Uncertainty, Misallocation and the Life-cycle Growth of Firms*

Eero Mäkynen† Oskari Vähämaa‡

June 26, 2026

Abstract

We propose a decomposition of static misallocation that distinguishes between idiosyncratic uncertainty and ex ante misallocation generated by tax-like distortions. Using profits-to-wage-bill and value-added-to-wage-bill ratios, we identify the two sources empirically. In the comprehensive Finnish firm-level data, uncertainty accounts for at least 40% of measured ex post misallocation, and the dispersion of prediction errors declines sharply with firm age. These empirical patterns are captured by a model of firm growth with learning, calibrated to standard life-cycle facts. The model implies that uncertainty reduces aggregate output and TFP by about 14%.

Keywords: firm dynamics, uncertainty, misallocation

JEL Codes: D24, E23, L11, O47

* We are indebted to William Kerr, Vincent Sterk, Akos Valentinyi and our discussant Paul Klein for detailed feedback. We would also like to thank Philippe Aghion, Niku Määttänen, Morten Ravn, Matti Sarvimäki, Petr Sedláček, Juhana Siljander, Juuso Välimäki and conference participants at CEF, EEA-ESEM, ESWC, HMRAD, Lisbon Macro Workshop and Nordic Summer Symposium in Macroeconomics for helpful comments and suggestions. Oskari Vähämaa acknowledges financial support from the Emil Aaltonen Foundation. First version: May 2021.

† University of Turku. Email: eeanma@utu.fi

‡ Corresponding author. University of Helsinki. Address: Arkadiankatu 7 (P.O. Box 17), 00014 University of Helsinki. Email: oskari.vahamaa@gmail.com

1 Introduction

How much aggregate output is lost because inputs are allocated to the wrong firms? A seminal paper by Restuccia and Rogerson (2008) illustrates that inefficient allocation of input factors across production units – misallocation – can have severe effects on TFP, a key determinant of cross-country differences in output.[1] To evaluate the empirical relevance of this channel, a popular indirect approach, pioneered by Hsieh and Klenow (2009), henceforth HK, is to measure marginal products of labor and capital using firm-level micro data.[2] If an economy's dispersion of marginal products, at least for the part that exceeds the US benchmark, reflects misallocation, the reallocation of input factors could lead to a significant increase in TFP and output. However, there are other factors that can generate dispersion in marginal products that are not necessarily directly related to inefficient allocation of input factors across producers.[3]

This paper studies one such source of dispersion: firms choosing their inputs before observing their current fundamentals. A firm that turns out to be more productive than expected will appear, ex post, to have too few inputs given its realized marginal product. Conversely, a firm that turns out to be less productive than expected will appear to have too many inputs. These outcomes generate dispersion in marginal products even if firms face no distortions at the time when inputs are chosen. Treating all such dispersion as ex ante misallocation therefore confounds imperfect information with distortions.

We show how firm-specific prediction errors can be separated from distortionary tax-like revenue wedges in a static HK-style framework by exploiting information in profits. Specifically, profits-to-wage-bill and value-added-to-wage-bill ratios together identify the two channels. The intuition behind our identification strategy is that tax-like distortions affect profits and the wage bill in a similar way, implying that the profits-to-wage-bill ratio is independent of the revenue wedge. The effects of imperfect information, by contrast, are asymmetric. Profits depend on both realized and expected fundamentals, while employment depends only on expected fundamentals. This asymmetry means that the profits-to-wage-bill ratio is informative about prediction errors. Finally, the value-added-to-wage-bill ratio together with prediction errors gives us the tax-like wedges. As a result, our accounting framework allows us to decompose ex post misallocation into variation in prediction errors

[1]See, e.g., Klenow and Rodríguez-Clare (1997) for the importance of TFP in explaining differences in output across countries.

[2]See, e.g., Bayer et al. (2018), Busso et al. (2013) and Cirera et al. (2020) for recent examples of this approach.

[3]For example, measurement errors, different production/demand structures and adjustment costs could affect the measured misallocation. See e.g. Bils et al. (2021), Rotemberg and White (2021), Gollin and Udry (2020), Haltiwanger et al. (2018), Bartelsman et al. (2013) and Asker et al. (2014).

(the uncertainty channel), variation in revenue distortions (ex ante misallocation), and the covariance between the two wedges.

We apply our framework to administrative data, collected for taxation purposes, covering nearly all Finnish firms over 1995–2012, and document two key empirical facts. First, uncertainty is as important as ex ante misallocation, both accounting for 44% of ex post misallocation. Second, uncertainty strongly diminishes with firm age. The dispersion of prediction errors is more than halved when one moves from new businesses to firms that have been operating for a decade. This decline mirrors the parallel decline in the dispersion of the HK revenue wedge.[4] Misallocation, on the other hand, is practically constant across age brackets. Decomposing the HK dispersion conditional on age requires a large set of firms with varying ages; thus, observing almost the whole population of firms is important.

We also examine how the HK revenue wedge relates to firm productivity. In line with earlier empirical literature, this relationship is negative. We find, however, that roughly half of it reflects measured prediction errors rather than ex ante distortions: high-productivity (low-productivity) firms tend to underestimate (overestimate) their realized productivity, contributing to the apparent "tax" pattern. We also document that firms with larger prediction errors are more likely to exit in subsequent periods. Together, these findings suggest that the prediction errors are related to firm fundamentals and firm behavior.

Our static decomposition leaves open the question of how much uncertainty actually matters for aggregate output and productivity. Answering this requires a dynamic model with learning and an endogenous age-size distribution. We set up a model with firm entry and exit, embedding the learning structure of Jovanovic (1982) in a general equilibrium framework similar to Hopenhayn and Rogerson (1993) and Melitz (2003).

The key features of our model are age-dependent uncertainty, tax-like wedges and convex adjustment cost. In line with our static exercise, firms choose inputs before observing their current-period TFP.[5] Moreover, this productivity is a combination of a persistent component and a transitory one, which firms are unable to separate from each other. They use Bayesian learning to form a forecast about their current persistent productivity. This structure implies that the dispersion of marginal products is smaller for older firms whose forecasts are more precise. We also allow for adjustment costs, which further generate age-dependent dispersion, as some authors have suggested this channel of misallocation might be important.[6] Tax-like revenue wedges, drawn at entry and persistent thereafter, capture ex ante distortions in our

[4]This is in line with Eslava et al. (2023), who observe that idiosyncratic distortions are particularly important for young plants in Colombian establishment-level data.

[5]With risk-neutral firms, this is equivalent to assuming one-period time-to-build for inputs, a setup used recently, e.g., by Boar et al. (2022).

[6]See, e.g., Eslava et al. (2023) or Asker et al. (2014).

baseline.

The model is empirically disciplined by targeting key life-cycle features of firm dynamics that are not directly related to misallocation: growth profiles of young firms, exit rates across age brackets, employment autocorrelations, and the size distribution. The calibrated model also fits the growth profiles of older firms, which are not used as targets.

The calibrated model reproduces our static decomposition. Applied to simulated data, the decomposition accounts for 93% of the observed variation in the HK revenue distortion, with uncertainty and ex ante misallocation shares close to their empirical counterparts and a decreasing age-dependent trend in uncertainty that quantitatively matches the data. This shows that the patterns we identify in the accounting exercise can emerge endogenously from a parsimonious learning model.

We then use the calibrated model to evaluate the aggregate consequences of uncertainty and misallocation. Under our baseline calibration, uncertainty alone reduces aggregate output and TFP by about 14%. This effect is robust across alternative modeling choices, including the calibration of the productivity process and the introduction of more flexible wedge structures with transitory and correlated components. Ex ante misallocation has a larger effect under our baseline calibration with permanent wedges, reducing output (TFP) by 41% (46%), but this number drops to a level comparable to the cost of uncertainty under more flexible wedge specifications. Adjustment costs alone reduce aggregate output and TFP by about 2%.

Literature. The paper relates to three branches of literature. First, it contributes to work that interprets static HK-style measures. A growing literature shows that measured TFPR dispersion can reflect forces other than misallocation; Haltiwanger et al. (2018) consider more general demand and production structures, Bils et al. (2021) and Rotemberg and White (2021) emphasize measurement error, Gollin and Udry (2020) separate measurement error, unobserved heterogeneity, and misallocation, and Baqaee and Farhi (2020) allow for flexible input-output linkages and varying substitutability in a non-parametric framework. Eslava et al. (2023) use price and quantity data to decompose wedges into input-price, markup, and residual components, and also document that measured wedges are larger for young plants. Our contribution to this branch is to use profits to separate idiosyncratic uncertainty from misallocation generated by tax-like wedges, and to show in comprehensive data from a developed country like Finland that uncertainty can be a quantitatively important source of ex post misallocation with a strong age-dependent trend.

Second, the paper is connected to structural work that quantifies specific sources of misallocation.[7] Especially relevant are structural analyses exploring the connection between

[7] Midrigan and Xu (2014) study financial frictions, Asker et al. (2014) analyze adjustment costs, and

misallocation and information frictions. David et al. (2016) and David and Venkateswaran (2019) study how imperfect information distorts investment and capital allocation. Feng (2025) is also closely related in documenting declining misallocation over the firm life cycle and interpreting it through learning. Relative to this work, our contribution is to bring direct empirical discipline from a static decomposition of measured wedges into a dynamic misallocation model. The accounting exercise separates uncertainty from ex ante wedges empirically; the structural model then quantifies how these channels matter for aggregate output and TFP in a setting with adjustment costs, entry, exit, learning, and selection.[8]

A few recent papers emphasize the connections between productivity investments and misallocation. Hsieh and Klenow (2014) allow firms' TFP to evolve endogenously over the life cycle.[9] In their approach, misallocation can severely discourage productivity investment and thus dampen aggregate TFP. In Bento and Restuccia (2017), firms can invest in productivity not just along the life cycle but also upon entry. Peters (2020) considers the effects of market power in generating misallocation; in his setup, the firms' market power is endogenous and evolves over the life cycle. In contrast, our mechanism does not require firms to underinvest in productivity, quality, or market access; measured wedges change with age because firms become better informed about their fundamentals.

Third, the paper relates to work on firm life-cycle growth and endogenous firm dynamics. Clementi and Palazzo (2016) and Sterk et al. (2021) study mechanisms that generate firm growth and exit over the life cycle, focusing on adjustment costs and ex ante heterogeneity, respectively. Our model, similar to Arkolakis et al. (2018), embeds the learning-and-selection mechanism of Jovanovic (1982) in a general equilibrium environment close to Hopenhayn and Rogerson (1993) and Melitz (2003). On the theoretical side, Tian (2022) studies firm dynamics under uncertainty. Gao and Zhang (2024) use direct subjective forecast data to discipline entrepreneurial learning, sharing our emphasis on empirical discipline of the learning mechanism.

The rest of the paper proceeds as follows. Section 2 develops the accounting framework and applies it to Finnish firm-level data, including the alternative identification strategy with materials and the robustness exercises. Section 3 presents the dynamic general equilibrium model with learning, entry, exit, adjustment costs, and revenue wedges. Section 4 calibrates the model, evaluates its fit, compares the model-implied and empirical decompositions, and reports the main aggregate counterfactuals. Section 5 studies correlated and

Bartelsman et al. (2013) emphasize allocation and selection with quasi-fixed factors and overhead labor.

[8] Selection can be important for the aggregate effects of micro-level misallocation, as shown by Yang (2021).

[9] They build on Atkeson and Burstein (2010) GE model. The approach is close to papers that examine the role of investments in organizational capital and customer base. See, e.g., Foster et al. (2016).

transitory wedges together with an alternative calibration of the productivity process. Section 6 concludes.

2 Accounting for Uncertainty and Misallocation

In this section, we develop an accounting framework to jointly measure misallocation and uncertainty in a static setup with a minimum amount of theory by just specifying a production function and a demand structure. We link our measure to the dispersion of the HK revenue wedge, a standard measure of misallocation when there are no distortions that affect capital and labor asymmetrically. Next, we apply our approach to Finnish firm-level data and show that uncertainty accounts for at least 40% of ex post misallocation. We also find that there is a strong age-dependent trend in uncertainty. Finally, we explore the sensitivity of our results.

2.1 A Theoretical Accounting Framework

This subsection develops a simple way to indirectly measure the wedges that uncertainty and tax-like revenue distortions can generate in static first-order conditions. Firms face CES demand and produce with Cobb-Douglas technology using labor and capital. They choose their inputs under imperfect information about the current-period fundamentals. To keep the exposition of the framework as simple as possible, we only focus on firms' uncertainty over their productivity. However, since we do not have data on quantities, this interpretation is observationally equivalent to assuming that demand is also uncertain. In addition, firms face idiosyncratic revenue distortions that are known to them at the time they make their production decisions.

There is a large number of firms, indexed by i, each producing a differentiated good. Individual goods are aggregated to a single final good with the CES aggregator. Thus, firm i in industry s at time t faces the isoelastic demand curve given by

$$y_{t,s,i} = \left(\frac{p_{t,s,i}}{P_t}\right)^{-\sigma} Y_t, \tag{1}$$

where $p_{t,s,i}$ is the price of the good, P_t is the price index and Y_t is the amount of the final good consumed.

