



TI 2009-043/3

Tinbergen Institute Discussion Paper

Forecasting Aggregate Productivity using Information from Firm-Level Data

*Eric J. Bartelsman**

Zoltan Wolf

VU University Amsterdam, and Tinbergen Institute;

** IZA Bonn.*

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at
<http://www.tinbergen.nl>.

Forecasting aggregate productivity using information from firm-level data

Eric J. Bartelsman^a, Zoltán Wolf^{b,1,2}

^a Vrije Universiteit Amsterdam, Tinbergen Institute and IZA Bonn

^b Vrije Universiteit Amsterdam and Tinbergen Institute

April 2009

¹*Address for correspondence:* Free University Amsterdam, Department of Economics, Room 2E-71, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. Tel: +31 20 598 6030. Fax +31 20 598 9870. *E-mail:* zwolf@feweb.vu.nl.

²We acknowledge the remote access facilities provided by Statistics Netherlands. We thank Dave Reifschneider of the Federal Reserve Board, Jerome Henry of the ECB and Joris de Wind of the DNB for providing information on current central bank productivity forecasting practices. Further thanks for comments from seminar participants at the Tinbergen Institute.

Abstract

This paper contributes to the productivity literature by using results from firm-level productivity studies to improve forecasts of macro-level productivity growth. The paper employs current research methods on estimating firm-level productivity to build time-series components that capture the joint dynamics of the firm-level productivity and size distributions. The main question of the paper is to assess whether the micro-aggregated components of productivity—the so-called productivity decompositions—add useful information to improve the performance of macro-level productivity forecasts. The paper explores various specifications of decompositions and various forecasting experiments. The result from these horse-races is that micro-aggregated components improve simple aggregate total factor productivity forecasts. While the results are mixed for richer forecasting specifications, the paper shows, using Bayesian model averaging techniques (BMA), that the forecasts using micro-level information were always better than the macro alternative.

JEL classification: C11, C14, C32, C33, D24, O12, O47

Keywords: economic growth, production function, total factor productivity, aggregation, firm-level data, Bayesian analysis, forecasting

1 Introduction

This paper contributes to the productivity literature by using results from firm-level productivity studies to improve forecasts of macro-level productivity growth. The paper employs current research methods on estimating firm-level productivity to build times-series components that capture the joint dynamics of the firm-level productivity and size distributions. The main question of the paper is to assess whether the micro-aggregated components of productivity—the so-called productivity decompositions—add useful information to improve the performance of macro-level productivity forecasts. To our knowledge, our study is a novel attempt to connect micro and macro level analysis whereby micro-level productivity estimates and decompositions of aggregate productivity provide additional information to be used in making macro forecasts.

Much of the recent research effort in the field of productivity, both at the micro and macro level, has been aimed at measurement. This research has achieved some notable results regarding several important issues - accounting for adjustment costs, variable factor utilization, intangible investments, etc. (e.g. Oliner et al. (2007), Basu and Kimball (1997), Kátay and Wolf (2008)); and issues connected to aggregation (e.g. Petrin and Levinsohn (2008)). Research into the practical matter of forecasting aggregate productivity has made much less progress. Forecasting total factor productivity (TFP) can be quite complicated, much more so than for well-understood primary economic measures. TFP essentially is computed as the ratio of output to inputs and measurement errors in either the numerator or the denominator do not cancel but exacerbate each other. A common finding of the productivity literature is that the yearly growth rates of TFP are often unstable or erratic, which makes forecasting exceedingly difficult.¹

Forecasting is further complicated by the fact that our theoretical understanding of aggregate productivity movements is less advanced than our understanding of, say, the labor market. Some theory posits that representative firms merely are able to pluck the fruits of movements in an exogenous productivity frontier. Consistent with this notion, macro-level productivity forecasting uses historical timeseries of aggregate productivity to

¹For this reason, most studies report annual percent changes calculated over several years depending on the dataset (for example: Corrado et al. (2006) and Oliner et al. (2007)).

uncover the underlying statistical process of TFP shocks. More recent work, grounded in empirical observation (Bartelsman and Doms (2000), Bartelsman et al. (2005)) lays the source of productivity change at explicit actions at the firm level, so that aggregate TFP change is the result of within-firm changes in productivity and between-firm movements in market share. Forecasting aggregate productivity movements thus requires theory and measurement of firm-level actions and productivity development as well as theory and measures of market selection and firm-size development. The main contributions on firm-level TFP measurement are Olley and Pakes (1996), Levinsohn and Petrin (2003), and a more recent study by Petrin and Levinsohn (2008). The classical papers on aggregation of TFP are Domar (1961) and Hulten (1978). However, the aggregation papers either assume that the firms operate in a frictionless world or assume that all firms are essentially identical 'representative firms'.

The theoretical links between firm-level actions and the interactions between firms and customers in a market that result in a non-degenerate firm size distribution have not been well modelled. Some early work attempts to provide explanations of how heterogeneous firm-level productivity interacts with resource allocation to determine aggregate growth. One example is Bartelsman et al. (2008a), who model the joint distribution of size and productivity. Another is Lentz and Mortensen (2005), who structurally estimate a model of creative destruction. This modelling does not yet provide a framework for using the full information in firm-level data for macro-level forecasting. Nevertheless, they provide the intuition for using micro-aggregated representations of the underlying relationships among size, productivity, firms' decisions, and market selection to aid in macro forecasting.

The micro-aggregated components of aggregate productivity growth may help explaining the aggregate for several reasons. First, the 'within' component or productivity growth among individual firms, is a simple weighted average of individual growth rates. It is the sole factor behind aggregate productivity dynamics in a neoclassical world. In our view, the within component can be expected to have forecasting power because it captures common behavior over the business cycle (such as factor hoarding), and steady factors at lower frequencies (such as technology diffusion). Second, the 'between' component is assumed to represent the underlying forces of reallocation of resources across

firms. It captures market selection mechanisms whereby more productive firms gain and less productive firms lose market share. Based on the findings of earlier work by Baily et al. (2001), Basu and Kimball (1997), and Basu et al. (1998) it may be expected to explain behave quite differently over the cycle than the within component, while over time it may reflect changes in demand conditions or policy environment. Third, the net entry component captures the entry and exit at the fringe, as well as the rejuvenation of industries through high-growth startup. If, for instance, the net entry term is large², it implies that entrants' productivity is larger than that of exiters, possibly indicating new opportunities arising through technological breakthroughs.

In our work, we do not attempt to put such structure or meaning on the contribution of the firm-level productivity decompositions to the macro forecasts. The literature on decomposing aggregate productivity growth has grown extensively over the past few years and indicates that there are many ways to define or modify the decompositions. Our primary objective is to improve forecasts of aggregate productivity growth using disaggregated information and not to favor one method of decomposition over another. To this end, we work with two standard decompositions (Olley and Pakes (1996), Baily et al. (2001)). We also introduce and estimate simple auxiliary equations, based on theoretical notions from the literature, for both individual productivity changes and evolution of firm-level market shares. The objective of the auxiliary models is to extract further information from the joint distribution of productivity and market share.

Our firm-level dataset, provided by Statistics Netherlands (CBS)³ consists of a large yearly panel of Dutch manufacturing firms over 1978-2004. More details can be found in Appendix 5.1.

The paper is organized as follows. Section 2 lays down the outline of the method we use to estimate TFP at the firm-level and describes the productivity decompositions

²The effects of firms' entry/exit decisions may become large. For example, Jaimovich and Floetotto (2008) showed in a structural IO model that the interaction between the number of firms and markups gives rise to endogenous procyclical movements in TFP. Their quantitative analysis suggest that about 40% of the variation in measured TFP in the U.S. is due to this interaction.

³The CBS 'Centrum voor Beleidsstatistiek' provides remote access facilities to the confidential data, under strict rules for disclosure.

used to generate the micro-aggregated productivity components used in the forecasting exercises. Section 3 discusses the various forecasting experiments and the metric used to evaluate their performance. The results of the various exercises are presented and the Bayesian Model Average outcomes show how the firm-level information improves the ability to forecast aggregate productivity. Section 4 concludes.

2 Firm-level TFP measures and productivity decompositions

In this paper, the empirical work is separated into measurement and aggregation procedures using micro-level panel data, and forecasting exercises using aggregated time-series. This section describes the work conducted in a secure computing environment using confidential firm-level information. The output of this section will be a collection of micro-aggregated time-series components of productivity that are analyzed outside of the secure environment, together with other macro-level timeseries, in order to forecast aggregate productivity. The forecasting analysis is described in section 3.

In this section, we provide a brief description of the methods used to construct firm-level TFP measures, as well as the methods by which the firm-level information is aggregated into micro-aggregated productivity components. To generate estimates of TFP at the firm-level we applied a standard procedure introduced by Olley and Pakes (1996) (OP) and modified by Levinsohn and Petrin (2003) (LP henceforth). The decompositions of aggregate productivity—i.e. micro-aggregates of firm-level productivity—are based on methods available in the literature, Olley and Pakes (1996) and Baily et al. (2001), and extended to allow a different dynamic interpretation.

2.1 Firm-level TFP measures

In estimating firm-level productivity measures, a main issue is controlling for endogeneity of the primary inputs into production. Consider a Cobb-Douglas production function

(indices i and t were dropped for simplification)

$$y = \beta_0 + \beta_l l + \beta_k k + \tau + \varepsilon \tag{1}$$

where y is log value-added, k is log capital, l is log labour, τ is log productivity and ε is assumed to be an *iid* disturbance. τ is unobservable by the econometrician but known⁴ to the firm. Since τ is in the information set on which the firm conditions its optimal choices of inputs, there will always be a nonnegative correlation between input factors and τ . This dependence renders simple OLS parameter estimates to be biased.

One of the key assumptions in the procedure to control for the endogeneity bias, is that capital is pre-determined, i.e. its level is chosen before production takes place. Hence, the orthogonality of k to the innovation in τ can be used to identify β_k . To solve the endogeneity problem with respect to freely variable labor, the method makes use of a proxy. The key assumption is that the proxy is monotonic in τ because then the proxy can be used to invert out the unobserved productivity shock. The main difference between OP and LP is that the former uses investment, the latter uses intermediate material use as a proxy.

There are, of course, caveats to both procedures, i.e. which proxy one should use. Akerberg et al. (2005) lists several arguments and suggests further modifications to better measure the parameters of the production function. Yet, we chose LP because (i) our forecasting exercises do not concern parameter-measurement in the first place, (ii) it is easily tractable and has strong intuitive arguments for identification, (iii) the data do not have direct estimates of investment needed to produce firm-level capital measures. We did not want to introduce more measurement error by applying another proxy for investment.

