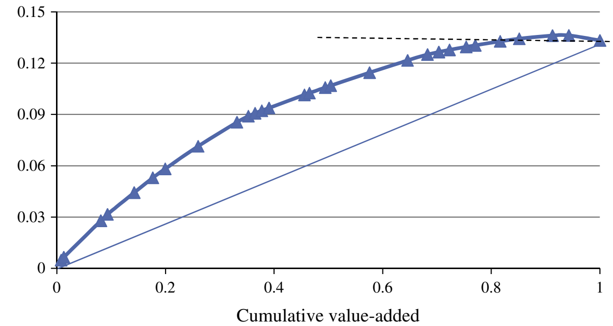

Note: 2-digit sectorsare ordered from high to low productivity growth from left to right. The line plots the cumulative contributions to manufacturing productivity growth against cumulative value added.

Fig. 9. Harberger sunrise diagram for cumulative productivity growth.

of value added are required to generate aggregate TFP growth. This is indicated on the graph by the intersection of the horizontal dashed line and the cumulative TFP growth line.[24]

Second, there is no overshooting. Only a single sector in China, tobacco processing (16), made a negative contribution to aggregate productivity growth over the full sample period. Several U.S. industries always detract from aggregate productivity growth, often accounting for almost half of manufacturing value-added. As a result, the U.S. curves rise steeply above the straight line at first to a level far above the endpoint and decline in the right part.

Third, only four Chinese industries contribute more than twice as much to aggregate productivity growth as their value-added share and two of them are very small.[25] Much of the area between the two lines in Fig. 9 is due to two sectors, transport equipment (37) and ordinary machinery (35), which are both large and experience very rapid TFP growth. Electric equipment (39) and food processing (13) stand out as large sectors with relatively weak productivity growth. More importantly, 18 of the 29 industries contribute almost proportionally to aggregate productivity growth, with a contribution between one-half and double of their value-added share.

To interpret the difference between the U.S. and Chinese patterns, Harberger's (1998) yeast versus mushroom distinction is useful: “yeast causes bread to expand very evenly, like a balloon being filled with air, while mushrooms have the habit of popping up, almost overnight, in a fashion that is not easy to predict.” He interpreted his mushroom-like figures for the United States as “the outcome of real cost reductions stemming from 1001 different causes.” Innovation-driven productivity growth comes to mind. He also expected the yeast process to fit best “with broad and general externalities, such as growth in knowledge or human capital, or brought about by economies of scale tied to the scale of the economy.” This seems consistent with the important role of broad based market liberalization in China. High rates of capital investment, imitation, and knowledge absorption from abroad have made high rates of productivity growth feasible in a wide range of sectors.

5. Conclusions

rawing on an unbalanced panel of firms that covers most of China's manufacturing sector, the purpose of this paper has been to examine the absolute size and the dynamics of productivity growth over a

[23] The higher labor input for these firms was entirely due to real wage increases that are weighted by one half. While it is possible that some restructured state firms exit and subsequently re-enter the sample with a different ID, this does not change the productivity versus input growth decomposition. It would only lead to classification of these firms from the entry–exit margin to the continuing firms group.

[24] Without the correction for rising wage costs in the preferred productivity estimate, the curve for China would be even flatter. Over the post-WTO period (2001–07), sectoral differentials in productivity growth are further diminished.

[25] These are, in descending order, production of timber (20), furniture (21), transport equipment (37), and office machinery (41). Their cumulative share of value added in 1998 was 9.3%, while they were responsible for almost one quarter of manufacturing productivity growth.

period that spans China's entry into the WTO. In our analysis, we have been especially attentive to a host of data and methodological issues.

By all indications, productivity growth has been very rapid, a finding that appears to be robust to a host of measurement issues. This finding is in sharp contrast to alternative perspectives such as Young's (2003) that suggest modest productivity growth outside of agriculture. Improvements in productivity of “continuing firms”, either as a result of restructuring efforts, or investments in capability building are an important part of the picture. Equally, if not more important, are the gains to creative destruction, i.e. entry and exit, that China's decentralized reforms have increasingly allowed. For our preferred estimate, the weighted average annual productivity growth for incumbents is 2.85% for a gross output production function and 7.96% for a value added production function over the period 1998–2007. This is among the highest compared to other countries. Productivity growth at the industry level is even higher, reflecting the dynamic force of creative destruction. Over the entire period, net entry accounts for over two-thirds of total TFP growth, with growth of entrants after entry making an important contribution to this. In all, TFP growth dominates input accumulation as a source of output growth and is responsible for two thirds of labor productivity growth.