The production technology of each firm is represented by a Cobb-Douglas production function of a firm's TFP, $z_{t,s,i}$, labor, $n_{t,s,i}$, and capital, $k_{t,s,i}$.

$$y_{t,s,i} = z_{t,s,i} n_{t,s,i}^{1-\alpha_s} k_{t,s,i}^{\alpha_s}, \tag{2}$$

where the capital intensity, α_s, is allowed to vary across industries. In the empirical section, we also consider a specification of the production function where the capital intensity is allowed to vary across firms and over time.

Given the demand system and the production function, firm i maximizes expected profits:

$$E(\pi_{t,s,i}) = Y_t^{\frac{1}{\sigma}} P_t (1-\tau_{t,s,i}) E(z_{t,s,i}^{\frac{\sigma-1}{\sigma}}) (n_{t,s,i}^{1-\alpha_s} k_{t,s,i}^{\alpha_s})^{\frac{\sigma-1}{\sigma}} - w_{t,s,i} n_{t,s,i} - R_t k_{t,s,i}, \tag{3}$$

where $\tau_{t,s,i}$ is a tax-like idiosyncratic distortion that firm i faces in period t. We allow wage rates, $w_{t,s,i}$, to be specific to each firm and time period because we use the wage bill as an observable variable in the empirical part of the paper. This approach allows the observed variation in wage bills to reflect either employment levels or wage rates. The rental rate, R_t, in contrast, is assumed to be common across firms in our baseline calculations (Section 2.3). In Section 2.5, we relax this assumption and use firm-specific implicit interest rates as a proxy for firm-time-specific rental rates.

From the first-order condition with respect to labor, we get the following familiar expression for the ex post distortion:

$$1 - \tau_{t,s,i}^{HK} \equiv (1-\tau_{t,s,i})(1-\varphi_{t,s,i}) = \frac{\sigma}{(\sigma-1)(1-\alpha_s)\frac{p_{t,s,i} y_{t,s,i}}{w_{t,s,i} n_{t,s,i}}}, \tag{4}$$

where $1-\varphi_{t,s,i}$ measures the prediction error of the firm,

$$1 - \varphi_{t,s,i} \equiv \frac{E(z_{t,s,i}^{\frac{\sigma-1}{\sigma}})}{z_{t,s,i}^{\frac{\sigma-1}{\sigma}}}.$$

To write this equation in terms of observable value added, $p_{t,s,i} y_{t,s,i}$, we have multiplied and divided the first-order condition by $z_{t,s,i}^{\frac{\sigma-1}{\sigma}}$. Note that this is the measure of revenue distortion in HK.

Equation (4) emphasizes that, without knowledge of the firm's information set at the time inputs are chosen, the observed ratio of realized value added, $p_{t,s,i} y_{t,s,i}$, and the wage bill, $w_{t,s,i} n_{t,s,i}$, is not sufficient to identify the tax-like distortion, $1-\tau_{t,s,i}$. Moreover, since capital is also chosen under uncertainty, the first-order condition with respect to it does not help to disentangle the prediction error from the revenue tax. The following proposition shows that this identification problem can be resolved by using additional information contained in realized profits.

Proposition 1 (Decomposition of the HK wedge) *Given the observable ratios of profits to the wage bill,* $\pi_{t,s,i}/(w_{t,s,i}n_{t,s,i})$*, and of value added to the wage bill,* $p_{t,s,i}y_{t,s,i}/(w_{t,s,i}n_{t,s,i})$*, the prediction error,* $1-\varphi_{t,s,i}$*, and the tax-like wedge,* $1-\tau_{t,s,i}$*, are separately identified by the following expressions:*

1. $$1-\varphi_{t,s,i}=\frac{\sigma}{\sigma-1}\left[\frac{1}{1+(1-\alpha_s)\frac{\pi_{t,s,i}}{w_{t,s,i}n_{t,s,i}}}\right],$$
2. $$1-\tau_{t,s,i}=(1-\varphi_{t,s,i})^{-1}\frac{\sigma}{(\sigma-1)(1-\alpha_s)\frac{p_{t,s,i}y_{t,s,i}}{w_{t,s,i}n_{t,s,i}}}.$$

The second expression follows directly from equation (4). It also illustrates why we sometimes refer to $\tau_{t,s,i}$ as a residual wedge. To derive the first expression, note that realized profits can be written in terms of the firm's optimal input choices:

$$\pi_{t,s,i}=(1-\tau_{t,s,i})p_{t,s,i}y_{t,s,i}(1-\frac{\sigma-1}{\sigma}(1-\varphi_{t,s,i})). \tag{5}$$

Next, note that equation (4) implies that

$$w_{t,s,i}n_{t,s,i}=(1-\alpha)(1-\tau_{t,s,i})(1-\varphi_{t,s,i})\frac{\sigma-1}{\sigma}p_{t,s,i}y_{t,s,i}. \tag{6}$$

Dividing equation (5) by equation (6) and solving for $1-\varphi_{t,s,i}$ gives the first part of the proposition.

Intuitively, both the wage bill and realized profits depend on $(1-\tau_{t,s,i})$ in a similar way. Thus, the ratio of profits to the wage bill is independent of the distortion responsible for ex ante misallocation. The effects of $(1-\varphi_{t,s,i})$, by contrast, are asymmetric for $\pi_{t,s,i}$ and $w_{t,s,i}n_{t,s,i}$. This occurs because optimal employment is determined by expected fundamentals, whereas realized profits depend on both expectations and realized fundamentals.[10]

Without capital frictions, variation in the log of the HK revenue distortion, calculated using equation (4), corresponds to variation in the log of TFPR, which is the standard indirect measure of misallocation.[11] In our case, this represents a measure of ex post misallocation.

The following remark highlights that this ex post measure can result from three distinct components.

[10]The realized fundamental, $z_{t,s,i}^{\frac{\sigma-1}{\sigma}}$, enters the wage-bill equation only because we have written it in terms of actual value added.

[11]See HK for details.

Remark 1 (Ex post misallocation and its components) *With the help of Proposition 1, we can decompose the measure of ex post misallocation,* $\mathrm{Var}(\ln(1-\tau^{HK}))$*, as*

$$\mathrm{Var}(\ln(1-\tau^{HK})) = \mathrm{Var}(\ln(1-\varphi)) + \mathrm{Var}(\ln(1-\tau)) + 2\,\mathrm{Cov}(\ln(1-\varphi), \ln(1-\tau)).$$

This decomposition shows that ex post misallocation can be separated into components reflecting uncertainty (the variance of the prediction error in logs), ex ante misallocation (the variance of the tax-like distortion in logs), and the covariance between the two.

In the empirical section, we also explore the association between wedges ($1-\tau^{HK}$, $1-\varphi$, and $1-\tau$) and productivity. To do so, we need to measure firm-level productivity. In the setting considered in this section, productivity is measured as

$$z_{t,s,i} = \frac{(p_{t,s,i} y_{t,s,i})^{\frac{\sigma}{\sigma-1}}}{n_{t,s,i}^{1-\alpha_s} k_{t,s,i}^{\alpha_s}} A_t, \tag{7}$$

where A_t is a common factor across firms that depends on aggregate prices P_t and aggregate output Y_t.

2.2 The Data

We analyze static misallocation in Finland using register data. Our identification of the different components of ex post misallocation, as specified in the previous section, is based on the ratios of value added to the wage bill and of profits to the wage bill. This approach is particularly suitable for our data, as these variables are collected for taxation purposes and are therefore measured with high precision and available for the vast majority of firms.

We use annual firm-level data from the Financial Statement Statistics for years 1989–2012, provided by Statistics Finland. However, to minimize measurement error, we focus on years 1995–2012, for which Statistics Finland utilizes the tax register data of businesses as its primary source of financial statement data. For this period, our data covers the vast majority of Finnish firms across industries, excluding the financial sector. The coverage varies between 95% and 99% of all Finnish firms. Earlier, during the years 1989–1994, Statistics Finland relied on its own survey, which had substantially lower firm coverage.[12] To determine the age of the firms, we also use the Business Register data at the establishment level for years 1989–2012. This data covers all firms that pay value-added tax or have paid employees.

We focus on industries 15–63 according to the NACE Rev. 2 classification. In addition to finance, insurance and real estate, we also exclude agriculture and mining. To concentrate

[12]Statistics Finland has retrospectively added administrative data for 1994. However, coverage in 1994 is not comparable to that in subsequent years.

on firms with meaningful balance-sheet information, and where the owners' labor is not the only source of work input, we restrict the sample to limited liability companies that, on average, employ more than one worker.

When reporting age-related patterns, we focus on firms that can be followed from their first year of operation. Given that our data spans the years 1995–2012, age-related results are reported for firms' first years of operation, up to a maximum of 17 years. For results that are not directly related to firm age, we include all firms that satisfy the other sample criteria, including those for which age cannot be determined and those that do not enter the data at the start of their life cycle.

The variables we use are value added, employment compensation (wages and salaries plus other personnel expenses), total profits, equity, and industry codes at the three-digit level. Unlike some other studies, we use value added as our baseline measure of output rather than revenues. An advantage of this choice is that fixed costs paid to other firms are netted out from the output measure. When we later examine the robustness of our results to wedges that do not directly affect profits, we also use revenues as an alternative measure of output.

In our static model, firms rent capital. To make the empirical measure of profits consistent with this assumption, we deduct the opportunity cost of a firm's own capital, computed as a 5% real interest rate times the firm's total equity, from total profits. We have also performed our accounting exercises using firm-specific implicit interest rates, computed by dividing a firm's interest payments by its borrowed capital. In addition, we consider a specification that uses observed profits directly, without deducting the opportunity cost of own capital, making the results independent of how a firm's own capital is measured. Both of these are reported in Section 2.5. A firm's age is defined as the year in which its first establishment appears in the business register data.

After applying our filtering criteria, we obtain around 740,000 firm-year observations for limited liability companies with at least one worker (see Appendix A for breakdowns of the sample by year and by firm age brackets). The average firm has 19.44 employees and is 7.07 years old. It generates value added of 1.19 million euros, which is substantially less than its mean sales of 4.91 million euros. Mean employment costs are 0.75 million euros, and mean profits are 0.51 million euros. All of these variables exhibit substantial dispersion around their means. Appendix A provides more detailed summary statistics for the variables we use.

We follow the baseline HK setup and attach an industry-specific, time-constant capital elasticity, α_s, to each firm using the labor shares at the 3-digit industry level. Later, in Section 2.5, we also consider firm-time-specific capital elasticities, $\alpha_{t,s,i}$. Finally, we assume a common σ for all firms. Given that our focus is on the variances of $\ln(1-\tau)$ and $\ln(1-\varphi)$, and the covariance between the two components, the exact value of σ does not affect our results regarding the decomposition.[13] We have also conducted our calculations using industry-specific fixed effects to account for variations in demand elasticities across industries. This adjustment does not affect our results, which are reported in the next sections.

2.3 Misallocation in Finland

We utilize the framework introduced earlier to analyze ex post misallocation and its decomposition for Finnish firms. In addition, we examine how these measures vary with firm age. We start by calculating the HK revenue wedge and its components using Proposition 1 for all limited liability companies with more than one employee in industries 15–63 over the period 1995–2012.

To improve reliability, we winsorize the ex post wedge at the 1% level. For observations affected by winsorization, we scale the ex ante wedge and prediction error according to their original contributions to the ex post wedge, ensuring the shares still sum to one. We then compute the variance of the (log) HK revenue wedge, the variance of its components, and their covariance, both for the pooled data and for age-dependent analyses.

Our empirical analysis reveals two key insights:

Finding 1. Uncertainty and ex ante misallocation are equally important in explaining ex post misallocation in pooled data.

Finding 2. Ex post misallocation and uncertainty exhibit strong declining trends over the firm life cycle, while ex ante misallocation remains relatively stable.

We discuss each of these patterns in detail below.

Table 1: The variance of the HK-style revenue wedge and its decomposition into uncertainty and ex ante misallocation.

Variable	Value	Share
$\mathrm{Var}(\ln(1-\tau_{HK}))$	0.206	
$\mathrm{Var}(\ln(1-\varphi))$	0.090	0.44
$\mathrm{Var}(\ln(1-\tau))$	0.090	0.44
$2\mathrm{Cov}(\ln(1-\varphi), \ln(1-\tau))$	0.025	0.12

[13] For the exploration of the association between productivity and decomposition components σ has an impact and there we use a value consistent with our calibration of the quantitative model below.

Table 1 highlights our first finding. It reports results for the pooled data, which includes all firm-year observations regardless of whether firm age is observed. From the table, we see that uncertainty accounts for 44% of the total variation. The variance of the residual wedge, our indirect measure of misallocation, also accounts for 44% of the variation in the HK-style misallocation measure. This suggests that both uncertainty and idiosyncratic revenue distortions play important roles in generating ex post misallocation and TFP losses.