2.1.1 Production function parameter estimates

Production function parameter estimates are presented for two concepts of production, namely value added (table 3) and gross output (table 4). Column 1 shows the industry codes—table 1 provides the legend—columns 2-4 contain the point estimates for K , L and

⁴at least up to its expected value

M , column 5 lists the value of a χ^2 test statistic for constant returns to scale, columns 5-6 show standard errors for $\hat{\beta}_k$ and $\hat{\beta}_l$, and the last column shows the number of firm-year observations. The parameter estimates and their variation across industries seem in line with findings in the literature. The $\hat{\beta}_k$ are found non-significant more frequently for the gross output production function, which was against our expectations. Since the output specification requires a more complex IV-technique together with a grid-search algorithm, it is less clear whether the difference between the two production concepts is explained by the estimation method or measurement issues. While we generated micro-aggregated TFP components for use in the forecast exercises for both the value added and gross output concepts, in the remainder of the paper we focus on results for value added.

2.2 Productivity Decompositions

Starting with the firm-level measures of productivity, accounting identities provide the link with aggregate productivity. In general, productivity (growth) at the aggregate level can be computed as a weighted average of firm-level productivity growth, with the weights related to firm size.⁵ The accounting identities can then be parsed into components that add up to aggregate productivity and that reflect movements in the firm-level productivity and reallocation processes. We start by looking at a simple decomposition Baily et al. (2001), and continue by building on a decomposition given by Olley and Pakes (1996). Both of these are augmented with a further refinement that models the 'push and pull' process by which market shares evolve depending on a firm's position in the firm-level productivity distribution.

The outputs of this part of the empirical exercise are sets of specifications that are to be used in the forecasting exercise. These specifications are based on binary choices

⁵If one wants measures of aggregate productivity growth that are derived in a welfare theoretic setting, many assumptions need to be made that likely do not hold for firms or industries (see e.g. Hulten (1978), Petrin and Levinsohn (2008)). Aggregated productivity computed following an accounting identity may therefore not reflect how aggregate welfare changes if firm-level productivity changes *and* all resources are reallocated instantaneously to their optimal use. Of course, frictions prevent such instantaneous reallocation and provide the reason that firm-level information may provide information about aggregate movements.

along four dimensions. We start with components based on value added or gross output measures of firm-level TFP as described above.⁶ Next, the accounting identities require a measure of market share or firm size. These may be constructed using the firm's share of industry output (either value added or gross output), or the firm's share of industry inputs (aggregating the factor inputs using the estimated output elasticities as weights). Next, the productivity components are computed using either of the two decomposition methods described below in equations 2 and 6. Finally, auxiliary models (equations 7 and 8 below) are estimated to capture productivity 'push and pull' effects. The models generate fitted values for individual firms' market shares and productivity growth. The productivity components are then based on either the actual or fitted values of these firm-level indicators.

2.2.1 Accounting Identities

The first dynamic decomposition, laid out by Baily et al. (2001), is the following:

$$\Delta\tau_t = \sum_{i \in C} \tilde{\phi}_i \Delta\tau_{it} + \sum_{i \in C} \Delta\phi_i (\tilde{\tau}_i - \tilde{\tau}) + \sum_{i \in E} \phi_{it} (\tau_{it} - \tilde{\tau}) - \sum_{i \in X} \phi_{it-1} (\tau_{it-1} - \tilde{\tau}), \quad (2)$$

where τ_{it} is the TFP-level of firm i in period t , ϕ_{it} is the market share of firm i in period t , $\tilde{\tau} = \frac{\tau_t + \tau_{t-1}}{2}$, Δ is the difference operator, C is the set of continuers, E is the set of entrants and X is the set of exiters in time t . The terms in the previous equation are often called within-, between-, entry-, and exit-terms, respectively.

In the second decomposition, we combine a static decomposition with equation 2. The static equation was introduced by Olley and Pakes (1996). The static decomposition is

$$\tau_t = \bar{\tau}_t + \sum_{N_t} (\phi_{it} - \bar{\phi}_t) (\tau_{it} - \bar{\tau}_t), \quad (3)$$

where $\bar{\tau}_t = \frac{1}{N_t} \sum_{N_t} \tau_{it}$, and $\sum_{N_t} (\phi_{it} - \bar{\phi}_t) (\tau_{it} - \bar{\tau}_t)$ is a covariance measuring the association between productivity and size. This decomposition informs us about characteristics of τ 's cross-section distribution. When the covariance-term is high, the weighted

⁶However, we present results based on value-added productivity for reasons outlined at the end of Section 2.1.1.

average is higher than the unweighted average. It also implies that the change in aggregate productivity is the sum of the change in the unweighted average productivity and change in the covariance-term. Differencing 3 yields:

$$\Delta\tau_t = \Delta\bar{\tau}_t + \sum_{N_t} (\phi_{it} - \bar{\phi}_t)(\tau_{it} - \bar{\tau}_t) - \sum_{N_{t-1}} (\phi_{it-1} - \bar{\phi}_{t-1})(\tau_{it-1} - \bar{\tau}_{t-1}), \quad (4)$$

i.e. aggregate productivity change is sum of the change in the unweighted average productivity and change in the covariance term. Entry and exit are implicitly accounted for here as the cross section sums run from $i = 1 \dots N_t$ and $i = 1 \dots N_{t-1}$ respectively. Hence, the weights are such that $\sum_{N_t} \phi_{it} = 1$ and $\sum_{N_{t-1}} \phi_{it-1} = 1$. Now express $\tilde{\tau}$ using 3:

$$\begin{aligned} \tilde{\tau} &= \frac{\tau_t + \tau_{t-1}}{2} = \frac{\bar{\tau}_t + \bar{\tau}_{t-1}}{2} + \frac{cov_t(\phi, \tau) + cov_{t-1}(\phi, \tau)}{2} \\ &= [\tilde{\tau}_t] + \tilde{cov}(\phi, \tau). \end{aligned} \quad (5)$$

We expect $\tilde{cov}(\phi, \tau)$ to reveal information about the cross section characteristics of the distribution. We base our analysis on 2, but we also want to introduce information about the size and productivity distributions. To this end, we plug 5 into the between term of 2 to get

$$\begin{aligned} \sum_C \Delta\phi_i(\tilde{\tau}_i - \tilde{\tau}) &= \sum_C \Delta\phi_i\tilde{\tau}_i - \sum_C \Delta\phi_i\tilde{\tau} \\ &= \sum_C \Delta\phi_i\tilde{\tau}_i - \sum_C \Delta\phi_i([\tilde{\tau}_t] + \tilde{cov}(\phi, \tau)), \end{aligned}$$

so 2 looks like⁷

$$\begin{aligned} \Delta\tau_t &= \\ & \sum_C \tilde{\phi}_i \Delta\tau_{it} + \sum_C \Delta\phi_i(\tilde{\tau}_i - [\tilde{\tau}_t]) - \sum_C \Delta\phi_i \tilde{cov}(\phi, \tau) \\ & + \sum_E \phi_{it}(\tau_{it} - \tilde{\tau}) - \sum_X \phi_{it-1}(\tau_{it-1} - \tilde{\tau}). \end{aligned} \quad (6)$$

⁷This manipulation does not affect the validity of the decomposition. Had we rewritten all the terms in 2, we would have got

The ϕ_{it} weights were defined at the beginning such that $\sum_{N_t} \phi_{it} = 1$, and hence $\sum_C \Delta\phi_i$ is not necessarily zero (it would be zero had we defined $\sum_C \phi_{it} = 1$). Therefore the covariance term is not killed by $\sum_C \Delta\phi_i$. However, the between-term is now different: it sums up the share-changes weighted by deviations from the time-average of the simple cross-section average. The covariance-term is shown separately.

Push and pull effects Equations 2 and 6 show the two main decompositions we used in the forecasting exercises. Both are identities separating aggregate productivity growth into dynamic components. In the case of 2 we have three components ('within', 'between' and a combined 'net-entry' contribution), and in the case of 6 we have a different between component and an additional 'cross-term'. We denote these components as $\Delta\tau_t^i$, $i \in \{dc1, dc2\}$ in the two different decompositions. Both of these approaches take the underlying firm-level market share indicator and productivity growth rate as given. In order to put more structure on the evolution of these variables we model two effects:

$$\begin{aligned} \Delta\tau_t &= \sum_C \tilde{\phi}_i \Delta\tau_{it} + \sum_C \Delta\phi_i (\tilde{\tau}_i - [\tilde{\tau}_t]) - \widetilde{cov}(\phi, \tau) \sum_C \Delta\phi_i \\ &\quad + \sum_E \phi_{it} (\tau_{it} - [\tilde{\tau}_t]) - \sum_X \phi_{it-1} (\tau_{it-1} - [\tilde{\tau}_t]) \\ &\quad - \widetilde{cov}(\phi, \tau) \left[\sum_E \phi_{it} - \sum_X \phi_{it-1} \right] \end{aligned}$$

where

$$\begin{aligned} & -\widetilde{cov}(\phi, \tau) \sum_C \Delta\phi_i - \widetilde{cov}(\phi, \tau) \left[\sum_E \phi_{it} - \sum_X \phi_{it-1} \right] \\ &= -\widetilde{cov}(\phi, \tau) \left[\sum_C (\phi_{it} - \phi_{it-1}) + \sum_E \phi_{it} - \sum_X \phi_{it-1} \right] \\ &= -\widetilde{cov}(\phi, \tau) \left[\sum_C \phi_{it} + \sum_E \phi_{it} - \sum_C \phi_{it-1} - \sum_X \phi_{it-1} \right] \\ &= -\widetilde{cov}(\phi, \tau) \left[\sum_{N_t} \phi_{it} - \sum_{N_{t-1}} \phi_{it-1} \right] = 0 \end{aligned}$$

a "pull" and a "push" effect. The specification is derived from the literature on frontier productivity, both theoretical (Acemoglu et al. (2002)) and empirical (Bartelsman et al. (2008b)). The approach essentially pins down a firm's position relative to the frontier. The pull-equation says that individual productivity growth depends positively on the distance from the frontier, in other words, firms further away are pulled more strongly towards it as technology spreads out:

$$\Delta\tau_{it} = \beta(\tau_t^F - \tau_{it}) + \eta_{it}, \quad (7)$$

where τ_t^F is frontier productivity and η_{it} is an autonomous or noise component. A case can be made for both $\beta < 0$ and $\beta > 0$. In the first case firms closer to the frontier can be thought to have greater absorption capacity⁸ (greater human capital, which is unmeasured, and hence it shows up as higher productivity), so they can take on new technologies and grow faster. In the $\beta > 0$ case - investigated by the studies above - less productive firms tend to grow faster. For instance, these firms can be thought as being new entrants with relatively high growth potential.⁹

Frontier productivity (τ_t^F) for each industry was calculated as the average of the top decile from a truncated distribution. We got rid of the top and bottom 1% of the observations to kill the possibly erratic effects on aggregates.¹⁰ We took a moving average of $\tau_{it}^m = \frac{\tau_{it} + \tau_{it-1}}{2}$ to further guard against extreme observations. The resulting truncated distribution was then used to calculate the average of the top decile.