These findings have important implications, not in the least for government policy.

First, even with declining opportunities for growth through accumulation of additional inputs, the high TFP growth estimates suggest that we can expect robust growth in manufacturing in the near future. Moreover, TFP growth will help to maintain profitability in the sector in the face of rising labor and other input costs.

Second, increasing competitive pressure and the absorption of foreign technology are often mentioned as drivers of TFP growth. Learning on the part of continuing firms as well as new entrants is critical to taking advantage of these sources of growth. For entrants, there are two dimensions to learning: first, identifying new opportunities that allow successful entry; and second, improvements in productivity subsequent to entry. Policies that facilitate learning among both kinds of firms are the key to sustained growth in the medium term. As Chinese firms narrow the technology gap with advanced countries, more of the learning will have to come from within firms.

Third, despite the dynamism we document, our results also point to continued constraints on the growth of some of the most productive of firms. Problems in the allocation of credit and biases in favor of larger firms with state-sector connections are potential reasons for this. With growth prospects on the extensive margin limited, new reforms to enhance efficient resource allocation still provide important growth potential. A policy of liberalizing entry and facilitating exit has already played an important role in reallocating resources to new firms. Constraints that sustain the remaining productivity differences among existing firms, including between the state and non-state sectors, will have to be tackled next. More work is needed to identify the exact nature of the constraints impeding reallocation, and how they can be best removed.

Appendix A. Supplementary data

Supplementary data to this article can be found online at doi:10.1016/j.jdeveco.2011.02.002.