Finally, in the pooled data the covariance between the prediction error and the revenue wedge takes a small positive value. Taken at face value, this suggests that firms with high $1-\tau$ are over-optimistic about their productivity. Alternatively, this pattern may also reflect omitted heterogeneity. Our robustness checks in Section 2.5 support this interpretation. For example, when we control for overhead labor by removing high level managers from the wage bill, the covariance term falls to approximately zero. Moreover, many other alternative specifications also explored in Section 2.5 yield small negative values for the covariance term. In Section 5 we use a model-based approach to examine another possibility, namely that this effect is driven by correlation between productivity and ex ante wedges.

In Appendices B and C, we further decompose the pooled results. Appendix B presents the decomposition across industries. The main observation from this exercise is that the results are remarkably stable across sectors. Uncertainty and ex ante misallocation are both important, each accounting for roughly 40–50% of the total variation. As in Table 1, the covariance term typically explains around 10%. The only notable exception is the electricity, gas, and water supply sector. This sector exhibits a higher level of ex ante misallocation, which is perhaps unsurprising given the tight regulation and limited competition in that industry.

Appendix C reports the decomposition for individual years. Once again, the results closely resemble those from the pooled data. The only notable deviation occurs in the first year of the sample, during which both uncertainty and ex ante misallocation appear elevated.

In this paper, we focus on misallocation measured using the wage bill. This choice is mainly motivated by our desire to minimize the role of measurement error. Our data, which comes from the Tax Administration's business taxation registers, provides high-quality information on value added, profits, and labor costs, all collected for taxation purposes. Using the same accounting framework, it is also possible to construct analogous measures based on capital income. While the register data also allows us to measure capital using balance-sheet information, these measures are more imprecise, as they are based on accounting values rather than market values. Moreover, unlike wages, firm-specific rental rates are not observable.

To assess the robustness of our results to this choice, in Appendix D, we redo our static accounting exercise using value-added-to-capital-income and profits-to-capital-income ratios under two measures of capital: total assets (net of short-term debt) and fixed capital (machinery, equipment, and structures), the latter being the measure used by HK. Both are combined with a common rental rate. The results show that both the level of ex post misallocation and its decomposition are sensitive to the capital measure used. Under both measures, the level of ex post misallocation is substantially higher than in our baseline. Taken at face value, the implied TFP gains from reallocation would fall well outside the range typically reported for developed economies like Finland. Turning to the decomposition, with total assets the level of uncertainty is broadly comparable to our baseline labor-based results, but ex ante misallocation dominates the decomposition. With fixed capital, the decomposition is much more balanced, in line with our baseline findings. The covariance term remains small under both measures. The sensitivity in the level of ex post misallocation primarily reflects the large and measure-dependent dispersion in log capital income, while the shift in the decomposition reflects differences in how profits co-move with each capital measure.

Returning to our baseline labor-based decomposition, we next explore the association between productivity and wedges by estimating the elasticity of the HK wedge, $1 - \tau^{HK}$, with respect to productivity, measured using equation (7). We estimate an OLS regression of $\ln(1 - \tau^{HK})$ on $\ln(z)$. The elasticity of the aggregate misallocation wedge with respect to productivity is typically found to be negative, suggesting that high-productivity firms or establishments face higher implicit "taxes." This is also the case in Finland, where we estimate an elasticity of -0.33. The negative association is stronger than what Hsieh and Klenow (2014) report for the US, but weaker than what they find for India and Mexico.[14]

In addition, we decompose the correlation between productivity and ex post wedges by estimating the elasticities of the prediction error and the ex ante wedge with respect to productivity. These elasticities are -0.16 and -0.17, respectively. This decomposition suggests that nearly half of the observed association between the aggregate wedge and productivity arises from prediction errors.

Finally, we explore life-cycle aspects of misallocation to illustrate the second key finding. We compute the variance of the HK revenue wedge conditional on firm age and redo the decomposition separately for each age group, reporting the variances of prediction errors and residual wedges. The results are shown in Figure 1, with shaded areas around each profile representing block-bootstrapped confidence intervals, where individual firms are used as blocks.

As noted earlier, there is a strong negative trend in ex post misallocation, shown by the red line with circles. The variance is almost twice as large for entrants as for firms older

[14]Hsieh and Klenow (2014) report elasticities between τ and z rather than between $1 - \tau$ and z, so their elasticities take positive values.

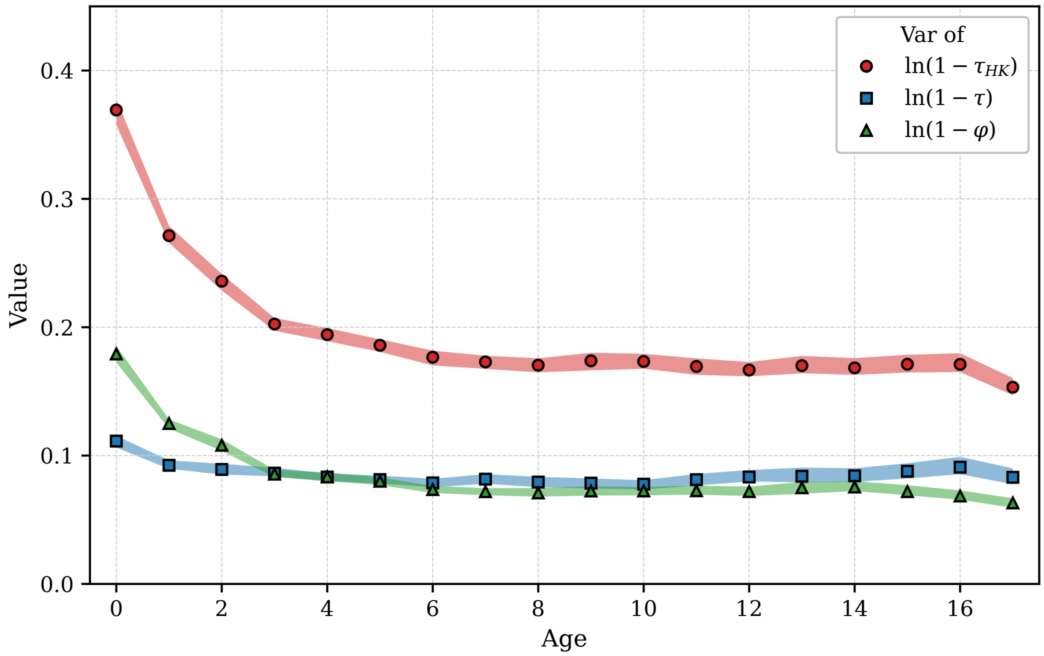

Figure 1: Uncertainty and misallocation for Finnish firms conditional on firm age.

Notes: Shaded areas represent 95% confidence intervals from a block bootstrap with individual firms as blocks.

than ten years. Most of this decline is accounted for by decreasing uncertainty, shown by the green line with triangles, which falls by more than half over the life cycle. In contrast, the dispersion of the residual wedge, shown by the blue line with squares, remains largely flat after a small initial decline during the first two years.

Since young firms also tend to be small, it is useful to investigate whether the reduction in uncertainty is, in fact, driven by increasing production size rather than by aging. This could happen, for example, if productivity and/or demand shocks tended to wash out more for larger firms. We evaluate this channel by first running regressions in which we explain the variables $1-\tau^{HK}$, $1-\tau$, and $1-\varphi$ using firm size. Afterwards, we take the residuals of these regressions and report their variances conditional on age. The results are given in the left panel of Figure 2. Comparing it with Figure 1, we see that the age-dependent trends are practically unaltered; thus firm size has a negligible effect on the results. In Appendix C, we report the life-cycle patterns of misallocation separately for different cohorts. This illustrates that the observed patterns are not driven by a single cohort.

In the right panel of Figure 2, we report the age-related accounting exercise for a "balanced panel," a subset of firms that survive at least until their tenth year. In our sample, these are firms founded between 1995 and 2002. The figure shows that uncertainty and ex

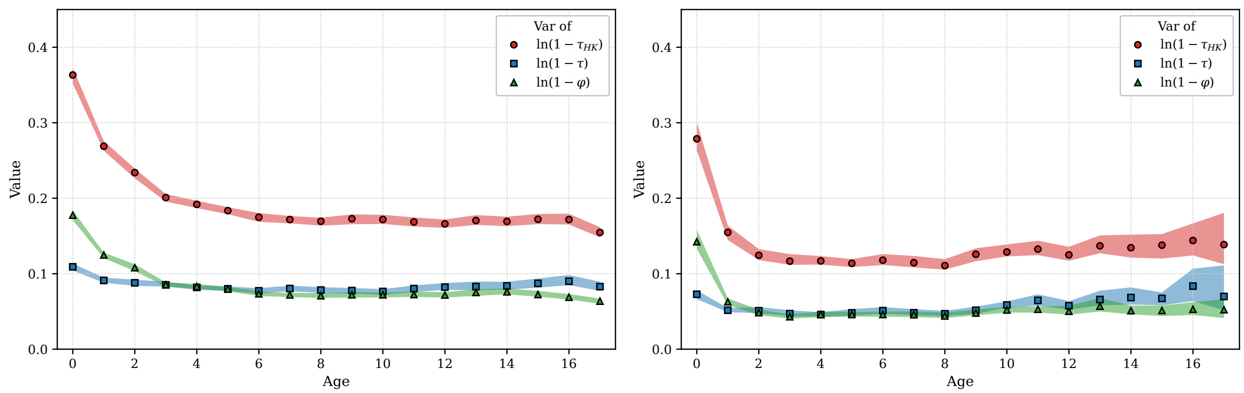

(a) Size controlled

(b) Balanced panel

Figure 2: Uncertainty and misallocation conditional on firm age. The left panel gives the result after controlling for firm size, the right panel shows results for a "balanced panel" that only contains firms surviving at least ten years.

Notes: Shaded areas represent 95% confidence intervals from a block bootstrap with individual firms as blocks.

post misallocation also decline with age for this group of survivors, with the reduction in uncertainty again being substantial. Comparing this figure with Figure 1 reveals that the convergence of uncertainty is faster for the group of survivors.

Motivated by the differences between Figure 1 and the right panel of Figure 2, we explore the association between prediction errors made by firms and their exit in the next period.[15] The results of a linear probability model (LPM) and a logit model based on our whole dataset are assembled in Table 2. The first two columns show the results when we only use the prediction error from the previous year, in addition to year and industry fixed effects, as an explanatory variable. The third and fourth columns include the prediction error from two periods before as an additional regressor.

As we can see, there is a statistically significant positive association between exits and prediction errors. For example, according to the linear probability model in column 1, a one-standard-deviation (0.29) increase in the prediction error leads to an increase of approximately 0.6 percentage points in the exit probability. Given that the unconditional exit probability in our data is 3.1%, this is a substantial increase. The inclusion of the second lag highlights the persistence of the relationship between exits and prediction errors.

The age decomposition hints that reducing uncertainty about fundamentals could play a significant role in understanding the life-cycle patterns of resource allocation. In addition, the decreasing trend in uncertainty, together with a positive association between prediction

[15]Given the administrative nature of our data, firms are highly unlikely to disappear from our data unless they have truly stopped existing as independent firms. Unfortunately, however, we are not able to distinguish between closures, mergers and acquisitions.

errors and exits, is in line with the Jovanovic (1982)-style mechanism of firm growth, in which learning plays an important role in explaining the observed up-or-out life-cycle patterns. In the next sections, we explore an alternative identification strategy and several robustness exercises. We also give further evidence of learning at the firm level by considering age dependency after controlling for firm-specific fixed effects.

Table 2: Exits and prediction errors.

	Dependent variable: Exit at the Beginning of Next Period			
	LPM	Logit	LPM	Logit
$\ln(1-\varphi)$	0.020	0.826	0.014	0.690
	(0.001)	(0.025)	(0.001)	(0.033)
$L(\ln(1-\varphi))$			0.021	0.993
			(0.001)	(0.032)
APE $\ln(1-\varphi)$		0.0236		0.0186
APE $L(\ln(1-\varphi))$				0.0268
Industry FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
N	640644	640564	552181	552019

Notes: For the linear probability models (LPMs), we report robust standard errors. L() refers to lagged value.

2.4 An Alternative Identification Strategy with Materials

Our baseline approach relies on the assumption that ex ante frictions affect profits directly, for example by acting as implicit or explicit taxes. A natural alternative is that these frictions operate primarily through input choices. Moreover, fixed costs not captured by the financial statement data could also affect the baseline profit-based identification. To assess these possibilities, we extend the production environment to include materials and assume that, unlike capital and labor, materials are chosen after firm-level productivity is realized.

The assumption that materials are chosen after uncertainty is resolved allows us to identify the ex ante wedge from the first-order condition with respect to materials. The ratio of the first-order conditions with respect to labor and materials can then be used to pin down the prediction error. As a result, this approach does not rely on profits for identification.