The push equation is assumed to encapsulate the market selection mechanism whereby less efficient firms are crowded out of the market. Assume the following specification:

$$\phi_{it} = \gamma_0 + \gamma_1(\tau_{it} - \bar{\tau}_t) + \varepsilon_{it}. \quad (8)$$

This equation reads: higher-than-average productivity firms end up gaining market share ($\gamma_1 > 0$), and lower-than-average productivity firms lose market share.

⁸cf. with the "second face" of research as called in the R&D literature

⁹The literature calls the $\beta > 0$ case - when less productive firms grow faster - as β -convergence. The concept was introduced by Robert Barro and Xavier Sala-i-Martin in their 1991 paper.

¹⁰I ignored the top and bottom 1% of the productivity distribution because it turned out that frontier time-series exhibited unplausibly large time-variation if these observations were included.

The use of equations 7 and 8 is that once we estimated parameters β, γ_0 and γ_1 , the fitted values $\widehat{\Delta\tau}_t$ and $\widehat{\phi}_{it}$ from these simple models inform about the underlying¹¹ catch-up components and shares. If they have explanatory power in our dataset then $\widehat{\Delta\tau}_t$ and $\widehat{\phi}_{it}$ should add to the forecasting power of the set of components that are computed using these fitted values rather than the actual data on firm-level productivity growth and market share.

Estimation of the auxiliary models The distance variables in both 7 and 8 are endogenous. For 7, the endogeneity problem is immediate. As for 8, a case can be made that a firm gains market share exactly because its productivity increased in the wake of a positive productivity shock. One way to get around it is to apply an IV estimator. The sample moment condition for $\widehat{\beta}$ is

$$E[z_i (\Delta\tau_i - \widehat{\beta}_{IV}(\tau_t^F - \tau_{it}))] = 0,$$

and the simple Anderson-Hsiao-type IV estimator is given by

$$\widehat{\beta}_{IV} = \frac{\sum_{i=1}^N \sum_{t=2}^T z_{it} \Delta\tau_{it}}{\sum_{i=1}^N \sum_{t=2}^T z_{it} (\tau_t^F - \tau_{it})} \quad \text{with } z_{it} = (\tau_{t-2}^F - \tau_{it-2}).$$

The instruments can also be Δz_{it} but that results in the loss of more observations. As for the push-equation (8), rewriting $\phi_{it} = x\gamma' + \varepsilon_{it}$ where $\gamma = [\gamma_0, \gamma_1]$ and $x = [1, \tau_{it} - \bar{\tau}]$.

¹¹in the sense of the simple econometric models 7 and 8.

The conventional IV-estimator again is

$$\begin{aligned}\hat{\gamma}_{IV} &= \frac{\sum_{i=1}^N z_i' \phi_i}{\sum_{i=1}^N z_i' x_i} = (Z'X)^{-1} Z'\Phi \text{ with} \\ z_i &= [(\tau_{i1} - \bar{\tau}_1), \dots, (\tau_{iT-1} - \bar{\tau}_{T-1})]', \\ \phi_i &= [\phi_{i2}, \dots, \phi_{iT}]' \text{ and} \\ Z &= [z_1' \dots z_N']', \Phi = [\phi_1' \dots \phi_N']'.\end{aligned}$$

The IV-results of equations 7 and 8 were in line with our expectations. β appeared significantly positive at conventional levels in all twelve industries when a value-added-based production function was used¹². Results were very similar in the case of gross output, except that β did not appear significant in the textiles industry.¹³ The γ 's were significantly positive for value-added-based productivity measures. The output-based specifications showed more ambiguous associations: γ_0 were negative in Nonmetallic mineral products and Electrical and optical equipment. Estimates were invariant to whether we measured firm size by input- or output-side indicators of firms.¹⁴

3 Forecasting

The previous section described the estimation and computation of a set of micro-aggregated productivity components that are to be used in conjunction with macro timeseries in order to forecast aggregate productivity. The basic idea in this section is to assess whether adding the timeseries components built up from the firm-level data improves the traditional forecasts that are made using only published aggregate timeseries. In this section we describe the experiments used in the assessment. First we argue that the proper metric

¹²The mean and standard deviation of β across the twelve industries was .039 and .014.

¹³The mean and standard deviation of β across the twelve industries was .032 and .016.

¹⁴That is, using input shares or value added/gross output shares did not affect point estimates at 4-digit precision.

is not to compare forecasts to actual annual TFP growth, but to the annual average trend growth rate. Next, we describe the 'horse races' comparing forecast performance of the macro timeseries to the forecast performance when a set of productivity components is used as well. The first horse race attempts to shed light on whether we can improve on the accuracy of aggregate productivity forecasts by forecasting microcomponents separately. The second investigates whether microcomponents better predict aggregate productivity in a multivariate framework. The third experiment uses the same multivariate framework but the target is the published aggregate.

3.1 Methodology

Yearly growth rates of total factor productivity are often found to be volatile. A common way to cope with the large variation is to analyze changes over a longer period. As an illustration, figure 1 shows yearly TFP-changes and table 5 shows average yearly changes for five-year periods for several countries.

Figure 2 depicts our micro-based and 'published' TFP measures for aggregate Dutch manufacturing. Micro-based aggregates are built using firm-level productivity estimates, 'published' aggregates, built from national accounts data, come from the EUKLEMS database (O'Mahony et al. (2007b)). We will refer to our firm-based TFP aggregates as micro-aggregated TFP hereafter. Although the two datasources characterize the same set of firms, the published and micro-aggregated TFP measures differ for a variety of reasons related to statistical practice. To trace these differences to underlying measures of inputs and outputs, figures 3 and 4 graph year-on-year growth rates for value added and labor.

Several stylized facts emerge from these graphs. First, aggregate yearly growth rates exhibit large variation in all countries (figure 1). Second, micro-aggregated measures generally vary more than published aggregates (figures 3, 4). Third, micro-aggregated dynamics are similar on the input and output side of the firm suggesting that yearly TFP growth numbers actually make sense despite their variation.¹⁵ Finally, micro-aggregated

¹⁵The same conclusion emerges from industry-by-industry analysis, not shown here. Micro-macro discrepancies are all the more obvious looking at level charts.

TFP-growth series - although more volatile - captures the dynamics of the published aggregate reasonably well (figure 2).

The most prominent of all the above observations is that yearly aggregate productivity growth exhibits large variation. Instead of focusing on yearly forecasts, we evaluate our estimates in 3-year forecast windows. We do not calculate MSE-type values but simply compare average growth rates implied by extrapolated yearly values. The relatively long time-span of the sample allowed us to evaluate results in four rolling forecast windows: 1998-2000, 1999-2001, 2000-2001, 2001-2003. These forecast windows are useful for at least two reasons. First, they allow us to investigate the effect of adding more data to our models. Second, they give insight about the predictive content of components $\Delta\tau_t^i$ over different phases of the business cycle¹⁶.

We make use of the two decompositions developed in 2.2, to reveal components that facilitate interpreting the dynamics of aggregate productivity. Both decompositions consist of three or four building blocks that aggregate to overall productivity growth. First, the within component shows the effects of firm-level productivity change. Second, the between component shows how the reallocation of resources among firms contributes to aggregate productivity growth. The third component encapsulates the effects of firm entry and exit. In decomposition 6 a cross-term component also is included.

The idea of the 'horse races' was to compare forecast performance of the aggregate timeseries to the forecast performance when a set of productivity components was used as well. The performance metric we chose was the average growth of the Hodrick-Prescott-trend¹⁷ of the aggregate, thus, for any two competing specifications we calculated the following:

$$\widehat{\Delta\tau}_{t+s}^{Ag} - \overline{\Delta\tau}_{t+s}^{HP} \text{ vs } \widehat{\Delta\tau}_{t+s}^{Mi} - \overline{\Delta\tau}_{t+s}^{HP}, \quad (9)$$

¹⁶Dutch manufacturing exhibited fast growth over the first forecast window (1998-2000). These were the last years of a high-growth-period, which started in the first half of the nineties. The second forecasting window (1999-2001) contains the turning point in the cycle with negative growth in 2001 but still strong activity over the entire three-year horizon. The remaining two windows were periods of downturns. See figure 3 for an illustration of these periods.

¹⁷Figure 5 shows the actual and HP-trend series of EUKLEMS aggregate productivity.

where $\widehat{\Delta\tau}_{t+s}^{Ag}$ denotes the forecast of the aggregate, $\widehat{\Delta\tau}_{t+s}^{Mi}$ denotes the forecast of micro-components, and $\overline{\Delta\tau}_{t+s}^{HP}$ denotes the average growth of the Hodrick-Prescott-trend of the aggregate, and s denotes the particular forecast window in which the forecast is evaluated.

We now describe the forecasting experiments. Experiment 1 was designed to investigate whether disaggregate information help forecast the aggregate, to be more specific, whether separate forecasts of disaggregates could improve on forecasting of the aggregate. Disaggregate information was encapsulated by the dynamic microcomponents implied by (2) and (6), which - in shorthand notation - specify aggregate productivity growth as

$$\Delta\tau_t^{Ag} = \Delta\tau_t^w + \Delta\tau_t^b + \Delta\tau_t^{ne} = \sum_{dc1} \Delta\tau_t^i \quad (dc1)$$

$$\Delta\tau_t^{Ag} = \Delta\tau_t^w + \Delta\tau_t^{b1} - \Delta\tau_t^{b2} + \Delta\tau_t^{ne} = \sum_{dc2} \Delta\tau_t^i. \quad (dc2) \quad (10)$$

In technical terms, we first estimated separate univariate autoregressive specifications for all the microcomponents $\Delta\tau_t^i$, $i \in \{dc1, dc2\}$ and the aggregate $\Delta\tau_t^{Ag}$.¹⁸ Then we computed the microcomponents forecasts from the disaggregate specifications $\widehat{\Delta\tau}_{t+s}^i$ such that

$$\widehat{\Delta\tau}_{t+s}^{Mi} = \sum \widehat{\Delta\tau}_{t+s}^i, \quad i \in \{dc1, dc2\}. \quad (11)$$

The aggregate forecast $\widehat{\Delta\tau}_{t+s}^{Ag}$ was computed from its own autoregressive model. Finally, we compared the performance of $\widehat{\Delta\tau}_{t+s}^{Ag}$ and $\widehat{\Delta\tau}_{t+s}^{Mi}$ using the metric in (9).