References

Ackerberg, D.A., Caves, K., Frazer, G., 2006. Structural Identification of Production Functions,. Working Paper December.
Ahn, S., Fukao, K., Kwon, H.U., 2004. The internationalization and performance of Korean and Japanese firms: an empirical analysis based on micro-data. Seoul Journal of Economics 17 (4), 439–482 Winter.
Aw, B.Y., Chen, X., Roberts, M.J., 2001. Firm-level evidence on productivity differentials and turnover in Taiwanese manufacturing. Journal of Development Economics 66 (1), 51–86 October.
Bai, C.-E., Hsieh, C.-T., Qian, Y., 2006. The return to capital in China. Brookings Papers on Economic Activity 2, 61–88.
Baily, M.N., Hulten, C., Campbell, D., 1992. Productivity dynamics in manufacturing plants. Brookings Papers: Microeconomics 4 (1), 187–267.
Bartelsman, E.J., Dhrymes, P.J., 1998. Productivity dynamics: U.S. manufacturing plants, 1972–1986. Journal of Productivity Analysis 9 (1), 5–34 January.
Bosworth, B., Collins, S.M., 2008. Accounting for growth: comparing China and India. The Journal of Economic Perspectives 22 (1), 45–66.
Brandt, L., Zhu, X., 2010. Accounting for China's Growth. Mimeo. University of Toronto.
Brandt, L., Rawski, T.G., Sutton, J., 2008. Industrial development in China. In: Brandt, L., Rawski, T.G. (Eds.), China's Great Economic Transformation. Cambridge University Press, New York.
Brandt, L., Van Biesebroeck, J., Zhang, Y., 2009. Creative Accounting or Creative Destruction? Firm-level Productivity Growth in Chinese Manufacturing. NBER Working Paper No. 15152, July.
Brandt, L., Tombe, T., Xiaodong Zhu, 2010. Factor Market Distortions Across Time, Space and Sectors in China. Mimeo. University of Toronto.
Caves, D.W., Christensen, L.R., Diewert, E.W., 1982a. The economic theory of index numbers and the measurement of input, output, and productivity. Econometrica 50 (6), 1393–1414 Nov.
Caves, D.W., Christensen, L.R., Diewert, E.W., 1982b. Multilateral comparisons of output, input, and productivity using superlative index numbers. The Economic Journal 92 (365), 73–86 March.
Chow, G.C., Li, K.-W., 2000. China's economic growth: 1952–2010. Economic Development and Cultural Change 51 (1), 247–256.
De Loecker, J., 2007. Product Differentiation, Multi-product Firms and Estimating the Impact of Trade Liberalization on Productivity. NBER Working Paper No. 13155, June.
De Loecker, J., Konings, J., 2006. Job reallocation and productivity growth in an emerging economy. Evidence from Slovenian manufacturing. European Journal of Political Economy 22, 388–408.
Dougherty, S., Herd, R., He, P., 2007. Has a private sector emerged in China's industry? Evidence from a quarter of a million Chinese firms. China Economic Review 18, 309–334.
Fernandes, A.M., 2007. Trade policy, trade volumes, and plant-level productivity in Colombian manufacturing industries. Journal of International Economics 71 (1), 52–71 March.
Foster, L., Haltiwanger, J., Krizan, C., 2001. Aggregate productivity growth: Lessons from microeconomic evidence. In: Hulten, C.R., Dean, E.R., Harper, M.J. (Eds.), New Developments in Productivity Analysis, NBER Studies in Income and Wealth. University of Chicago Press, Chicago, IL.
Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, firm turnover, and efficiency: selection on productivity or profitability? The American Economic Review 98 (1), 394–425 March.
Griliches, Z., Mairesse, J., 1983. Comparing productivity growth: an exploration of French and U.S industrial and firm data. European Economic Review 21, 89–119 March–April.
Groves, T., Hong, Y., McMillan, J., Naughton, B., 1994. Autonomy and incentives in Chinese state enterprises. Quarterly Journal of Economics 109 (1), 183–209.
Haltiwanger, J.C., 1997. Measuring and analyzing aggregate fluctuations: the importance of building from microeconomic evidence. Federal Reserve Bank St. Louis Review 79 (3), 55–77 May–June.
Harberger, A.C., 1998. A vision of the growth process. The American Economic Review 88 (1), 1–32 March.
Holz, C. “Measuring Chinese Productivity Growth, 1952–2005.” Mimeo. Hong Kong University of Science and Technology. 2006.
Hopenhayn, H.A., 1992. Entry, exit, and firm dynamics in long run equilibrium. Econometrica 60 (5), 1127–1150 September.
Hsieh, C.-T., Klenow, P.J., 2009. Misallocation and manufacturing TFP in China and India. Quarterly Journal of Economics 74 (4), 1403–1448 November.
Jefferson, G., Rawski, T.G., 1994. Enterprise reform in Chinese industry. The Journal of Economic Perspectives 8 (2), 47–70.
Jefferson, G.H., Rawski, T.G., Zhang, Y., 2008. Productivity growth and convergence across China's industrial economy. Journal of Chinese Economic and Business Studies 6 (2), 121–140 May.
Olley, G.S., Pakes, A., 1996. The dynamics of productivity in the telecommunications equipment industry. Econometrica 64 (6), 1263–1297.
Park, A., Yang, D., Shi, X., Jiang, Y., 2010. Exporting and firm performance: Chinese exporters and the Asian financial crisis. The Review of Economics and Statistics 92 (4), 833–842 November.
Pavcnik, N., 2002. Trade liberalization, exit, and productivity improvement: evidence from Chilean plants. Review of Economic Studies 69 (1), 245–276.
Perkins, D.H., Rawski, T.G., 2008. Forecasting China's economic growth. In: Brandt, L., Rawski, T.G. (Eds.), China's Great Economic Transformation. Cambridge University Press, NY.
Petrin, A., Levinsohn, J., 2006. Measuring aggregate productivity growth using plant-level data. Working Paper, University of Minnesota, Department of Economics, April.
Tybout, J.R., Westbrook, M.D., 1995. Trade liberalization and the dimensions of efficiency change in Mexican manufacturing industries. Journal of International Economics 39, 53–78.
Van Biesebroeck, J., 2007. Robustness of productivity estimates. Journal of Industrial Economics 55 (3), 529–569 September.
Van Biesebroeck, J., 2008. The sensitivity of productivity estimates: revisiting three important productivity debates. Journal of Business and Economic Statistics 26 (3), 321–338 July.
World Bank, 2007. World Bank. *Vietnam Development Report 2006: Business*. Joint Donor Report to Vietnam Consultative Group Meeting, Hanoi. December.
Young, A., 2003. Gold into base metals: productivity growth in the People's Republic of China during the reform period. Journal of Political Economy 111 (6), 1220–1261.
Zheng, Jinghai, Bigsten, Arne, Hu, Angang, 2006. Can China's growth be sustained? A productivity perspective. Working Paper in Economics, Gothenburg University, Department of Economics.