To see this, let us write down the firm's two-stage problem:

$$\max_{k_{t,s,i},\, n_{t,s,i}} \left\{ \mathbb{E}\left[\max_{m_{t,s,i}} (1-\tau_{t,s,i})\, p_{t,s,i} y_{t,s,i} - p^m_{t,s,i} m_{t,s,i} \right] - w_{t,s,i} n_{t,s,i} - R_t k_{t,s,i} \right\},$$

where $p_{t,s,i} = P_t \left(\frac{y_{t,s,i}}{Y_t}\right)^{-1/\sigma}$ and $y_{t,s,i} = z_{t,s,i} k^{\alpha_s}_{t,s,i} n^{\xi_s - \alpha_s}_{t,s,i} m^{1-\xi_s}_{t,s,i}$.

The first-order condition with respect to materials implies

$$(1-\tau_{t,s,i}) \left(1 - \frac{1}{\sigma}\right) (1-\xi_s) p_{t,s,i} y_{t,s,i} = p^m_{t,s,i} m_{t,s,i}, \tag{8}$$

Solving for $1-\tau_{t,s,i}$ yields

$$1 - \tau_{t,s,i} = \frac{\sigma}{\sigma - 1} \frac{1}{(1-\xi_s)\, \frac{p_{t,s,i} y_{t,s,i}}{p^m_{t,s,i} m_{t,s,i}}}. \tag{9}$$

Substituting equation (8) into the firm's problem and taking the first-order condition with respect to labor gives

$$(1-\tau_{t,s,i}) \left(1 - \frac{1}{\sigma}\right) (\xi_s - \alpha_s) \underbrace{\frac{\mathbb{E} z_{t,s,i}^{\frac{1-\frac{1}{\sigma}}{\left(1-\frac{1}{\sigma}\right)(1-\xi_s)}}}{z_{t,s,i}^{\frac{1-\frac{1}{\sigma}}{\left(1-\frac{1}{\sigma}\right)(1-\xi_s)}}}}_{\equiv (1-\varphi_{t,s,i})} p_{t,s,i} y_{t,s,i} = w_{t,s,i} n_{t,s,i},$$

where $1 - \varphi_{t,s,i}$ captures the prediction error induced by the gap between expected and realized productivity. Taking the ratio of the first-order conditions for labor and materials yields

$$1 - \varphi_{t,s,i} = \frac{1-\xi_s}{\xi_s - \alpha_s} \frac{1}{\frac{p^m_{t,s,i} m_{t,s,i}}{w_{t,s,i} n_{t,s,i}}}. \tag{10}$$

Equations (9) and (10) therefore allow us to recover the ex ante wedge and the prediction error using revenue, labor costs and material costs.

To implement this approach in the data, we use a measure of firms' material costs in addition to the variables used previously. Our measure of these costs includes all intermediate inputs, excluding external services. Moreover, instead of using value added as the measure of firm output, $p_{t,s,i} y_{t,s,i}$, we now measure output as value added plus material costs. In line with the baseline calibration, we set the capital elasticity, α_s, and the material elasticity, $1 - \xi_s$, using labor and material shares at the 3-digit industry level.

Table 3: Decomposition of ex post misallocation under alternative identification using materials.

Variable	Value	Share
$\text{Var}(\ln(1-\tau_{HK}))$	0.320	
$\text{Var}(\ln(1-\varphi))$	0.737	2.30
$\text{Var}(\ln(1-\tau))$	0.183	0.57
$2\text{Cov}(\ln(1-\varphi), \ln(1-\tau))$	-0.601	-1.88

The results for the alternative approach are reported in Table 3. Adding materials increases the variation in the output measure, leading to an increase in measured ex post misallocation. In line with this, ex ante misallocation is also scaled up by a similar proportion. There is, however, a substantially larger increase in the uncertainty component, which now exceeds ex post misallocation. This is made possible by a large negative covariance term.

Figure 3 reproduces the decomposition of ex post misallocation conditional on firm age. As was the case with the baseline results, the uncertainty component exhibits a strong age-dependent trend, while ex ante misallocation remains relatively stable over the life cycle. Ex post misallocation and uncertainty track each other somewhat less closely than in the baseline results, reflecting the larger role of the covariance term in this specification.

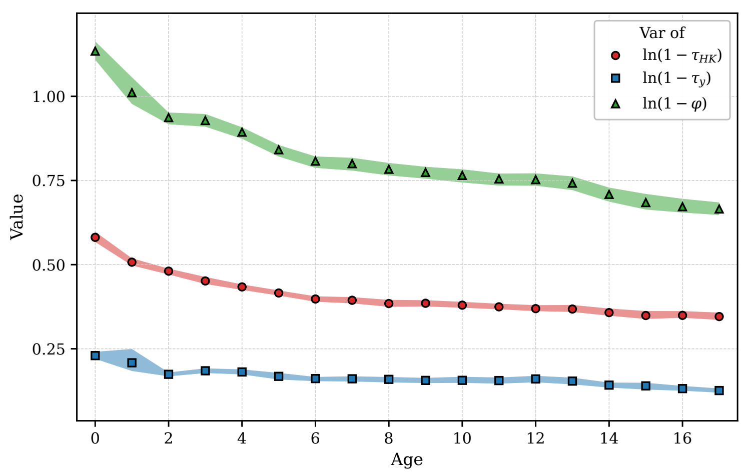

Figure 3: Age profiles under alternative identification using materials.

Notes: Shaded areas represent 95% confidence intervals from a block bootstrap with individual firms as blocks.

2.5 Robustness

Our baseline indirect approach suggests that uncertainty plays a substantial role in generating ex post revenue misallocation and in explaining its age-dependent decline. The alternative identification strategy with materials appears to support this view.

In this section, we explore the robustness of these results to different forces that might generate variation in the observed profits-to-wage-bill and value-added-to-wage-bill ratios. Unless otherwise noted, the exercises below build on our baseline profit-based approach. We start by allowing for production heterogeneity across firms within industries and across time. We then relax the assumption of homogeneous interest rates. After that, we consider a specification that is agnostic about the role of capital in production. Next, we exploit the panel dimension of our data by introducing firm-specific fixed effects. This allows us to control for differences in unobserved firm heterogeneity. We also use the firm fixed-effects setup to provide some indicative evidence that forecasting precision improves within firms as they age. Finally, we explore the potential role of overhead labor by excluding upper-level managers from firms' measured labor input.

The results of these exercises are reported in Table 4 and Figure 4. At the end of this section, we briefly discuss results from an extended setup that combines our baseline profit-based approach with the alternative approach using materials to allow for labor distortions in addition to output distortions (see Appendix E for details).

In the previous sections, we used industry-specific capital elasticities. However, as David and Venkateswaran (2019) have highlighted, heterogeneity in production technologies, in the form of varying capital intensities, can be important for observed misallocation. To take this into account, we consider firm-time-specific capital elasticities, $\alpha_{t,s,i}$, identified from capital-labor ratios at time t at the firm level. The first-order conditions of our firm's problem imply that

$$\frac{1-\alpha_{t,s,i}}{\alpha_{t,s,i}} = \frac{w_{t,s,i}n_{t,s,i}}{R_t k_{t,s,i}}.$$

Note that in the presence of relative distortions that affect capital and labor asymmetrically, this approach exaggerates the variation in the capital elasticity. However, together with the alternative extreme of a common industry-specific α_s, this gives us information about the sensitivity of our results to production heterogeneity.

The results with production heterogeneity are shown in row b of Table 4. Row a of the table reproduces the baseline results from the previous section. A comparison of these rows reveals that allowing for production heterogeneity reduces ex post misallocation by 26%, from 0.206 to 0.153. Next, looking at the decomposition, we see that the covariance term is now negative and close to zero when heterogeneity is allowed. In line with this, the relative

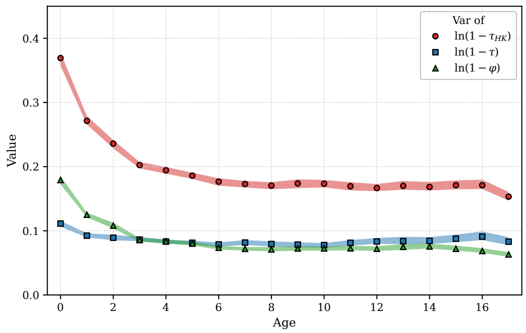

(a) Baseline

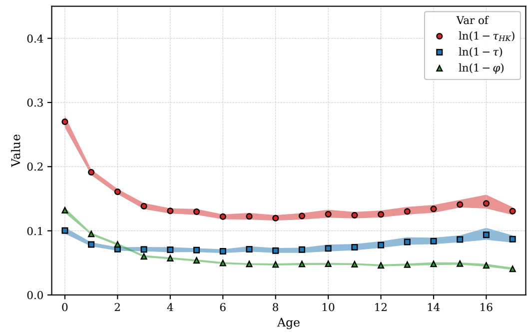

(b) Firm-specific α

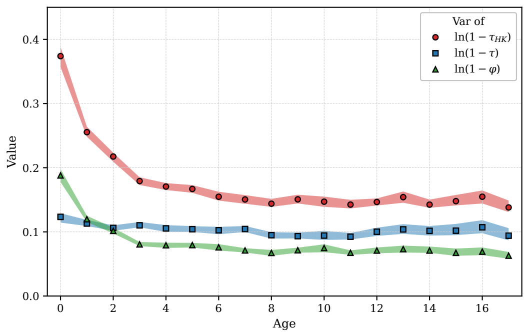

(c) Het. interest rate

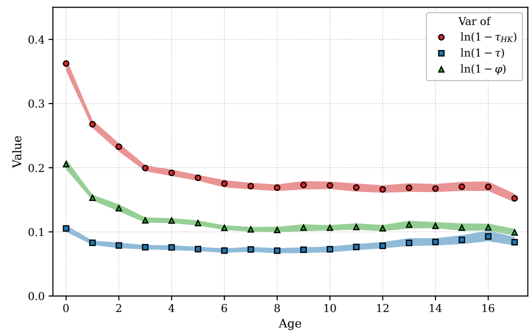

(d) Agnostic capital

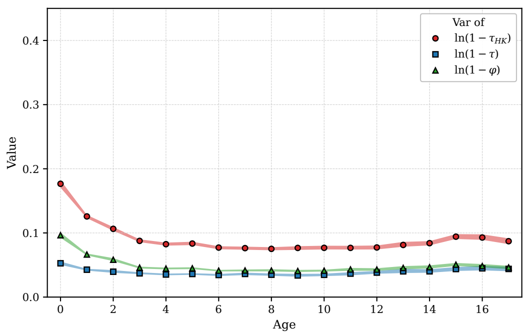

(e) Fixed effects

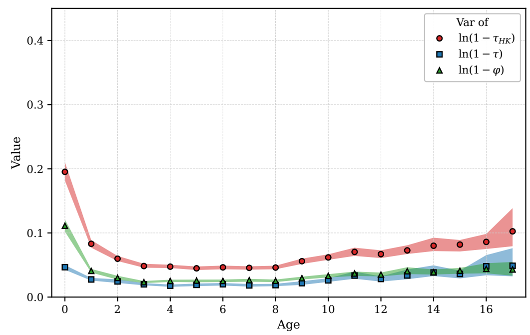

(f) Fixed effects (Balanced panel)

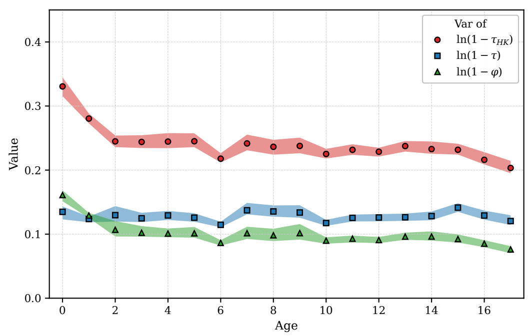

(g) Blue Collars

Figure 4: Age patterns for alternative specifications.

Notes: Shaded areas represent 95% confidence intervals from a block bootstrap with individual firms as blocks.

Table 4: Sensitivity of static misallocation and its components.

	$\text{Var}(\ln(1-\tau_{HK}))$	$\text{Var}(\ln(1-\tau))$	$\text{Var}(\ln(1-\varphi))$	$2\text{Cov}(\ln(1-\tau), \ln(1-\varphi))$
a) Baseline	0.206	0.090 [44%]	0.090 [44%]	0.025 [12%]
b) Firm-specific α	0.153	0.084 [55%]	0.062 [41%]	0.007 [5%]
c) Het. interest rate	0.187	0.109 [58%]	0.088 [47%]	-0.010 [-5%]
d) Agnostic capital	0.204	0.086 [42%]	0.124 [61%]	-0.006 [-3%]
e) Fixed effects	0.097	0.042 [43%]	0.052 [54%]	0.003 [3%]
f) Blue collar	0.244	0.131 [54%]	0.104 [43%]	0.008 [3%]

Notes: Relative contributions are given in brackets.

contributions of uncertainty and ex ante misallocation (reported in brackets) are elevated.

In panel a of Figure 4, we again reproduce the baseline life-cycle aspects of misallocation and its components, while in panel b we use firm-time-specific elasticities. Comparing the two figures reveals that even though the starting values are lower with production heterogeneity, the general patterns are similar: ex post misallocation is halved during the first ten years, while the uncertainty component decreases even more in relative terms. The life-cycle profile of ex ante misallocation is relatively similar to that in the baseline results, though there seems to be a slight increase in ex ante misallocation after the first 10 years.