A related question is how the dynamic components $\widehat{\Delta\tau}_{t+s}^i$ perform relative to other disaggregate information. We investigated this issue in experiment 1', which amounted to rerunning experiment 1 but using conventional statistical industry components as disaggregates. We formed four groups within manufacturing following EUKLEMS-practice. These are consumer goods, intermediate goods, electrical machinery and other investment

¹⁸We selected the order of all the $AR(p)$ models based on the *BIC* rather than the *AIC* because *BIC* more heavily penalizes the loss of degrees of freedom. The *BIC* suggested that for our yearly observations the $AR(1)$ specifications - relative to $AR(2)$ - were rich enough to capture the dynamics of $\Delta\tau_t^i$ and $\Delta\tau_t^{Ag}$. We did not estimate models with $p > 2$ because we wanted to preserve degrees of freedom.

goods industries.¹⁹ Then, just like above, we forecasted aggregate productivity using disaggregate information, which was encapsulated by these four industry components. In short, our $\widehat{\Delta\tau}_{t+s}^{Mi}$ is now implied by

$$\widehat{\Delta\tau}_{t+s}^{Mi} = \sum_{i \in I} \widehat{\Delta\tau}_{t+s}^i, \quad (12)$$

where $I = \{cons, inter, inv, elec\}$ and the $\widehat{\Delta\tau}_{t+s}^i$ are all generated by separate univariate autoregressive specifications.

The set of specifications in the experiments above is spanned by three dimensions. First, we have two decompositions shown in equations 2 and 6. Second, the values for the shares ϕ_{it} were calculated using output- or input²⁰-side measures of firms. Third, our auxiliary models in section 2.2.1 allowed us to use either measured and fitted values of $\Delta\tau_t$ and ϕ_{it} in the decompositions. Overall, the three dimensions yielded eight specifications for the microcomponents and the aggregate forecasts. The industry components forecasts were evaluated for four specifications as there we had only one decomposition.

Experiment 2 was intended to investigate whether microcomponents could improve aggregate forecasts in a multivariate framework. This experiment attempts to shed light on two issues. First, it is of interest to see how results are affected in a multivariate forecast model, which is common in macro-practice. Second, it is also important to know how results are affected by combining micro-aggregated information with published aggregate information. To this end, we estimated vector autoregressions where the regressand is a vector of logarithmic differences of value added, capital, labor and microaggregated TFP: $y_t = (\Delta va_t, \Delta k_t, \Delta l_t, \Delta\tau_t^{Ag})$.²¹ Regressors included some of y_t 's own lags together with lagged microcomponents. We denote the latter as $\Delta\tau_{t-q}^i$, $i \in \{dc1, dc2\}$. Data for Δva_t , Δk_t , Δl_t came from the published EUKLEMS database, while $\Delta\tau_t^{Ag}$ and $\Delta\tau_{t-q}^i$ were constructed from firm-level data using equations 2 and 6. In short, and using the generic notation $VAR(p, q)$ for a VAR with p endogenous lags and q lags of microcomponents $\Delta\tau_t^i$,

¹⁹See Table 1 for classification details.

²⁰These are the fitted values from production function estimation: the weighted average of primary inputs where the weights are the production function parameters.

²¹The components of y_t are suggested by the production function.

we estimated $VAR(1, 2)$, $VAR(2, 1)$ and $VAR(1, 1)$ using least squares. Correspondingly, the set of specifications described under the experiment 1 were extended by the order of the VARs.

We asked a similar question in experiment 3. The difference with experiment 2 is that we are now interested in whether microcomponents improve published aggregate forecasts in a multivariate framework. This amounted to replacing $\Delta\tau_t^{Ag}$ by $\Delta\tau_t^{EUKLEMS}$ in y_t above. This experiment simulates a more realistic forecasting environment because, at time t , microaggregated information $\Delta\tau_t^{Ag}$ is generally not available to the researcher whereas $\Delta\tau_t^{EUKLEMS}$ is. The purpose here was to investigate the relationship between the published aggregate $\Delta\tau_t^{EUKLEMS}$ and the lagged microcomponents $\Delta\tau_t^i$ and see whether (lagged) microcomponents have marginal forecasting power over the lagged published aggregate.

3.2 Forecast results

Our evaluation method amounts to comparing the performance of the microcomponents specifications and their aggregate alternatives by calculating the deviations from our Hodrick-Prescott benchmark given by 9.

In experiment 1, forecasts are specified by equations 11 (and 12 for experiment 1'). Tables 6 and 7 show forecast performance in the first and the last forecast window for value-added-based productivity.²² The first column of each table identifies whether direct measures or fitted values were used for $\Delta\tau_{it}$ and ϕ_{it} in the specifications. This column also shows whether the shares ϕ_{it} were calculated using output- or input-side measures of firms. Columns 2 and 3 contain the microcomponents forecasts specified in (11) for the two decompositions specified in (10). Column 4 shows the forecasts for industry components specified in (12). Columns 5 and 6 show the aggregate forecasts $\widehat{\Delta\tau}_{t+s}^{Ag}$ for total manufacturing. The first entry in column 2 of table 6 says that the microcomponents forecast (using observed productivity and observed output-shares) is 2.4 percentage points lower than the HP-trend over 1998-2000. All the other entries are to read analogously.

²²We do not present results for all the windows to save on space. Detailed results for the other forecast windows are available upon request.

The corresponding aggregate forecast, the first entry in column 5, is -3.2% . These two numbers imply that the aggregate alternative to microcomponents is 0.8 percentage point farther away from the observed HP-trend. If we compare numbers in the previous manner we can conclude that aggregate forecasts would have missed to detect the productivity acceleration of the late nineties whereas several of our specifications would have not. It is also clear that all microcomponents forecasts beat the aggregate forecasts. This conclusion carried over to all remaining forecast windows, Table 7 shows it for the 2001-2003 window.

As for industry components (column 4 in tables 6 and 7), we found that their forecasts were closer to aggregate ones. More importantly, industry components could *not* beat the micro-aggregated components forecasts in any of the four forecast windows. This is an interesting result and suggests that the decompositions of aggregate productivity among incumbents and entrants, and among within and between contributions are more informative about aggregate productivity dynamics than the productivity dynamics of industry subgroups.

These results are remarkable for several reasons. First, they suggest that microcomponents' predictive power is strong not only in particular periods but over all phases of the business cycle (see footnote 16 in Section 3.1 for details). This is in contrast with Hendry and Hubrich (2006), who found that disaggregate forecasts based on subaggregates of the consumer price index performed worse than aggregate forecasts.²³ We also found that decomposition 6 produced more precise forecasts than decomposition 2, suggesting that introducing information about the cross section distribution proved to be useful.

Experiment 2 was designed to reveal the predictive power of microcomponents over aggregate productivity growth in a VAR-framework. $\widehat{\Delta\tau}_{t+s}^{Ag}$ was generated from VARs where no lagged microcomponents were present in the set of explanatory variables and $\widehat{\Delta\tau}_{t+s}^{Mi}$ was generated from VARs with microcomponents. We note the main conclusions here without presenting detailed results.²⁴ First, the superiority of microcomponent-VARs were not obvious. They performed better in some specifications and forecast windows,

²³Upon reflection, this was to be expected. Micro-level pricing behavior depends in an important way on (expected) aggregate prices and the individual differences may cancel. Feedback from the aggregate to individual productivity behavior likely is small.

²⁴Numbers are available upon request.

while in others they did not. Second, the issue of model selection proved to be critical. An obvious criterion to choose from a set of VARs is the Bayesian Information Criterion (BIC). We found that using BIC to choose between a microcomponent-VAR and an aggregate VAR may lead to serious errors. For instance, choosing a microcomponent-VAR based on BIC could lead to errors up to 1 – 2 absolute percentage points as measured by our metric in equation 9. On the other hand, not choosing a microcomponent-VAR based on BIC could lead to errors up to 1 absolute percentage points by our metric. Furthermore, we always found a microcomponent-VAR in our specification-set, which performed better²⁵ than the aggregate-VAR. The overall message of experiment 2 is that (i) we can always find a microcomponent-VAR performing better than the aggregate-VAR, even in our restricted specification set; (ii) model selection is critical but ambiguous; (iii) selecting a particular VAR based on BIC may lead to large errors.

To explore the practical question whether microcomponents can better predict the published aggregate we replaced $\Delta\tau_t^{Ag}$ by $\Delta\tau_t^{EUKLEMS}$ in experiment 3. As above, we only note the main conclusions here. Results resemble those in experiment 2 in that BIC and our performance measure did not yield the same rank in our specification-set. For instance, BIC favored a VAR featuring input shares in the first forecast window with a 1.2 percentage deviation, however, measured shares led to only .08 percentage point deviation (not shown here). We also found that we could always do better than the aggregate-VAR with at least one microcomponent-VAR specification. However, just like in experiment 2, we did not find clear superiority of microcomponents.

3.3 Bayesian analysis

This section outlines the Bayesian forecasting approach to experiment 3. It has at least three advantages over the Frequentist exercise. First, it is easy to compare and test different specifications even if they are not nested. This is an important advantage because we saw that model selection is a critical issue when microcomponents are combined with macro-information in a multivariate framework. Second, combining forecasts is simple and intuitive in a Bayesian framework. Third, the Frequentist analysis does not

²⁵in terms of our performance measure in 9

fully account for uncertainty. We could calculate confidence bands around the forecasts but those would only reflect forecast uncertainty. A Bayesian extension reflects both parameter uncertainty and the inherent randomness of forecasting.

This section outlines the Bayesian approach assuming we retrieved a posterior distribution for the coefficients and error covariance matrix of a VAR. Denote them by β and V , respectively. A technical description of the estimation and forecasting can be found in Appendix 5.2.

The main object of interest is the predictive density

$$p(y_{T+1}|y_T, M) = \int_0^\infty \int_{-\infty}^\infty p(y_{T+1}|V, \beta, M)p(V, \beta|y_T, M)d\beta dV, \quad (13)$$

where y_s is observed data up to time s , V is the covariance matrix of shocks and β is the parameter matrix in the VAR, and M denotes the model we are using. The predictive density coherently integrates (i) the uncertainty about β , V and (ii) the intrinsic uncertainty about the future y_{T+1} , conditional on the history y_T of observed data and model M . $p(V, \beta|y_T, M)$ and $p(y_{T+1}|V, \beta, M)$ are retrieved from a customized Gibbs-sampler (MCMC posterior simulator see, e.g. Koop (2003)). An illustrative example of the descriptives of $p(y_{T+s}|V, \beta, M)$, $s = 1, 2, 3$ in our case is shown by figure 6 and tables 8 and 9. These performance measures for value added/output, capital, labor and productivity are similar in nature to the ones in experiments 2 and 3. In the sequel we report performance measures only for TFP. Deviations from average HP-trend growth were computed using the means of the distributions of individual growth trajectories, not the means of the yearly distributions.