A branch of recent literature has highlighted the substantial variation in interest rate spreads and its implications for aggregate outcomes (see e.g., Bai et al., 2018, Cavalcanti et al., 2021 or Gilchrist et al., 2013). Even though we focus on symmetric wedges, this variation could still be relevant for our decomposition since it could alter the observed ratios of monopoly profits to the wage bill. Up to this point, when generating our profit measure, we have assumed that the rental rates of capital are equal across firms. However, heterogeneous default risk, for instance, could mean that the required return on a firm's own capital also varies across firms. We explore the importance of this channel by using firm-time-specific implicit interest rates as a proxy for the required return on own capital. That is, we take the profit measure from financial statement data and subtract the firm's own capital multiplied by a firm-time-specific interest rate from it. When calculating these interest rates, we divide interest payments by debt.

Borrowing costs are also heterogeneous among Finnish firms; the standard deviation of implicit interest rates is 2.5% (with the mean rate being 2.1%). However, when we recalculate our decomposition using the profit measure that takes this into account, our decomposition results are only slightly altered.[16] From row c of Table 4, we see that the role of ex ante wedges is somewhat larger, as ex ante misallocation now accounts for 58% of the total variation. The relative contribution of uncertainty is 47%. In line with the previous robustness exercise, the covariance term is now negative and close to zero. Panel c of Figure 4 reports the evolution of misallocation and its components conditional on firms' age with heterogeneous interest rates. The profile of ex ante misallocation is in line with the previous results. As to uncertainty, the starting values are somewhat higher than in the benchmark case, while the reduction in the first years is a bit faster.

In our baseline static framework, both labor and capital were assumed to be variable inputs that could be freely chosen in each period. However, the structural model in Section 3 features labor as the only variable input of production, corresponding to an economy where $\alpha_s = 0$. To explore whether the treatment of capital in the accounting framework matters for our results, and to examine how constraints on capital adjustment, such as adjustment costs (e.g., Asker et al., 2014) or credit rationing (e.g., Buera et al., 2011), might affect them, we proceed by being agnostic about how the level of capital is determined and only assume that labor is a variable input of production.

Conditional on the level of capital, the first-order condition with respect to labor is still given by equation (4). Defining a new measure of profits as monopoly profits plus rental rates on capital, $\pi^*_{t,s,i} \equiv \pi_{t,s,i} + R_t k_{t,s,i}$, we have

$$\pi^*_{t,s,i} = (1 - \tau_{t,s,i}) p_{t,s,i} y_{t,s,i} - w_{t,s,i} n_{t,s,i}$$
$$\pi^*_{t,s,i} = (1 - \tau_{t,s,i}) p_{t,s,i} y_{t,s,i} \left(1 - \frac{\sigma - 1}{\sigma} (1 - \alpha_s)(1 - \varphi_{t,s,i})\right).$$

Dividing this expression by wage bill, $w_{t,s,i} n_{t,s,i}$, and solving for $1 - \varphi_{t,s,i}$ gives

$$1 - \varphi_{t,s,i} = \frac{\sigma}{\sigma - 1} \frac{1}{1 - \alpha_s} \frac{1}{1 + \frac{\pi^*_{t,s,i}}{w_{t,s,i} n_{t,s,i}}}.$$

We can again insert this expression into part (ii) of Proposition 1 to solve for the ex ante

[16] With identical samples, this specification would yield the same ex post misallocation as the baseline, since the value-added-to-wage-bill ratios are unchanged. The differences observed reflect variations in sample composition, as this variant requires additional variables to be observable for a given firm-year observation.

wedge.

Unlike in our baseline specification, where capital intensity cannot be absorbed by industry fixed effects, the uncertainty measure here depends on capital intensity only in a multiplicative fashion. Thus, after log transformations, industry differences in capital use can be controlled for with fixed effects, allowing us to recover ex post misallocation and its composition without assumptions about how capital is used or about industry differences in capital intensity. Moreover, since π^* corresponds directly to the observed profit measure in the financial statement data without the need to deduct the opportunity cost of own capital, the results are also independent of how a firm's own capital is measured.

Row d of Table 4 assembles the results of our indirect approach when we do not take a stance on how the level of capital is chosen and use observed profits directly. Given that the value-added-to-wage-bill ratios are not altered, the measure of ex post misallocation is practically unchanged.[17] The most notable difference compared to the previous results is the slightly elevated level of uncertainty. The level of ex ante misallocation, on the other hand, is unaltered. The covariance term again takes a small negative value. Panel d of Figure 4 gives the life-cycle profiles for this case. From it, we see that uncertainty goes up for all ages and closely tracks the profile of ex post misallocation. The age pattern of ex ante misallocation again resembles that of the baseline results.

We also utilize the panel dimension of our data to control for unobserved firm heterogeneity. To do this, we take the (log) wedges generated by our baseline approach and run regressions where we use these wedges as dependent variables and firm fixed effects as explanatory variables. We take the residuals of these regressions, given that a firm has at least two observations, and report the variation in the residuals. This approach allows us to control for, for example, heterogeneity in markups or fixed costs. On the downside, we are also likely to throw away relevant variation in the data. For uncertainty, this happens, for example, if, due to learning, firms' prediction errors are persistent and firms that are over-optimistic about their productivity exit after observing a few bad signals in a row. Thus, it seems likely that the fixed-effect setup gives downward-biased results for ex post misallocation and its components.

The results of the fixed effect setup are shown in row e of Table 4. As expected, compared to the benchmark results, the variation in ex post wedge is now reduced. The same is true for all of its components. Interestingly, the drop is smaller for uncertainty than for ex ante misallocation. Despite these somewhat uneven reductions, the relative importance of uncertainty and ex ante misallocation is still in line with the previous results. Contrary to

[17] As with the heterogeneous interest rate specification, the small difference in ex post misallocation relative to the baseline reflects variations in sample composition.

the previous robustness exercises, the covariance term takes a small positive value. Panel e of Figure 4 illustrates the age-related patterns with fixed effects. The broad patterns in ex post misallocation and uncertainty are similar to the baseline results though with reduced level effects.

To illustrate the role of within-firm learning versus selection, we repeat the fixed effect exercise (panel f of Figure 4) for a balanced panel, i.e., for firms that survive at least until their tenth year. From this we see that, even after controlling for unobserved firm heterogeneity, uncertainty is declining for surviving firms in the early parts of their life cycle. There is also a slight upward trend in ex post misallocation and its components for these firms in the latter years, though widening confidence bands and low levels make conclusions somewhat uncertain.

In row f of Table 4 and in panel g of Figure 4 we explore the sensitivity of our results to overhead labor. We use the socio-economic group variable from the FOLK data, again provided by Statistics Finland. At the firm-year level, we calculate the ratio of wages paid to upper-level managers to total wages paid to workers for whom the socio-economic group is observed. We then adjust the labor cost variable downward in line with this ratio and redo our baseline decomposition. The change in how we measure labor input increases the observed ex post misallocation somewhat, to 0.244. In line with this, both uncertainty and ex ante misallocation also increase. For ex ante misallocation, the increase also translates into a higher relative contribution. Unlike the other two components, the contribution of the covariance term declines both in absolute and relative terms when compared against the baseline. The socio-economic group variable is available for 1995 and 2000, and annually from 2004 onward. This limited coverage generates more volatility in the life-cycle profiles. Otherwise, the observed patterns are similar to the other specifications.

Finally, in Appendix E we extend our approach to include labor distortions. To do this, we consider a more general setting, where we combine our profit- and material-based approaches. As in Section 2.4, the first-order condition for materials allows us to identify the output wedge. In line with our baseline approach, we can then use the information embedded in profits to identify the prediction error. The labor input wedge is then identified from the first-order condition of labor. When the prediction error does not correlate with the output wedge, this approach potentially underplays the role of uncertainty and overplays the role of ex ante wedges if firms also face some uncertainty regarding their use of materials.

The results of the extended approach (see Table 17 and Figure 8 in Appendix E) indicate, similarly to the previous exercise with materials, an elevated level of uncertainty and a strongly negative covariance term. However, contrary to Section 2.4, there is also a substantial upward shift in ex ante misallocation, measured by the variance of the combined wedge. Given that ex ante misallocation now reflects not just the role of output wedges, but also the contribution of labor wedges, this result could be seen as a sign of the importance of labor distortions in Finland. For example, highly centralized wage setting with strong union coverage, together with rigid labor protection practices, could generate substantial misallocation. As for the age patterns, uncertainty still has a strong downward trend, while ex ante misallocation stays relatively stable.

3 Model

Motivated by the results from our static decomposition of misallocation, we set up a general equilibrium model in which firms have to choose their inputs without full information about the current-period productivity. To allow for the observed age-dependent trend in the prediction errors, we add Jovanovic (1982)-style learning mechanism to a general equilibrium framework, similar to Hopenhayn and Rogerson (1993) and Melitz (2003).

3.1 Households

There is a mass $\bar{N}$ of risk-neutral, infinitely lived households that derive utility from consumption and supply labor inelastically. The behavior of households can be summarized with a representative household whose preferences are given by $\sum_{t=0}^{\infty} \beta^t C_t$, where C_t is a consumption basket of individual goods aggregated with the CES aggregator such that

$$C_t = \left(\int_{\Omega_t} c_{i,t}^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \tag{11}$$

where Ω_t is the amount of goods available. The household owns the firms and thus the budget constraint is given by

$$\int_{\Omega_t} p_{i,t} c_{i,t} di = w_t \bar{N} + \Pi_t, \tag{12}$$

where Π_t are aggregate profits and w_t is the wage rate. We focus on stationary equilibrium and, thus, from now on, we drop the time indexes. To ease the notation, we also drop the firm index i. Moreover, we use labor as the numéraire.

3.2 Incumbent

There is an endogenous measure of incumbent firms denoted by Ω. Each firm produces a unique good and faces a demand in line with (11) and (12),

$$y = \left(\frac{p}{P}\right)^{-\sigma} Y, \tag{13}$$

where $P = \left(\int_\Omega p^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ is the ideal price index associated with the CES aggregator.

The production function of a firm is given by linear technology

$$y = e^z n, \tag{14}$$

where e^z is the firm's TFP and n the number of employees hired by the firm. The firm-specific TFP in the current period is given in logs as

$$z = z_p + z_{tr}. \tag{15}$$

In this equation, z_p is a persistent productivity component that follows an AR(1) process for each firm :

$$z_p = \rho z_{p,-} + \epsilon_p, \tag{16}$$

where $0 < \rho < 1$ and $z_{p,-}$ is the value of the persistent component in the previous period. The innovation term, ϵ_p, is distributed according to $\epsilon_p \sim N(0, \sigma^2_{\epsilon_p})$. For a new firm, the initial persistent component is drawn from a normal distribution with mean zero and variance $\frac{\sigma^2_{\epsilon_p}}{1-\rho^2}$. Finally, the other component in equation (15), z_{tr}, is temporary productivity, which is drawn from a normal distribution in each period. We assume that $z_{tr} \sim N(0, \sigma^2_{z_{tr}})$.

A firm observes z but is unable to decompose it. In line with the analysis in the previous section, we assume that the firm needs to choose its employment before it observes z. This implies that in each period t, the firm needs to form an estimate of the persistent productivity component conditional on the history of observed values of z up to $t-1$. We denote this prediction by m. As more information is accumulated, m becomes more precise.

The firm uses Bayesian learning to update its expectation. Given the log-normality, we get the standard Kalman filter with the following recursive representation for the prediction,

m, and the variance of the firm's expectation of its permanent productivity, Σ:

$$m' = \rho m + K(z - m) \tag{17}$$

$$K = \frac{\rho \Sigma}{\Sigma + \sigma^2_{z_{tr}}} \tag{18}$$

$$\Sigma' = \frac{\rho^2 \Sigma \sigma^2_{z_{tr}}}{\Sigma + \sigma^2_{z_{tr}}} + \sigma^2_{\epsilon_p} \tag{19}$$

(See, e.g., Ljungqvist and Sargent, 2018, for details). We assume that all firms start with common priors, i.e., with the unconditional persistent productivity distribution:

$$m_0 = 0, \quad \Sigma_0 = \frac{\sigma^2_{\epsilon_p}}{1 - \rho^2}.$$

Since all firms share this initial condition and Σ evolves deterministically, Σ is a function of firm age alone, and we use age a as a sufficient statistic for Σ in the firm's recursive problem below.

Given this structure, the firm's expected value for m' before observing productivity z is distributed as

$$E(m' \mid m, \Sigma) \sim N(\rho m, \rho \Sigma K).$$

In addition, the distribution of z conditional on productivities up to $t - 1$ is given as

$$z = m + (z_p - m) + z_{tr} \sim N(m, \Sigma + \sigma^2_{z_{tr}}). \tag{20}$$

Following, e.g., Restuccia and Rogerson (2008), we do not explicitly model the sources of misallocation, but assume that these frictions can be summarized by an idiosyncratic distortion, $1 - \tau$, that appears in the firm's profit maximization problem as a revenue tax would. For simplicity, we assume that the firm-specific distortion only contains a permanent component. That is,

$$\ln(1 - \tau) = \tau_p, \tag{21}$$

where τ_p is normally distributed and drawn upon entry. In addition, operating firms also have to pay periodic fixed costs, c_f, paid in labor.

Given the demand structure, the productivity process and the revenue wedge, a firm's objective is to maximize its lifetime profits by making an optimal exit/stay decision and, conditional on staying, by choosing current-period employment. As stated earlier, we assume that this decision is made before the current-period productivity is known. Thus, the firms' marginal productivities will differ due to the idiosyncratic revenue wedge and expectation error. In addition to these channels, we also allow for convex adjustment costs that could also be a potential explanation for the age-dependent trend in misallocation.

The intra-period timing is summarized in Figure 5. At the beginning of a period, an incumbent firm chooses whether it wants to exit or not. A firm that exits avoids the fixed cost and only pays the adjustment cost of scaling employment to zero. A firm that decides to continue pays periodic fixed costs, c_f . Next, it chooses its employment. If a continuing firm decides to adjust its scale, the firm has to pay an adjustment cost, $\lambda(\frac{n-n_-}{\bar{n}})^2\bar{n}$, where $\bar{n} = \frac{n+n_-}{2}$. Unlike fixed costs (and the entry cost introduced in the next subsection), adjustment costs do not consume labor; they are a resource cost measured in numéraire units. This drives a wedge between aggregate output and consumption. After choosing its employment level, the firm produces and observes its current-period (combined) productivity. Finally, at the end of the period, the firm may be forced to exit due to an exogenous shock that happens with probability γ.

Taken together, the firm's problem at the beginning of a period can be summarized by the following recursive expression:

$$V(m, a, n_-, \tau_p) = \max\{W(m, a, n_-, \tau_p), -2\lambda n_-\}, \tag{22}$$

where the relevant state variables are the firm's belief about its permanent productivity, m; the firm's age, a; its employment in the previous period, n_-; and the revenue distortion, τ_p. The value of staying, $W(\cdot)$, is given by

$$\begin{aligned} W(m, a, n_-, \tau_p) = \Bigg[& \max_n Y^{\frac{1}{\sigma}} P e^{\tau_p} E_z[(e^z)^{\frac{\sigma-1}{\sigma}}] n^{\frac{\sigma-1}{\sigma}} - n - c_f - \lambda(\frac{n-n_-}{\bar{n}})^2 \bar{n} \\ & + \beta(1-\gamma) \int V(m', a+1, n, \tau_p) dF(m'|\rho m, \rho\Sigma K) \Bigg]. \end{aligned} \tag{23}$$

The expectation $E_z[\cdot]$ appears inside the revenue expression because employment n is chosen before the current period's productivity z is realized. The solution to this Bellman equation gives an exit policy $x(m, a, n_-, \tau_p)$, which takes value 1 if the firm chooses to exit and 0 if the firm chooses to continue, and an employment policy $n(m, a, n_-, \tau_p)$.

3.3 Entry

There is a continuum of potential entrants. Each of them has to pay an entry cost, c_e, paid in labor, if they want to start operating. After paying the entry cost, a new firm starts as an incumbent firm in the next period with age zero and belief $m = 0$. Its initial persistent productivity z_p is drawn from the unconditional distribution $N(0, \sigma_{\epsilon_p}^2/(1-\rho^2))$. The entrant

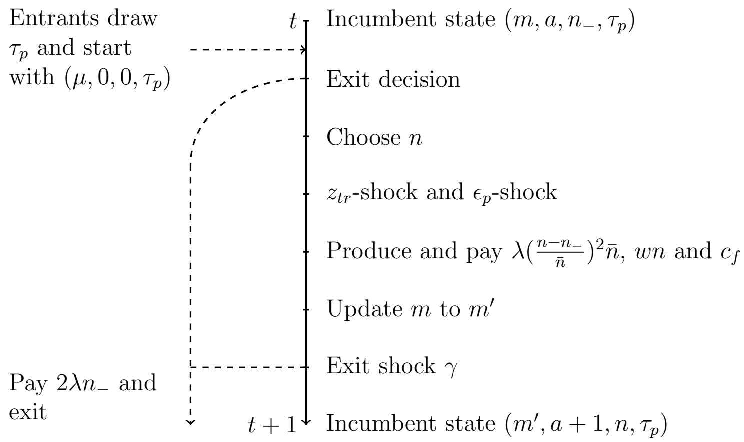

Figure 5: Intra-period timing.

learns its permanent revenue wedge, τ_p, only after entry. The amount of entrants is such that the expected gains are equal to the entry cost:

$$c_e = \beta \int V(0, 0, 0, \tau_p) H(d\tau_p), \tag{24}$$

where $H(\cdot)$ is the distribution of permanent revenue wedges across potential entrants.

3.4 Stationary Equilibrium

Using firms' exit and hiring policies, $x(\cdot)$ and $n(\cdot)$, we can define the evolution of the firm distribution measured at the beginning of each period, $\Psi(dm, a, dn_-, d\tau_p)$. At the stationary equilibrium, $\Psi(dm, a, dn_-, d\tau_p)$ is given by

$$\begin{aligned}\Psi(\mathcal{M}', A', \mathcal{N}, T_p) = \sum_{a|a+1\in A'} \int_{(m,n_-,\tau_p)|n(\cdot)\in\mathcal{N},\tau_p\in T_p)} \\ (1-\gamma)Q_a(m, \mathcal{M}')(1 - x(m, a, n_-, \tau_p))\Psi(dm, a, dn_-, d\tau_p) \\ + M\mathbb{I}(0 \in A')\mathbb{I}(0 \in \mathcal{M}')\mathbb{I}(0 \in \mathcal{N}) \int_{T_p} H(d\tau_p),\end{aligned} \tag{25}$$

where $Q_a(m, \mathcal{M}')$ is the age-dependent transition function for beliefs, each (m, a, n_-, τ_p) is such that $n(m, a, n_-, \tau_p) \in \mathcal{N}$ and $\tau_p \in T_p$. Moreover, M denotes the measure of entrants; $\mathbb{I}(a + 1 = 0)$ is an indicator function taking the value of one if $0 \in A'$; $\mathbb{I}(m' = 0)$ is getting the value of one if $0 \in \mathcal{M}'$; and $\mathbb{I}(n = 0)$ takes the value of one if $0 \in \mathcal{N}$.

Given the measure of firms, we can express the labor demand as

$$N = \sum_a \int [n(m, a, n_-, \tau_p) + c_f]\Psi(dm, a, dn_-, d\tau_p) + c_e M. \tag{26}$$

The stationary equilibrium can be defined with policy functions, a price index, aggregate output, a stationary distribution of firms and a mass of entrants such that

1. the policy rules $x(\cdot)$ and $n(\cdot)$ solve the firm's problem given by (22) and (23),
2. the price level and aggregate output are such that the free entry condition holds,
3. the stationary measure of firms is given by (25)
4. the mass of new entrants is such that the labor market clears, i.e., N, given by (26), is equal to fixed labor supply $\bar{N}$.

Aggregate output is defined as $Y = \left(\int_\Omega y^{(\sigma-1)/\sigma}\right)^{\sigma/(\sigma-1)}$. Using this, we report two aggregate measures in the quantitative analysis: labor productivity, $Y/\bar{N}$, and aggregate TFP, Y/N_{prod}. Here $N_{prod} = \sum_a \int n(m, a, n_-, \tau_p)\Psi(dm, a, dn_-, d\tau_p)$ is production labor, which excludes labor absorbed by the fixed costs and the entry cost.

4 Quantitative Analysis

In this section, we match our model with the Finnish firm-level data. We then redo our static calculations using simulated data and explore the importance of different frictions for the determination of the aggregate output and TFP.

4.1 Calibration Strategy

The parameters governing the preferences of the representative household are calibrated directly. We use the simulated method of moments to fix the rest of the parameters (directly) related to the firms' problem by setting our model to match the prime observable features of the Finnish firm-level data: life-cycle growth patterns, the firm-size distribution and the selection patterns. We calculate our targets from the same data set we used in Section 2.

In line with our data, the model's period is set to one year. We assume a 5% real interest rate and thus fix β at 0.95. We follow HK and set the elasticity of substitution, σ, at 3. In addition, as stated earlier, we have normalized the productivity process by setting its unconditional mean to zero. We also normalize the aggregate labor supply to $\bar{N} = 1$.

After this, we are left with eight parameters to calibrate internally: the persistence and innovation variance of the productivity process, ρ and $\sigma^2_{\epsilon_p}$; the variance of transitory productivity, $\sigma^2_{z_{tr}}$; the variance of the revenue wedge, $\sigma^2_{\tau_p}$; the entry cost, c_e; the fixed costs, c_f; the adjustment cost parameter, λ; and the exogenous exit rate, γ. We fix these parameters using the simulated method of moments, minimizing the sum of squared relative distances between moments computed from the Finnish firm-level data and their model-generated counterparts, with the identity weighting matrix.

We target 14 moments calculated from the Finnish firm-level data using the same sample restrictions as in Section 2. Our first set of targets summarizes the growth profiles of young firms: the mean and variance of employment growth rates for surviving firms in two age brackets (ages 0–2 and 3–5). The mean growth rates and growth-rate variances for older firms (ages 6–10 and 11–15) are used for validation. Next, we target exit rates across the age brackets (ages 0–2, 3–5, 6–10 and 11–15). In addition, we target the first- and second-order autocorrelations of employment. Finally, we target the size distribution of firms using five employment brackets ($0 \leq n < 5$, $5 \leq n < 10$, $10 \leq n < 20$, $20 \leq n < 50$ and $50 \leq n$), giving four free moments. Due to a complicated equilibrium setup, the parameter values are defined jointly. However, next, we give a heuristic argument about which statistics are most relevant for which parameters.

Let us start by considering parameters relating to the productivity process. A noisier signal,i.e., higher $\sigma^2_{z_{tr}}$, amplifies uncertainty and the option value of waiting, making firms more reluctant to exit early in their life cycle. This reduces the gap between the exit rates of different age brackets. The variance of innovations in the persistent technology component, $\sigma^2_{\epsilon_p}$, increases the weight that firms give to new information, as well as the prevalence of big innovations in the observed process, $z - m$. These effects boost the variance of growth rates for all firms. The persistence of the AR(1) process, ρ, also increases the variation of observed innovations, and thus the variances of growth rates increase for all age groups. However, unlike $\sigma^2_{\epsilon_p}$, ρ also has a substantial effect on the mean growth rate of new firms and the variance of the employment distribution. The first- and second-order autocorrelations of employment provide further discipline on the productivity process parameters: since employment is driven largely by productivity beliefs, its autocorrelation structure carries much of the same identifying content as the variance and autocovariances of productivity would (a logic we exploit more directly in Section 5).

Increasing the entry cost c_e reduces entry and thus competition, boosting prices. This increases the optimal size of all firms, which in turn fosters the growth rate of young firms. The shape of the employment distribution is sensitive to fixed costs c_f. Increasing fixed costs raises endogenous exits and tilts the firm distribution towards large firms. The exit rates of older firm are primarily informative about the exogenous exit rate γ, since most low-productivity selection has already occurred by that stage of the life cycle. A higher variation in revenue distortions, $\sigma^2_{\tau_p}$, increases the first two moments of the employment distribution. Finally, adjustment costs λ mainly affect the autocovariance of employment beyond what the productivity process alone would generate.

Before moving on to the fit of the model, it is worth highlighting that our targets do not include ex post misallocation nor its components. Our goal is to see whether a learning model that is set to reproduce the basic life-cycle facts of firms' growth is naturally able to generate uncertainty and misallocation patterns in line with our accounting exercise.

4.2 Fit of the Model

Table 5 reports the data targets and the model counterparts, with shaded rows marking validation moments that were not used as calibration targets. The associated parameter values are given in Table 6.

Overall, the model fits the Finnish data quite well, especially taking into account the over-identification. Mean growth rates decline from around 13–16% in ages 0–2 to near zero by ages 3–5 in both the data and the model. The variance of growth rates for ages 0–2 is also closely matched, while the variance for ages 3–5 is somewhat lower in the model than in the data. The exit-rate profiles for all four age brackets are matched precisely. Regarding the other targets, the model captures the autocorrelations of employment and the firm-size distribution well.

We use the mean growth rates for older surviving firms (ages 6–10 and 11–15) as well as the variances of these growth rates for a validation exercise. Both in the model and in the data, average employment growth rates of older firms are close to zero. In addition, the model generates dispersion in these growth rates similar to the observed dispersion.

In line with the different growth patterns for young and old firms present in the data and in the model, our calibration implies substantial learning over the life cycle. One way to summarize this is to look at the ratio $\frac{\Sigma_\infty+\sigma^2_{z_{tr}}}{\Sigma_0+\sigma^2_{z_{tr}}}$, i.e., the variance of the forecast error in productivity for old firms with converged uncertainty relative to the entrants' variance. With our parametrization, this ratio is 57%.