We estimated and forecasted several VARs with various lag-order in the four forecast windows. As an example, we present detailed results for the first window (1998-2000) in Table 10. All numbers are to read as explained in Section 3.2. For instance, the first entry of column 2 of Table 10 says that measured output-side shares in a $VAR(1, 2)$ yielded .3 percentage point higher productivity growth forecasts than the average HP-trend.

What is different is that Table 10 gives a Bayesian Model Average (BMA) forecast. The average is taken over decompositions and structures and is weighted by the model's Predictive Bayes Factor (PBF). The logic of averaging is the following. The predictive

likelihood $p(y_{T+s}|y_T, M_i)$ can be calculated in time $T + s, s \geq 1$. It is a real number given by evaluating $p(\cdot|y_T, M_i)$ at the *observed* y_{T+s} (Geweke and Amisano (2008), van Dijk et al. (2007)). The idea is to calculate the probabilities that the observed y_{t+s} was generated by model i . We evaluated them using the posterior simulator deployed to estimate β and V . That is, we computed the real number $R^{-1} \sum_{r=1}^R p(y_{T+s}|y_T, M, \theta^{(r)})$, where $\theta^{(r)} = (\beta^{(r)}, V^{(r)})$ and $r = 1 \dots R$ are the indices of the elements of the chain after convergence. Given these values, we constructed model weights using PBFs. We assumed the same model-priors, so the weights were given by the predictive likelihoods²⁶.

The first four rows of Table 10 already show that some of the microcomponents give more accurate forecasts than the aggregate alternative, and some do not. However, their PBF-weighted average is more accurate. It is remarkable that, in every forecast window and every VAR specification - not shown here -, the PBF-weighted (over decompositions and structures) average forecast always proved to be more accurate than the aggregate alternative even if some specifications performed worse.

We extended this way of model averaging over the entire set of specifications we have investigated so far. That is, the two decompositions and four structures and the three lag-order specifications ($VAR(2, 1)$, $VAR(1, 2)$, $VAR(1, 1)$) yielded 24 microcomponent-VARs and three aggregate-VARs in each forecast window.

Table 11 summarizes the result of model-averaging over the entire set of specifications. Several interesting conclusions emerge from these results. First, individual performance measures of VARs showed that forecasts were much more similar than in the Frequentist exercise²⁷. Second, BMA-forecasts using micro components were always better than the macro-alternative, whether averaging over only structures and decompositions or over the entire set of specifications. It holds true even if individual specifications performed less well. This is an interesting result and may be explained by the BMA-logic: every model is as good as its forecasts. Since we weighted forecasts by their predictive likeli-

²⁶after appropriate transformation.

²⁷This has probably to do with our strong prior on the variance of the shocks. We specified a restricted-variance shock-distribution for reasons explained in Appendix 5.2. As a result, the estimated means of simulated distributions differ less than in the frequentist exercise. Differences, or improvements should be evaluated in light of this.

hoods, specifications with less explanatory power were assigned smaller weights whereas specifications with more explanatory power were assigned larger weights.

4 Conclusions and Extensions

The aim of this paper is to build aggregate productivity forecasts using information on firm-level productivity and market share evolution. Our work builds upon the literature on measuring and analyzing firm-level productivity, but our question was not how to measure productivity at the firm level but rather how to apply these estimates in forecasting. To our knowledge, our study is a novel attempt to connect micro and macro level analysis whereby micro-level productivity estimates are used to build aggregates and forecasts.

We carried out several forecasting experiments for different specifications and models. The result from these horse-races is that micro-aggregated components improve simple aggregate total factor productivity forecasts. While the results are mixed for richer forecasting specifications, the paper shows, using Bayesian model averaging techniques (BMA), that the forecasts using micro-level information were always better than the macro alternative.

The paper does not address the important question of choosing the optimal forecasting tool. One natural extension in this direction is to investigate the performance of microcomponents in structural timeseries forecasting models. These signal extraction techniques are relevant because they impose explicit structure on the dynamic behavior of trends, cycles and irregular terms. As a consequence of the richer structure, we may expect them to capture observed dynamics better than simple autoregressive specifications. Another area for future work is choosing the optimal benchmark. The Hodrick-Prescott-filter is known to perform poorly at the sample endpoints, leading to the development of other real-time filtering techniques in recent years. Some alternatives are based on univariate approaches (see Wildi (1998), for an early example) others are based on multivariate approaches (see Stock and Watson (2002) for a dynamic factor modelling example). Although these issues are important and relevant aspects of our forecast evaluation exercises, they are beyond the scope of the paper and left for future research.

References

- ACEMOGLU, D., P. AGHION and F. ZILIBOTTI (2002), “Distance to Frontier, Selection and Economic Growth”, *Discussion Paper 3467*, Center for Economic Policy Research.
- ACKERBERG, D. A., K. CAVES and G. FRAZER (2005), “Structural Identification of Production functions”, .
- BAILY, M. N., E. J. BARTELSMAN and J. HALTIWANGER (2001), “Labor Productivity: Structural Change and Cyclical Dynamics”, *The Review of Economics and Statistics*, vol. 83, no. 3, pp. 420–433.
- BARTELSMAN, E. J. and M. DOMS (2000), “Understanding productivity: lessons from longitudinal microdata”, *Journal of Economic Literature*, vol. 38, no. 3, pp. 569–594.
- BARTELSMAN, E. J., J. HALTIWANGER and S. SCARPETTA (2005), “Measuring and analyzing cross-country differences in firm dynamics”, *draft*, NBER/CRIW conference, Washington DC., March 2005.
- BARTELSMAN, E. J., J. HALTIWANGER and S. SCARPETTA (2008a), “Cross Country Differences in Productivity: The Role of Allocative Efficiency”, *mimeo*.
- BARTELSMAN, E. J., J. HASKEL and R. MARTIN (2008b), “Distance To Which Frontier? Evidence on Productivity Convergence from International Firm-Level Data”, *Discussion Paper 7032*, Centre for Economic Policy Research.
- BASU, S., J. FERNALD and M. KIMBALL (1998), “Are Technology Improvements Contractionary?”, *Discussion Papers 625*, International Finance.
- BASU, S. and M. S. KIMBALL (1997), “Cyclical Productivity with Unobserved Input Variation”, *Working Paper 5915*, National Bureau of Economic Research.
- CORRADO, C., J. BEAULIEU, E. J. BARTELSMAN and P. LENGERMANN (2006), “Modeling Aggregate Labor Productivity and a Disaggregate Level: New Results for US sectors and industries”, *mimeo*.

- VAN DIJK, H. R., F. RAVAZZOLO and M. VERBEEK (2007), “Predictive gains from forecast combinations using time varying model weights”, *Econometric institute report*, Econometric Institute, Erasmus University Rotterdam.
- DOMAR, E. D. (1961), “On the Measurement of Technological Change”, *The Economic Journal*, vol. 71, no. 284, pp. 709–729.
- GEWEKE, J. and G. AMISANO (2008), “Comparing and evaluating Bayesian predictive distributions asset returns”, *Working Paper 969*, European Central Bank.
- HENDRY, D. F. and K. HUBRICH (2006), “Forecasting Economic Aggregates by Disaggregates”, *Working Paper 589*, European Central Bank.
- HULTEN, C. R. (1978), “Growth Accounting with Intermediate Inputs”, *The Review of Economic Studies*, vol. 45, no. 3, pp. 511–518.
- JAIMOVICH, N. and M. FLOETOTTO (2008), “Firm Dynamics, Markup Variations and the Business Cycle”, *Journal of Monetary Economics*, vol. 55, pp. 1238–1252.
- KÁTAY, G. and Z. WOLF (2008), “Driving Factors of Growth in Hungary - a Decomposition Exercise”, *Working Paper 6*, Magyar Nemzeti Bank.
- KOOP, G. (2003), *Bayesian Econometrics*, Wiley and Sons.
- LENTZ, R. and D. T. MORTENSEN (2005), “An Empirical Model of Growth through Product Innovation”, *Working Paper 11546*, National Bureau of Economic Research.
- LEVINSOHN, J. and A. PETRIN (2003), “Estimating Production Functions Using Inputs to Control for Unobservables”, *Review of Economic Studies*, vol. 70, no. 2, pp. 317–341.
- OLINER, S. D., D. E. SICHEL and K. J. STIROH (2007), “Explaining a Productive Decade”, *Finance and Economics Discussion Series 63*, Board of Governors of the Federal Reserve System.
- OLLEY, G. S. and A. PAKES (1996), “The Dynamics of Productivity in the Telecommunications Equipment Industry”, *Econometrica*, vol. 64, no. 6, pp. 1263–1297.

- O'MAHONY, M., B. V. ARK and G. YPMA (2007a), "The EUKLEMS Productivity Report", *mimeo*, 2007 march.
- O'MAHONY, M., M. P. TIMMER and B. V. ARK (2007b), "EUKLEMS Growth and Productivity Accounts: Overview November 2007 Release", *mimeo*.
- PETRIN, A. and J. LEVINSOHN (2008), "Measuring Aggregate Productivity Growth Using Plant-level Data", *mimeo*.
- STOCK, J. H. and M. W. WATSON (2002), "Forecasting Using Principal Components from a Large Number of Predictors", *Journal of the American Statistical Association*, vol. 97, pp. 1167–1179.
- WILDI, M. (1998), "Detection of Compatible Turning-Points and Signal-Extraction for Non-Stationary Timeseries", *Operations Research Proceedings 1998*, pp. 298–199.

5 Appendix

5.1 The data

5.1.1 CBS data

General description The microlevel database consists of a large panel of Dutch manufacturing firms. The database is under close surveillance by the Centrum voor Beleidsstatistiek²⁸. The time dimension of the panel spans between 1978 and 2004.

To compile our working data, we first established longitudinal links in the database because the CvB changed firm identifiers as of 1993. We used correspondences between firm identifiers included in the underlying datafiles. There was no documentation about such aspects of firm dynamics as mergers, acquisitions and split-ups so we could not control for these transformations.

²⁸Dutch Statistics Office

The industry classification system also changed in 1993. Since both the 1974-based ('old') and the 1993-based ('new') industry codes were available in 1990-1992, we generated a correspondence between the two industry classifications. We could then use this correspondence to generate deflators that originated from the EUKLEMS dataset, which in turn is based on the new industry classification.

Table 1 explains names of industry codes. In order not to lose observations and maintain a reasonable size, we merged industries 23 and 24.

Table 2 describes how many firms were used in the decompositions, the last columns of tables 3 and 4 show the how many observations were used in TFP-estimation. The last columns in tables 3 and 4 show firm-year observations whereas table 2 treats each firm as a single observation.