As a second validation exercise, we explore how average employment evolves over the life cycle once we account for selection through exits, not just survivor growth. To be more precise, we calculate average log employment growth from age a to $a+1$ and decompose it as

$$\underbrace{(\overline{\log n}^S_{a+1} - \overline{\log n}_a)}_{\text{Total}} = \underbrace{(\overline{\log n}^S_{a+1} - \overline{\log n}^S_a)}_{\text{Survivor growth}} + \underbrace{(\overline{\log n}^S_a - \overline{\log n}_a)}_{\text{Selection}},$$

Table 5: Targets and model counterparts.

Moment	Data	Model
By Age Group		
Mean growth rate, age 0–2	0.13	0.16
Mean growth rate, age 3–5	0.02	0.01
Mean growth rate, age 6–10	0.00	0.00
Mean growth rate, age 11–15	−0.01	−0.01
Variance of growth rate, age 0–2	0.16	0.17
Variance of growth rate, age 3–5	0.10	0.06
Variance of growth rate, age 6–10	0.07	0.05
Variance of growth rate, age 11–15	0.07	0.05
Exit rate, age 0–2	0.05	0.04
Exit rate, age 3–5	0.03	0.03
Exit rate, age 6–10	0.02	0.03
Exit rate, age 11–15	0.02	0.02
Autocorrelation		
First-order autocorrelation of log employment	0.96	0.97
Second-order autocorrelation of log employment	0.94	0.93
By Employees Group		
Share of firms with $1 \leq$ employees < 5	0.57	0.57
Share of firms with $5 \leq$ employees < 10	0.20	0.21
Share of firms with $10 \leq$ employees < 20	0.11	0.11
Share of firms with $20 \leq$ employees < 50	0.07	0.07
Share of firms with employees ≥ 50	0.05	0.03

Notes: Shaded rows indicate moments that are not targeted in the calibration. Growth rates are computed as log differences of employment for surviving firms. Sample restrictions follow Section 2.

where $\overline{\log n_a}$ is the average log employment of all firms at age a, $\overline{\log n_a}^S$ is the average log employment of those that survive to age $a+1$, and $\overline{\log n_{a+1}}^S$ is the average log employment of those same survivors at age $a+1$. The survivor component corresponds to our calibration's growth-rate moment (average log difference in employment for survivors), while the selection component reflects the contribution of exits to measured employment growth.

Table 6: Parameter values.

Description	Parameter	Value
Entry cost	c_e	12.66
Fixed cost	c_f	0.79
Std. of innovations to AR(1)	σ_{ϵ_p}	0.16
Persistence of AR(1)	ρ	0.94
Signal noise	$\sigma_{z_{tr}}$	0.38
Tax-like wedges	σ_{τ_p}	0.55
Adjustment costs	λ	0.11
Exogenous exit	γ	0.02

Table 7: Decomposition of life-cycle employment growth in data and model.

	Growth		Selection		Total	
Age group	Data	Model	Data	Model	Data	Model
0–2	0.13	0.16	0.08	0.04	0.21	0.20
3–5	0.02	0.01	0.09	0.03	0.11	0.04
6–10	0.00	-0.01	0.07	0.02	0.07	0.01
11–15	-0.01	-0.01	0.06	0.02	0.05	0.01

Notes: For each age bracket, the table reports the decomposition Total = Growth + Selection. Total is the difference between survivors' average log employment at the next age and the full population's average log employment at the source age. Growth is the difference between survivors' average log employment at the next age and at the source age (the targeted growth moment). Selection is the difference between survivors' and the full population's average log employment at the source age. The Growth column for ages 0–2 and 3–5 reports targeted moments; the remaining cells are not directly targeted. The data sample is restricted to firms tracked from birth.

Table 7 gives the growth decomposition. As discussed above, the model generates employment growth rates for surviving firms that are in line with the data. The model also replicates the observed total employment growth for the youngest age bracket. However, the selection component generated by the model is weaker than that in the data for all age brackets, causing total employment growth in the model to fall short of that in the data at older ages. Given that the model replicates the targeted exit rates for all age brackets, the most likely explanation is that exogenous exit, which affects firms equally, dominates the endogenous exit channel ($\gamma = 0.02$ in our calibration).

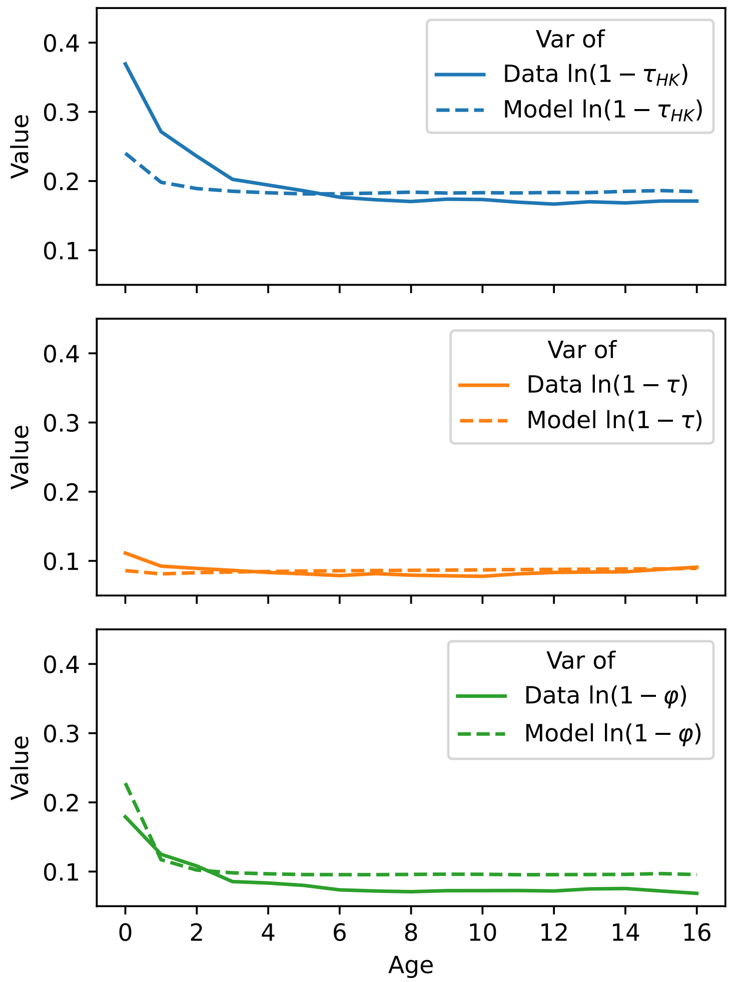

Figure 6: Uncertainty and misallocation conditional on age: model vs. data.

4.3 Static Misallocation in the Simulated Data

We now explore the misallocation patterns in simulated data generated by our model and compare them with the ones we observed in the Finnish data. In this regard, we redo our calculations of static misallocation that allowed us to decompose ex post misallocation into ex ante misallocation and uncertainty. We use the benchmark empirical results based on the pooled data given in Section 2.3 as the comparison point for our model.

Table 8 reports the indirect measure of ex post misallocation and its decomposition in data generated with our calibrated model. The variance of the HK revenue wedge is 0.191. In the Finnish data with our baseline setup, it was 0.206 when looking at all industries and years jointly. Thus, our model is able to explain around 93% of the observed variation in the HK revenue wedge.

Looking at the components of ex post misallocation, we observe that, without targeting the decomposition, the model generates levels of uncertainty and ex ante misallocation that are very close to the ones observed in the pooled data. The model-implied variance of $1-\varphi$ is 0.10. In the pooled data, this was 0.090. The variance of the tax-like wedge is 0.093, while it was 0.090 in Section 2.3. The main difference between the model and the empirical results is in the covariance term, which was 0.025, while here it is essentially zero. Due to this, the relative importance of uncertainty and ex ante misallocation is slightly higher in the model when compared to the empirical counterpart.

Table 8: The variance of the HK-style revenue wedge and its decomposition into uncertainty and ex ante misallocation, calculated from the simulated data.

Variable	Value	Share of $\mathrm{Var}(\ln(1-\tau_{HK}))$
$\mathrm{Var}(\ln(1-\tau_{HK}))$	0.191	1.00
$\mathrm{Var}(\ln(1-\tau))$	0.093	0.49
$\mathrm{Var}(\ln(1-\varphi))$	0.100	0.52
$2\mathrm{Cov}(\ln(1-\tau),\ln(1-\varphi))$	-0.002	-0.01

In Figure 6, we replicate the exercise shown in Figure 1. We calculate the variance of the HK revenue wedge conditional on the age of the firms and decompose it into a component reflecting uncertainty and the dispersion in the residual wedge. To facilitate comparisons with the empirical decomposition, we also reproduce the results given in Figure 1.

Starting with uncertainty, Figure 6 reveals that our model generates an age-dependent pattern that is broadly similar to the one in the data. The largest deviation between the model and the data is for the new firms: 0.228 versus 0.179. For old firms, the model generates a level of uncertainty that is slightly higher than in the data. For example, at age ten, the variance of $1-\varphi$ is 0.096 in the model and 0.072 in the data. In line with the aggregate results in Table 8 and the fact that there is no strong age-dependent trend in the observed ex ante misallocation, the life-cycle profile of the ex ante misallocation produced by the model is also close to the empirical one.

Finally, we see that the clearest quantitative difference between the model and the data is in ex post misallocation in the early life cycle. For example, in the first year, the model generates about 62% of the observed ex post misallocation (0.228 vs. 0.369). Given the model's ability to replicate the two independent components of ex post misallocation, namely uncertainty and ex ante misallocation, relatively well, the difference between the empirical accounting exercise and the one implied by the model arises from the covariance term.

As Section 2.5 already highlighted, one obvious explanation for this could be unmodeled heterogeneity, for example, in financial costs or production technology.

The baseline model also reproduces the negative relationship between the HK wedge and firm productivity. We estimate the elasticity of $\ln(1-\tau^{HK})$ with respect to $\ln(z)$ in the simulated data to be -0.471, compared to -0.33 in the data. Since the baseline assumes ex ante wedges are independent of productivity, almost all of the model's elasticity arises from the prediction-error channel and a small selection effect. In Section 5 we explore an extension that allows ex ante wedges to be correlated with productivity.

Table 9: Exits and prediction errors in the model.

	Dependent variable: Exit at the Beginning of Next Period	
	LPM	LPM
Constant	0.021	0.020
	(0.000)	(0.000)
$\ln(1-\varphi)$	0.047	0.038
	(0.000)	(0.000)
$L(\ln(1-\varphi))$		0.019
		(0.000)

Notes: L() refers to lagged value.

In the empirical section, more optimistic firms (i.e., firms with a higher prediction error) were more likely to exit in the following periods. This is also true in the model. Table 9 illustrates using a linear probability model. In the first column, we again predict exit at the beginning of the next period using $1-\varphi$. The second column also includes the prediction error of the previous period as an explanatory variable. In line with the empirical section, the lagged prediction error also matters for exits.

4.4 The Quantitative Significance of Uncertainty, Misallocation and Adjustment Costs

In order to evaluate how uncertainty and other factors affect output and productivity, we set up a benchmark economy without distortions or uncertainty, where in each period firms first observe their productivity components and then choose their labor input. Thus, there is neither learning nor intra-period uncertainty. Moreover, there are no adjustment costs or tax-like distortions. Otherwise, we use our calibrated parameter values. We normalize the aggregate output of this economy to 100 and then illustrate how adding frictions affects it in relative terms. Since the labor supply is fixed to one, this also gives the responses of aggregate labor productivity. The upper panel of Table 10 assembles the results of these counterfactuals. In all cases, we use the parameter values given in Table 6 to fix the levels of different frictions. Since part of the labor force is used for entry and fixed costs, aggregate labor productivity ($Y/\bar{N}$) and aggregate TFP (Y/N_{prod}) could differ in non-trivial ways. Given this, we also report the relative aggregate TFP responses in the lower panel of the table.

From Table 10, we see that uncertainty alone reduces output and TFP by about 14%. The relative contribution of uncertainty is similar when added on top of misallocation: output falls from 58.7 to 50.8 and TFP from 53.8 to 46.2, a 14% drop in both cases.

Even though in the model, as well as in the data, uncertainty and ex ante misallocation were equally important for ex post misallocation, their impact on aggregate output and TFP is quite different. Compared to the benchmark, the introduction of tax-like distortions leads to a 41% reduction in output and a 46% drop in TFP.

The main reason why uncertainty is substantially less costly than ex ante misallocation in the model is that uncertainty distorts resource allocation only temporarily. In our calibration, firms learn their current persistent productivity relatively fast and adjust their size accordingly. When learning happens more slowly, the costs associated with uncertainty increase substantially. By the same logic, allowing misallocation itself to have a transitory component substantially reduces its aggregate cost. Under the alternative specifications explored in Section 5, the cost of misallocation is comparable to the cost of uncertainty.

Finally, adjustment costs alone reduce aggregate output and TFP by 2%. The relative drop is of the same magnitude when adjustment costs are added to an economy that already has ex ante misallocation in place. However, when uncertainty is present, the additional effect of adjustment costs becomes small or slightly positive. One potential explanation for this is that adjustment costs shift employment towards older firms, which, in the presence of uncertainty, use labor more effectively than younger firms.