Variables The dataset did not have observed capital so we used depreciation as a proxy for the capital stock. We deflated nominal depreciation using an industry-specific implicit capital deflator that was calculated from EUKLEMS data. The other firm-specific micro-variables were: number of employees, output, value added, input materials, energy and payroll.

Outliers We filtered outliers using the interquartile-range method. This method is a standard way of dealing with skewed distributions. For TFP-estimation, we considered an observation to be outlier if either of the following was true:

$$\begin{aligned}x_o &< q25 - 1.5 * IQR \\x_o &> q75 + 1.5 * IQR,\end{aligned}$$

where $x_o \equiv \{K/Y; M/Y; WEXP/L\}$ and K denotes capital, M denotes input material purchases, $WEXP$ denotes wage expenditures, Y denotes either value added or output depending on which model we used. $q25$ and $q75$ stand for the lower and upper quartiles of the distribution of the variables in the set x_o , and $IQR \equiv q75 - q25$.

For aggregation, we used slightly different filtering. Since depreciation is not observed in the first six years of the panel, we had to backcast individual productivity growth

rates in 1978-1983. The backcasting procedure was based on the (geometric) average of non-outlier value-added/revenue labor-productivity growth. Using these "imputed" productivity numbers we carried out a secondary outlier filtering in order to get rid of implausibly large or small productivity-observations. The logic of this secondary filtering was the same as above with the exception that we deployed energy instead of depreciation.

5.1.2 The EUKLEMS data

History and general description The EUKLEMS "... project aims to create a database on measures of economic growth, productivity, employment creation, capital formation and technological change at the industry level for all European Union member states from 1970 onwards. This work will provide an important input to policy evaluation, in particular for the assessment of the goals concerning competitiveness and economic growth potential as established by the Lisbon and Barcelona summit goals. The database should facilitate the sustainable production of high quality statistics using the methodologies of national accounts and input-output analysis. The input measures will include various categories of capital, labour, energy, material and service inputs. Productivity measures will be developed, in particular with growth accounting techniques. Several measures on knowledge creation will also be constructed. Substantial methodological and data research on these measures will be carried out to improve international comparability. There will be ample attention for the development of a flexible database structure, and for the progressive implementation of the database in official statistics over the course of the project. The database will be used for analytical and policy-related purposes, in particular by studying the relationship between skill formation, technological progress and innovation on the one hand, and productivity, on the other. To facilitate this type of analysis a link will also be sought with existing micro (firm level) databases. The balance in academic, statistical and policy input in this project is realized by the participation of 15 organizations from across the EU, representing a mix of academic institutions and national economic policy research institutes and with the support from various statistical offices and the OECD."²⁹

²⁹<http://www.euklems.net>

Structure There have been two EUKLEMS releases so far. We used the one that was published in March 2008. The files contain 62 variables, including basic data and growth accounting variables. The datafiles are structured to follow an industry classification list which corresponds with the NACE³⁰ list.

Variables For a detailed description of the variables, see O’Mahony et al. (2007b) and O’Mahony et al. (2007a). This project used growth accounting variables in forecasting exercises and deflators in the nonparametric identification procedure. To be more specific, the following variables were used in forecasting experiments.

GO_QI: Gross output, volume indices, 1995 = 100,
VA_QI: Gross value added, volume indices, 1995 = 100,
LAB_QI: Labour services, volume indices, 1995 = 100,
CAP_QI: Capital services, volume indices, 1995 = 100,
TFPva_I: TFP (value added based), 1995=100,
TFPgo_I: TFP (gross output based) growth, 1995=100,
GO: Gross output at current basic prices (in millions of local currency),
VA: Gross value added at current basic prices (in millions of local currency).

The following were used for firm-level calculations at the first stage of the project:

GO_P: Gross output, price indices, 1995 = 100,
VA_P: Gross value added, price indices, 1995 = 100,
II_P: Intermediate inputs, price indices, 1995 = 100,
CAP: Capital services (in millions of local currency),
CAP_QI: Capital services, volume indices, 1995 = 100,

the last two were used to calculate implicit capital deflators for every industry.

³⁰NACE – Classification of Economic Activities in the European Community.

5.2 Bayesian analysis

5.2.1 Notation

Consider the following unrestricted VAR (1) specification

$$\begin{aligned}
 y_{1t} &= \phi_{11}y_{1t-1} + \cdots + \phi_{1p}y_{nt-p} + \varepsilon_{1t} \\
 y_{2t} &= \phi_{21}y_{1t-1} + \cdots + \phi_{2p}y_{nt-p} + \varepsilon_{2t} \\
 &\vdots \qquad \qquad \qquad \vdots \qquad \qquad \qquad \vdots \\
 y_{nt} &= \phi_{n1}y_{1t-1} + \cdots + \phi_{np}y_{nt-p} + \varepsilon_{nt}
 \end{aligned}$$

in a matrix representation as

$$y_t = \Phi y_{t-1} + \varepsilon_t$$

where y_t is an $n \times 1$ vector containing the variables of the VAR, Φ is a $n \times p$ matrix reflecting the dynamic structure of the VAR and ε_t is an *iid* mean-zero normal random vector with covariance matrix V . This last assumption says there is no autocorrelation left in the error term but the errors of different equations at time t can be correlated.

A more convenient representation of the above VAR is the following. Let $y_t = (y_{1t} \dots y_{nt})'$ be the t -th observation vector of dimension $(n \times 1)$ in the sample, $x_t = (y_{1t-1} \dots y_{nt-p})'$ be the $n \times p$ -vector containing the p lags of the n variables. Collect the ϕ coefficients equation-by-equation in the $(n \times p)$ -vectors β_1, \dots, β_n . The error vector stays the same $\varepsilon_t = (\varepsilon_{1t} \dots \varepsilon_{nt})'$. Then the above system can be written as

$$\begin{aligned}
 y_t &= X_t \beta + \varepsilon_t \\
 X_t &= I_n \otimes x_t = \begin{pmatrix} x_t & 0 & \cdots & 0 \\ 0 & x_t & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & \cdots & \cdots & x_t \end{pmatrix} \\
 \beta &= \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_n \end{pmatrix}, \forall t = 1 \dots T
 \end{aligned}$$

This implies that $y_t \in \mathbb{R}^{n \times 1}$, $\beta \in \mathbb{R}^{(n^2 \times p) \times 1}$, $X_t \in \mathbb{R}^{n^2 \times p}$ and $\varepsilon_t \in \mathbb{R}^{n \times 1}$.

5.2.2 Bayesian representation

The final objective of the forecasting exercise is to describe our beliefs about the joint evolution of future observations and parameters based on our sample up to time T . In other words, we want to characterize the posterior predictive density. We can write this as

$$p(y_{T+1}\dots y_{T+H}, V, \beta|y_T), \quad (14)$$

i.e. the joint posterior density of future observations and parameters conditional on observed data. Using standard laws and the direct consequence of Bayes' theorem we deconvolve (14) as

$$p(y_{T+1}\dots y_{T+H}, V, \beta|y_T) = p(y_{T+1}\dots y_{T+H}|V, \beta, y_T) p(V, \beta|y_T). \quad (15)$$

The first term of the right hand side represents the beliefs about the future realizations of the VAR, the second term is the posterior density of the parameters. Rewriting 15 gives an operational representation of y^{T+1} 's posterior predictive density:

$$p(y_{T+1}|y_T) = \int_0^\infty \int_{-\infty}^\infty p(y_{T+1}|V, \beta) p(V, \beta|y_T) d\beta dV \quad (16)$$

such that we integrate out the VAR-parameters β and the covariance matrix of the error vector ε_t .

5.2.3 Prior elicitation

To understand we first characterize the posterior $p(\beta, V|y_T)$ in 16. In the usual manner, this posterior can be written as

$$p(\beta, V|y_T) \propto p(y_T|\beta, V) p(\beta, V). \quad (17)$$

The first term on the right hand side of 17 is the Gaussian likelihood. The second term is the joint prior. We choose a natural conjugate joint prior specification because - when combined with a Gaussian likelihood - it yields full conditional posterior distributions

of the same family³¹. The most obvious natural conjugate prior for V is the Inverse Wishart (IW) distribution. In short, this prior elicitation is a matrix-generalization of the Gamma-prior in the univariate case. Similarly, choosing a normal prior for β results in a normal full conditional posterior. To sum up the priors, we write

$$\begin{aligned} V &\sim IW(\underline{V}, \underline{df}) \\ \beta &\sim N(\underline{\beta}, \underline{Q}) \end{aligned} \tag{18}$$

where the positive definite \underline{V} matrix plays the role of the scaling parameter and \underline{df} is degrees of freedom. The probability density function of the IW -distributed $n \times n$ matrix V looks like this:

$$p(V) = \frac{|\underline{V}|^{\underline{df}/2} |V|^{-(\underline{df}+n+1)/2} e^{-\text{trace}(\underline{V}V^{-1})/2}}{2^{\underline{df} \cdot n/2} \Gamma_n(\underline{df}/2)}$$

where $\Gamma_n(\underline{df}/2)$ is the multivariate gamma function.

Using 18, we can write out the joint prior as

$$p(\beta, V) \propto N(\underline{\beta}, \underline{Q}) IW(\underline{V}, \underline{df}), \tag{19}$$

and the joint posterior in 17 as

$$p(y_T | \beta, V) N(\underline{\beta}, \underline{Q}) IW(\underline{V}, \underline{df}). \tag{20}$$

Prior parameters $\underline{\beta}$, \underline{Q} , \underline{V} and \underline{df} A frequently used prior for macro VARs is the so called Litterman-Minnesota prior. Our priors are similar to but slightly different from this well-known prior. In the Litterman-Minnesota case a VAR is written in levels and the prior on the coefficient matrix is such that the first lag has a prior mean of 1, the other lags have prior mean zero. This prior assumes that there are n random walks in the system.

³¹See, for instance, Koop (2003).

Our endogenous variables value added/output, capital, labor and productivity are implied by the production function specifications. Our decompositions characterize aggregate productivity growth that we think can be described by the first difference of a random walk. Correspondingly, we take zero prior means, i.e. $\underline{\beta} = \mathbf{0} \in \mathbb{R}^{n \times 1}$ for the individual coefficients. We also assume that individual parameters are independent so that \underline{Q} is diagonal.

Our data approximate yearly growth rates by log-differences. A 0.3 log point change around zero is approximately equivalent to a growth rate of 0.26, a -0.3 log point change corresponds to a growth rate of -0.36. Moreover, if we move further away from zero along the real line, the log differences are even poorer approximations of the real growth rates. In terms of aggregate productivity and the other endogenous variables, these relatively small log-changes imply already so large shifts on the original scale of the variables that we expect it to happen very rarely. So it seemed reasonable to restrict prior variances of the shocks and the coefficients.

Therefore, the diagonal elements of the error variance-covariance matrix \underline{V} were set to 10^{-2} . The non-diagonal elements of \underline{V} were set to zero because we did not have a strong prior on the contemporary covariance structure of the errors. In terms of a scalar zero-mean normal distribution, this corresponds with 99% of the probability mass falling into the interval $[-0.258; +0.258]$. On the one hand, this is the range where there is a small chance of atypical shocks. On the other, with this prior variance setting, the majority of the support corresponds to the interval $[-0.2; +0.2]$, where the log approximation works reasonably well and where we would expect the principal part of the typical aggregate growth rates. The same argument holds for parameter variances \underline{Q} . Roughly speaking, nontypical β -draws would blow up/shrink already implausible values for y_t in the next period, which may result in forecasts of log-shifts changes that would imply implausibly large changes in our endogenous variables.

\underline{df} was also set to zero as this corresponds to the non-informative prior on the degrees-of-freedom parameter.

5.2.4 Posterior distributions, estimation and forecasting

Using the joint posterior we can develop a Gibbs-sampler based on the full conditional posterior distributions of β and V (see Koop (2003) for more details).

Full conditional distributions: $p(\beta|V, y_T)$ and $p(V|\beta, y_T)$ The full conditional distribution of the parameter vector β is proportional to a multivariate normal distribution. It is so because if V and y_T are nonrandom, the only thing that is random in the joint posterior (20) is β . $p(\beta|V, y_T)$ can be written as

$$\begin{aligned} \beta|V, y_T &\sim N(\bar{\beta}, \bar{Q}) \\ \bar{\beta} &= \bar{Q} \left(\underline{Q}^{-1} \underline{\beta} + \sum_{t=1}^T X_t' V^{-1} y_t \right) \\ \bar{Q} &= \left(\underline{Q}^{-1} + \sum_{t=1}^T X_t' V^{-1} X_t \right)^{-1}. \end{aligned}$$

The full conditional distribution of the error covariance matrix is proportional to the product of the Gaussian likelihood and the $IW(\underline{V}, df)$ prior about V .³² We use the fact that this product is also proportional to an $IW(\underline{V}_1^{-1}, df_1)$ -density with a new scale matrix $\underline{\Omega}_1^{-1}$ and degrees of freedom parameter df_1 :

$$V|\beta, y_T \sim IW(\bar{V}, \bar{df}),$$

where the updated scale and the degrees-of-freedom parameters \bar{V} and \bar{df} are calculated as

$$\begin{aligned} \bar{V} &= \left(\underline{V}^{-1} + \sum_{t=1}^T (y_t - X_t \beta) (y_t - X_t \beta)' \right)^{-1} \\ \bar{df} &= T + \underline{df}. \end{aligned}$$

The Gibbs-sampler amounts to iterating the following algorithm until convergence:

³²Again, based on the joint posterior distribution given by (20).

1. Give starting values to $\mu_0 = (\beta^{(0)}, V^{(0)})$.
2. Simulate $\beta^{(m+1)}$ from $p(\beta|V^{(m)}, y_T)$ and then simulate $V^{(m+1)}$ from $p(V|\beta^{(m+1)}, y_T)$.
3. Set $m = m + 1$ and go to step 2.

Forecasting After the Markov-Chain has converged we can consider the elements of the chain (after discarding the burn-in observations and thinning the chain) as sample from the joint posterior. Now we turn to the first term of the right hand side of (15): $p(y_{T+1} \dots y_{T+H}, V, \beta | y_T)$. Conditional on $y = (y_1, \dots, y_T), \beta$ and V , ε is distributed normally along the future path of the VAR, i.e., y_{T+k} is also conditionally normally distributed. So in order to obtain draws from the forecast distribution of y_t we have to simulate future paths of y_{t+i} $i = 1 \dots k$. To this end, we take each element, indexed by (m) , of the Markov-chain after the burn-in value ($m \geq R$) and thinning. A trajectory at time T can be constructed at the (m) -th element of the Markov-chain as:

1. Draw the random vectors $\varepsilon_{t+1}^{(m)}, \varepsilon_{t+2}^{(m)}, \dots, \varepsilon_{t+k}^{(m)} \sim N(0, V^{(m)})$
2. Calculate

$$\begin{aligned} y_{t+1}^{(m)} &= \beta^{(m)} y_t + \varepsilon_{t+1}^{(m)} \\ y_{t+2}^{(m)} &= \beta^{(m)} y_{t+1}^{(m)} + \varepsilon_{t+2}^{(m)} \\ &\vdots \\ y_{t+k}^{(m)} &= \beta^{(m)} y_{t+k-1}^{(m)} + \varepsilon_{t+k}^{(m)} \end{aligned}$$
3. Go to the $(m + 1)$ -th element of the Markov-chain.
4. Calculate the appropriate moments of the forecast distribution from $y_{t+i}^{(m)}$ $i = 1 \dots k$, $m \geq R$.

In this manner, we have a simulated trajectory for each element of the Markov-chain so we can calculate moments of the forecast distribution. The forecast distribution approximated this way fully reflects parameter uncertainty.³³

³³If we forecast the system at the mean or median of the posterior error covariance matrix V (parameter vector β), the forecast distributions miss the uncertainty in the coefficient vector β (error covariances V).

Tables and figures

Figure 1: Tfp-growth in overall manufacturing in different countries.

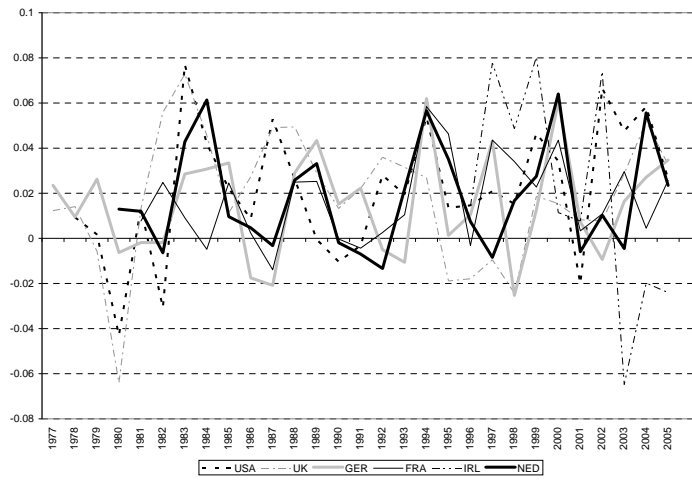


Figure 2: Tfp growth by micro-based and EUKLEMS-based aggregate; overall manufacturing.

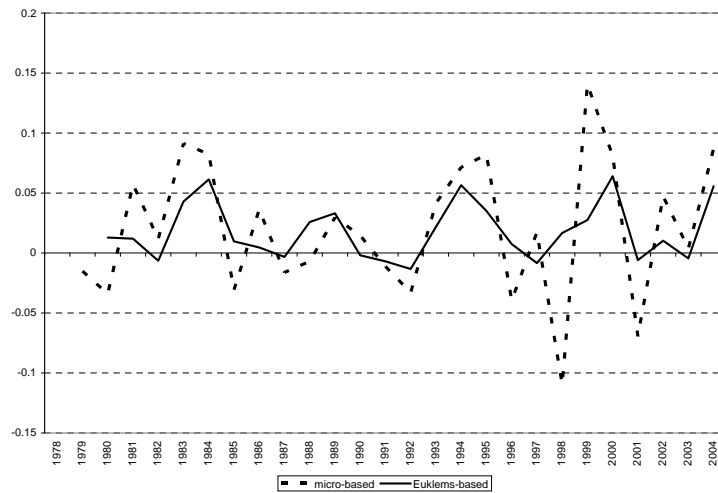


Figure 3: Value added growth by micro-based and EUKLEMS-based aggregates; overall manufacturing (solid line: micro-based aggregate; dashed line: EUKLEMS aggregate).

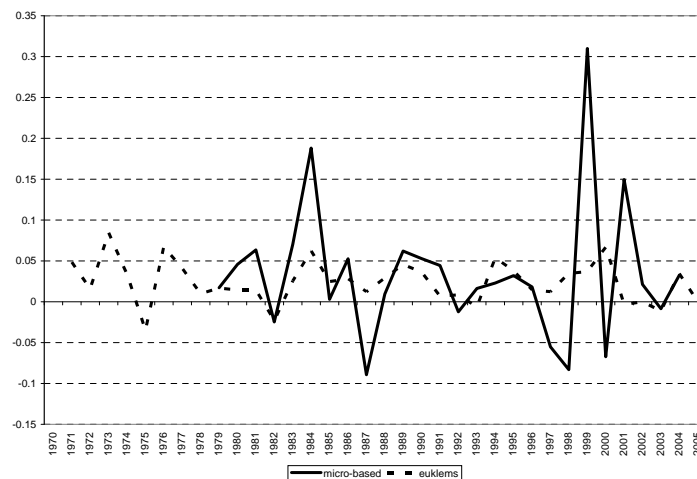


Figure 4: Labor input growth by micro-based and EUKLEMS-based aggregates; overall manufacturing (solid line: micro-based aggregate; dashed line: EUKLEMS aggregate).

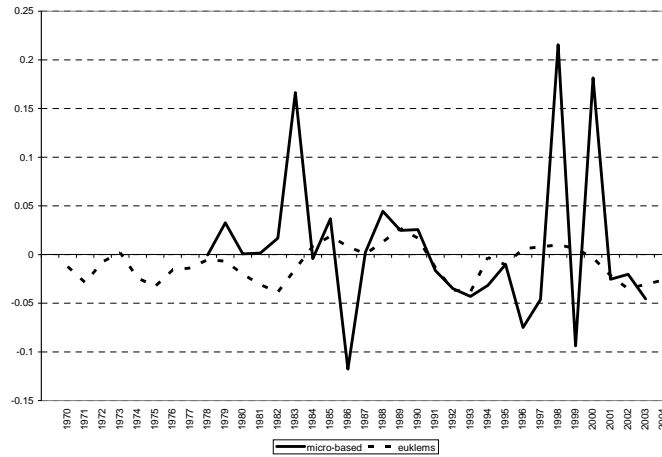


Figure 5: EUKLEMS aggregate tfp index (1995=100) and its HP-trend

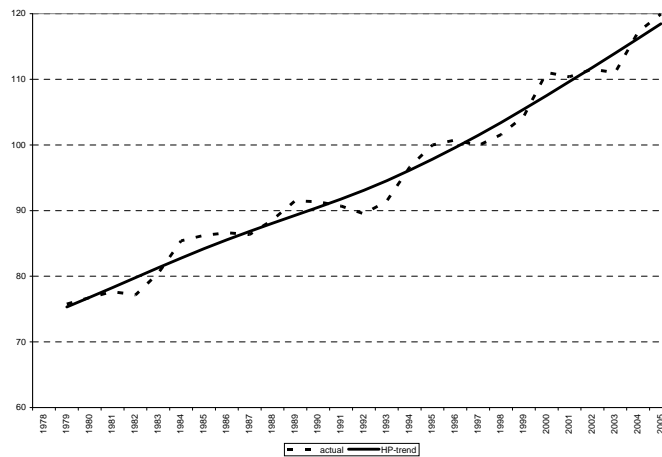


Figure 6: Forecast distributions of yearly tfp growth rates on a three-year horizon (1998-2000).

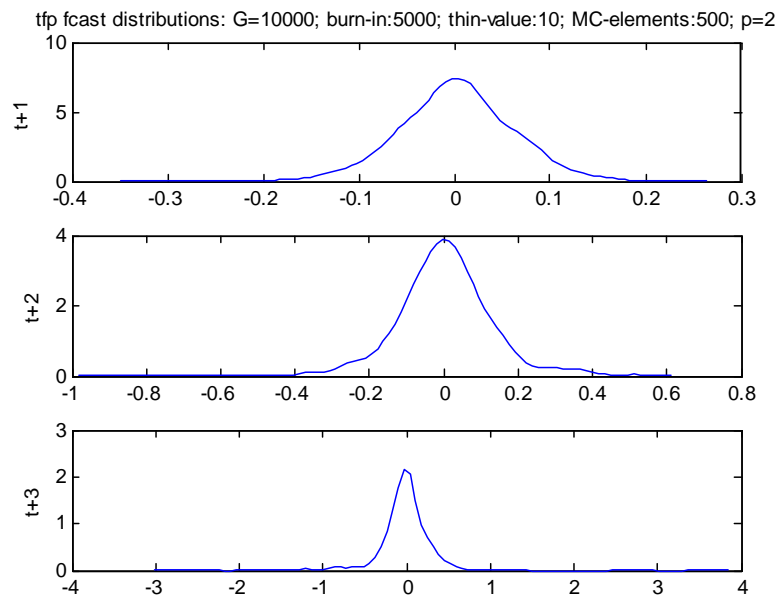


Table 1: Key to industry codes

Manufacturing classification	
<i>industry code</i>	
Consumer manufacturing	
15t16	Food products, beverages and tobacco
17t19	Textiles, textile products, leather and footwear
36t37	Manufacturing nec; recycling
Intermediate manufacturing	
20	Wood and products of wood and cork
21t22	Pulp, paper, paper products, printing and publishing
23	Coke, refined petroleum products and nuclear fuel
24	Chemicals and chemical products
25	Rubber and plastics products
26	Other non-metallic mineral products
27t28	Basic metals and fabricated metal products
Investment goods, excluding hightech	
29	Machinery, nec
34t35	Transport equipment
Electrical machinery and communication services	
30t33	Electrical and optical equipment

Source: Table 4 in O'Mahony et al. (2007a).

Table 2: Number of firms in the analysis

industry code	original	va model	output model
15t16	5994	4284	4421
17t19	3823	2503	2429
20	1321	1040	1038
21t22	5908	4064	4224
24	1585	1088	1132
25	1713	1307	1321
26	2076	1395	1472
27t28	7617	5117	5246
29	4804	3358	3367
30t33	3296	2140	2197
34t35	2291	1504	1534
36t37	3743	2205	2265
overall			
manufacturing	44171	30005	30646

Table 3: Estimation results - Value added model

industry code	L	K	M	ctrs	seL	seK	nobs
15t16	0.52	0.22	.	581.607	0.007	0.014	19745
17t19	0.54	0.23	.	390.991	0.008	0.011	9457
20	0.58	0.12	.	431.116	0.017	0.015	4624
21t22	0.60	0.13	.	660.952	0.009	0.006	20196
24	0.44	0.23	.	101.387	0.016	0.023	6390
25	0.53	0.24	.	84.735	0.019	0.015	6800
26	0.52	0.21	.	191.860	0.013	0.016	7094
27t28	0.65	0.12	.	1056.116	0.008	0.008	24428
29	0.62	0.12	.	458.341	0.012	0.010	18231
30t33	0.65	0.16	.	211.017	0.013	0.010	8008
34t35	0.72	0.06	.	183.449	0.014	0.010	6861
36t37	0.59	0.13	.	322.399	0.013	0.012	9033

Table 4: Estimation results - Output model

industry code	L	K	M	ctrs	seL	seK	nobs
15t16	0.25	0.12	0.65	2.6439	0.013	0.019	19745
17t19	0.40	0.16	0.42	2.9865	0.019	0.022	9457
20	0.31	0.01	0.66	0.0296	0.012	0.007	4624
21t22	0.36	0.07	0.56	0.0657	0.008	0.033	20196
24	0.23	0.13	0.59	2.0937	0.017	0.153	6390
25	0.33	0.02	0.59	1.8747	0.011	0.047	6800
26	0.34	0.02	0.69	0.0368	0.010	0.152	7094
27t28	0.45	0.01	0.54	0.0002	0.008	0.003	24428
29	0.42	0.06	0.53	0.4290	0.007	0.018	18231
30t33	0.47	0.05	0.63	2.8980	0.013	0.036	8008
34t35	0.44	0.01	0.55	0.0194	0.015	0.000	6861
36t37	0.38	0.1	0.37	6.7575	0.013	0.020	9033

Table 5: Average tfp-growth in overall manufacturing in different countries, source: EU-KLEMS database (average percent per annum rates over period indicated in first column)

	USA	UK	GER	FRA	IRL	NED
1981-1985	0.024	0.039	0.018	0.012	#N/A	0.024
1986-1990	0.015	0.034	0.010	0.008	#N/A	0.012
1991-1995	0.022	0.019	0.014	0.022	#N/A	0.018
1996-2000	0.026	-0.004	0.020	0.028	0.046	0.021
2001-2005	0.035	0.026	0.015	0.015	-0.006	0.016

Table 6: Out-of-sample forecast performance, percentage point deviation from HP-trend of aggregate manufacturing productivity growth, estimation sample: 1978-1997, forecast period: 1998-2000

structure	microcomp.		industry comp.	aggregate	
	dc1	dc2		dc1	dc2
observed VA-shares	-2.4	-2.2	-3.5	-3.2	-3.3
observed input shares	-0.9	-0.9	-2.8	-3.1	-3.1
fitted VA shares & tau	-1.7	-1.4	-4.7	-4.6	-4.7
fitted input shares & tau	-1.2	-0.9	-3.3	-3.3	-3.4

dc1: equation 2; dc2: equation 6; microcomp. forecast is implied by equation 11; industry comp. is implied by equation 12; the aggregate forecast was computed from its own autoregressive model.

Table 7: Out-of-sample forecast performance, percentage point deviation from HP-trend of aggregate manufacturing productivity growth, estimation sample: 1978-2000, forecast period: 2001-2003

structure	microcomp.		industry comp.	aggregate	
	dc1	dc2		dc1	dc2
observed VA-shares	-2.9	-1.7	-2.9	-3.2	-3.6
observed input shares	-2.1	-2.6	-3.0	-3.0	-3.0
fitted VA shares & tau	-5.0	-5.6	-8.5	-8.8	-8.9
fitted input shares & tau	-3.3	-3.9	-5.5	-5.7	-5.7

dc1: equation 2; dc2: equation 6; microcomp. forecast is implied by equation 11; industry comp. is implied by equation 12; the aggregate forecast was computed from its own autoregressive model..

Table 8: Predictive distribution statistics over forecast horizon 1998-2000, macro VAR on growth rates

	Variable	Mean	Stdev	NSE	Median
t+1					
	va/q	-0.0088	0.0803	0.0025	-0.0093
	cap	-0.005	0.0655	0.0021	-0.0048
	lab	-0.0008	0.0527	0.0017	0.0003
	tfp	0.0037	0.0651	0.0021	0.006
t+2					
	va/q	0.0208	0.1774	0.0056	0.0137
	cap	0.0214	0.1302	0.0041	0.0136
	lab	-0.0021	0.119	0.0038	-0.0047
	tfp	0.0072	0.1408	0.0045	0.0011
t+3					
	va/q	-0.0531	0.4408	0.0139	-0.006
	cap	-0.0604	0.3486	0.011	-0.0105
	lab	0.0054	0.25	0.0079	0.0011
	tfp	-0.0353	0.3568	0.0113	-0.0061
Average					
	va/q	-0.0137	0.2328	0.0074	-0.0005
	cap	-0.0147	0.1814	0.0057	-0.0006
	lab	0.0008	0.1406	0.0044	-0.0011
	tfp	-0.0081	0.1876	0.0059	0.0004

Gibbs-iterations:10000; burn-in:5000; thin-value:10; MC-elements:500.

Growth rates are approximated by logarithmic differences

Table 9: Descriptives of overall trajectory distributions over forecast horizon 1998-2000, macro-VAR on growth rates

Variable	Mean	Stdev	NSE	Median
va/q	-0.0071	0.11	0.0035	-0.0012
cap	-0.0091	0.101	0.0032	-0.0039
lab	0.0036	0.0751	0.0024	-0.002
tfp	-0.0031	0.0969	0.0031	-0.0017

G=10000; burn=5000; thin=10; MC-elements:500.

Table 10: Bayesian forecast results: average value added tfp-growth rates for overall EU-KLEMS manufacturing (percentage point deviation from HP-trend, estimation sample: 1978-1997, forecast period: 1998-2000)

structure	VAR(1,2)*		VAR(2,1)**	
	dc1	dc2	dc1	dc2
VAR - micro & macro				
va/q shares	0.3	1.5	1.0	6.0
input shares	-0.9	2.3	-1.2	0.4
fitted va/q shares & tau	0.4	-1.5	-0.9	2.4
fitted input shares & tau	0.6	0.0	-0.1	4.7
<i>PBF-weighted average of fcasts:</i>		0.3		0.3
VAR - macro		-2.3		-2.3

*VAR(1,2): 1st lag of endogenous variables (growth rates of aggregate capital services, labor services, tfp, value added), 1st and 2nd lags of exogenous variables. Exogenous variables are lagged values of components in the micro-based decompositions (ie.

**VAR(2,1): 1st and 2nd lags of endogenous variables (growth rates of aggregate capital services, labor services, tfp, value added), 1st lag of exogenous variables. Exogenous variables are lagged values of components in the decompositions (ie. $X(t-1)$).

Table 11: Bayesian Model Averages of forecast results: value added tfp-growth rates for overall manufacturing, percentage point deviations from HP-trend. Averages are calculated using all models used in a forecast window

Forecast horizon	VAR-microcomponents	VAR-aggregate
1998-2000	-0.7	-1.8
1999-2001	-1.1	-1.9
2000-2002	-1.4	-1.9
2001-2003	-1.5	-2.1