5 Correlation Between Productivity and Measured Wedges

A large body of research emphasizes the importance of correlated wedges; if high-productivity firms appear to be taxed while low-productivity firms are subsidized, the aggregate costs of misallocation increase. Importantly, empirical evidence shows that this negative correlation can be quite strong for many countries (see, e.g., Bento and Restuccia, 2017, for estimates

Table 10: The aggregate effects of uncertainty, misallocation and adjustment costs.

	No uncertainty	Uncertainty
Panel A. Output		
No distortions or adjustment costs	100.00	86.37
Adjustment cost	97.72	86.00
Misallocation	58.72	50.76
Misallocation and adjustment costs	57.82	51.34
Panel B. TFP		
No distortions or adjustment costs	100.00	86.01
Adjustment cost	97.92	85.48
Misallocation	53.77	46.15
Misallocation and adjustment costs	52.96	46.76

Notes: Uncertainty refers to a case where firms do not know their persistent productivity and choose their labor before they know the current-period productivity. In order to shut down tax-like distortions, we set $\sigma_{\tau_p} = 0$.

of the elasticity of wedges with respect to productivity).

Earlier, we demonstrated that this pattern also holds for Finland. We further showed that our baseline model generates a similar correlation between ex post wedge and productivity. However, in the model this arises mainly from the co-movement between prediction error and productivity, whereas the empirical results in Section 2 suggest a more balanced contribution: both the correlation between realized productivity and prediction error, and the correlation between productivity and the ex ante wedge, play a role. In this section, we examine the sensitivity of our quantitative results to a more flexible specification of distortions, one that allows ex ante wedges to be correlated with productivity while still ensuring that firms remain uncertain about their productivity. We also explore the sensitivity of our baseline results to an alternative calibration of the productivity process that builds on the autocovariance structure of observed productivity.

5.1 Extended Model

In line with Section 3, firm productivity depends on transitory and persistent components (see equations (15) and (16)), which the firm cannot decompose. However, now we assume ex ante wedges of the form

$$\ln(1-\tau) = \tau_p + \tau_z.$$

As in the baseline model, τ_p is an observable, normally distributed permanent wedge drawn upon entry. The new component, τ_z, is defined as

$$\tau_z = \tau_{tr} + \iota z_{p,-},$$

where $\tau_{tr} \sim N(0, \sigma^2_{\tau_{tr}})$ is a transitory shock drawn each period. The second term ties the wedge to the previous period's permanent productivity. This specification avoids the need for an additional state variable, z_-, at the beginning of the period and ensures that the ex ante wedge, observable at the beginning of the period, does not depend on still-unrealized current productivity. To preserve the learning structure over the life cycle, we assume that the firm observes only τ_z but not its components. Transitory shock τ_{tr} plays a dual role in this setup: it allows for a flexible wedge structure and, by obscuring $z_{p,-}$ in the composite signal τ_z, preserves the firm's inference problem.

Thus, the firm now receives two signals about its permanent productivity: τ_z at the beginning of the period (after the exit decision) and z at the end of the period. The structure leads to sequential Kalman filtering. At the start of the period, the firm has a prediction about its permanent productivity in the previous period, $m_{-|z_-}$, with variance $\Sigma_{-|z_-}$. This prediction is based on information up to the point of observing productivity z_-. Using this information, the firm makes the exit/stay decision. Conditional on continuing, the firm observes τ_z, updates its prediction of the previous period's productivity, $m_{-|\tau,z_-}$, and forms a prediction of current productivity to decide employment. Finally, at the end of the period, the firm observes z and updates its prediction of current productivity. The Bellman equation for the incumbent and the relevant Kalman filtering equations are provided in Appendix F. The rest of the model structure follows the setup in Section 3.

5.2 Alternative Calibrations

The introduction of the more flexible ex ante wedge structure complicates the identification of model parameters. To proceed, we adopt an alternative calibration strategy in which we first fix the parameters governing the information friction (ρ, σ_{ε_p} and $\sigma_{z_{tr}}$) using measured productivity (together with firms' exit behavior). We then calibrate the remaining parameters by targeting the same set of observables as in Section 4. Under this approach, the size distribution and growth rates play a particularly important role in pinning down the wedge parameters.

We start by constructing a firm-level productivity measure following HK. Given the production function, demand structure and demand elasticity introduced in Section 2, firm-level TFP can be recovered (up to an industry-specific constant) by dividing value added, raised to the power of the markup, by a composite input index (see equation (7)).

Without endogenous exits, we could use the autocovariance function of productivity, z, at lags 0, 1 and 2 to identify ρ, σ_{ε_p} and $\sigma_{z_{tr}}$. Our solution to the complications arising from selection follows Guvenen et al. (2021), who estimate an earnings process for workers while accounting for mobility between employment and nonemployment, allowing the probability of nonemployment to vary with individual productivity and worker age. We adapt this approach to the firm setting and use the method of simulated moments to jointly estimate a logistic exit function (in which the probability of exit depends on firm productivity and age) together with the autocovariances of the productivity process. The details of this approach, along with the empirical targets, are provided in Appendix F.

This approach leads to the following parameterization: $\rho = 0.98$, $\sigma_{\varepsilon_p} = 0.14$ and $\sigma_{z_{tr}} = 0.48$. Compared to the parameter values in Section 4, the process for z_p is more persistent, and the noise component is more volatile. With these parameters fixed, we calibrate the remaining parameters for three different model variants.

The first model (M1) corresponds to the original specification. It allows us to assess the sensitivity of the results in the previous section to an alternative calibration in which the information friction is directly tied to the evolution of the productivity process. M1 also serves as a benchmark for analyzing the effects of more flexible distortions. The second model (M2) introduces the temporary ex ante wedge but imposes $\iota = 0$. Finally, the third model (M3) incorporates both temporary wedges and correlated wedges in addition to the permanent component.

Tables 19 and 20 in Appendix F report the models' fits to the empirical targets and the corresponding parameter values. Overall, the model fits are quite good and broadly consistent with the baseline results. However, one noticeable deviation concerns the growth rates of young firms: the alternative calibration strategy produces too rapid growth during the first three years of their life cycle, with too much variation across firms. This pattern holds for all three models considered and reflects the increased persistence in productivity under the new parameterization.

Examining the parameter values reveals that both entry and adjustment costs rise to offset the increased heterogeneity in productivity. A key difference between M1 and the other two models is that introducing more flexible wedges substantially reduces the variance of the permanent wedge component. In M3, the data favors a modest negative correlation between the ex ante wedge and persistent productivity, with $\iota = -0.1$.

5.3 Results for Alternative Calibrations

Table 21 in Appendix F reports the static misallocation generated by our models. Our baseline model with the alternative calibration, M1, produces more ex post misallocation than observed in the data. Consistent with this, both ex ante misallocation and uncertainty are higher than in the empirical results reported in Section 2.

The models with more flexible wedge structures, either without correlated wedges (M2) or with correlated wedges (M3), generate ex post misallocation levels that align well with the data. Across the three models, the levels of uncertainty and ex ante misallocation are closest to their empirical counterparts for M3. However, like M2, M3 generates somewhat too high uncertainty and too low ex ante misallocation. The elevated role of uncertainty across all models is in line with the alternative calibration strategy increasing the volatility of the production process, which in turn creates more variation in growth rates.

All three models (M1-M3) produce a small negative covariance term between prediction errors and ex ante wedges. The negative sign comes from adjustment costs, which depress $\pi/(wn)$ while raising $py/(wn)$ for growing firms; the static decomposition then assigns them higher $\ln(1-\varphi)$ but lower $\ln(1-\tau)$, generating negative covariance between measured uncertainty and the residual ex ante wedge. The adjustment costs are higher in M1-M3 than in the baseline because the alternative calibration estimates a more volatile productivity process, which requires higher adjustment costs to match the calibrated growth volatility.

In M3, the $\iota < 0$ mechanism partially offsets this negative covariance: high-productivity firms tend to underestimate their productivity in the first years of operation while simultaneously facing higher ex ante wedges, generating positive comovement for this group. As a result, M3 produces the covariance term closest to zero among the three models, though the adjustment-cost channel still slightly dominates.

For all models, there is a negative correlation between productivity and the ex post wedge. The elasticities of the HK wedge with respect to productivity are -0.40, -0.37 and -0.39 for M1, M2 and M3, respectively, compared with -0.33 in the data. Similar to the baseline model, this pattern largely reflects the association between productivity and prediction errors. The prediction-error component accounts for 83%, 95% and 77% of the total elasticity of the HK wedge with respect to productivity in M1, M2 and M3 respectively, measured as the ratio of the prediction-error elasticity to the total elasticity.[18]

However, the covariance between productivity and the ex ante wedge also plays a non-trivial role. In M1 and M2 this arises from selection: firms with low productivity and high "taxes" are more likely to exit. The selection effect is strongest in M1, where the

[18] The elasticities of $\ln(1-\varphi)$ with respect to z are -0.33, -0.35 and -0.30 for M1, M2 and M3, respectively.

variance of permanent wedges is largest. In M3, the selection effect is amplified by the direct correlation between the ex ante wedge and lagged productivity. Nevertheless, even in M3 the contribution of this correlation remains smaller than in the data, where the two components contributed roughly equally. Matching that pattern would require a more negative value of ι.

Table 11: Frictions and aggregate TFP across alternative models.

	M1	M2	M3
Panel A. Output			
No distortions	100.00	100.00	100.00
Misallocation and adjustment costs	53.95	88.00	88.18
Misallocation, adjustment costs and uncertainty	46.27	74.36	80.44
Panel B. TFP			
No distortions	100.00	100.00	100.00
Misallocation and adjustment costs	53.84	87.74	88.13
Misallocation, adjustment costs and uncertainty	45.68	73.27	80.15

Notes: M1 refers to the original model under the alternative calibration strategy; M2 adds temporary wedges; M3 additionally allows ex ante wedges to be correlated with productivity.

Table 11 illustrates how misallocation and uncertainty affect aggregate output and TFP. Each model is first solved without uncertainty, adjustment costs or misallocation wedges. We then compute the relative decline in aggregate output (TFP) when ex ante misallocation, with adjustment costs, or all distortions are introduced.[19]

The first thing that stands out from the table is that the role of misallocation declines substantially once non-permanent wedges are introduced. For M1, the relative drop in output and TFP due to misallocation wedges and adjustment costs is about 46%, which is in line with the baseline results in Section 4. For M2 and M3, the effect of ex ante misallocation, with adjustment costs, is only 12%. This highlights the impact of the decline in the variation in the persistent wedge between M1 and the two other models.

The second noteworthy observation is that the alternative calibration strategy does not substantially change the aggregate effects of the uncertainty channel: the relative drop in output and TFP is now about 9–16% compared with 14% in the baseline model.

[19]Given that uncertainty is fundamentally related to the ex ante wedge in M3, we do not report the effects of uncertainty alone.

6 Conclusions

In this paper, we develop a method to quantify the level of idiosyncratic uncertainty and distinguish it from firm-level ex ante wedges. The approach builds on a minimum amount of theory, just specifying demand structure and production function. In our setup, the two key empirical ratios to pin down uncertainty and ex ante misallocation are profits-to-wage-bill and value-added-to-wage-bill. We then explore the importance of uncertainty and ex ante misallocation for the observed ex post misallocation utilizing Finnish administrative data that gives us annual high-quality information of nearly all Finnish firms. According to our findings, uncertainty plays at least as important a role as ex ante misallocation, accounting for at least 40% of ex post misallocation across our specifications. Moreover, we also show that there is a strong decreasing age-dependent trend in uncertainty.

To understand these empirical results, we set up a life-cycle model of firm dynamics where firms are uncertain about their fundamentals. We match the model with the growth profiles, size distribution and selection patterns in the Finnish data and show that the model produces patterns of uncertainty and ex ante misallocation in line with our empirical observations. We then use the calibrated model to evaluate the importance of uncertainty and misallocation for aggregate output and TFP.

According to our baseline counterfactuals, uncertainty reduces aggregate output and TFP by about 14%. This effect is robust to alternative modeling choices, including the calibration of the productivity process and the introduction of more flexible wedge structures with transitory and correlated components. Ex ante misallocation has a larger effect under our baseline calibration with permanent wedges, reducing output (TFP) by 41% (46%), but this number drops to a level comparable to the cost of uncertainty under the more flexible wedge specifications.

Given the substantial output losses associated with uncertainty across all our specifications, it is natural to ask whether and how policy can affect it. It seems plausible that idiosyncratic uncertainty, at large, is a part of the economic environment and is not mainly directly driven by policies. However, this does not mean that some policies might not indirectly affect uncertainty. For example, any policies that affect the age distribution of firms also influence the observed variation in prediction errors, as older firms generally make more precise forecasts. Thus, startup grants and other government support for new businesses also indirectly contribute to misallocation through their effect on the age distribution.

References

Arkolakis, C., T. Papageorgiou, and O. A. Timoshenko (2018). Firm learning and growth. *Review of Economic Dynamics 27*, 146–168.

Asker, J., A. Collard-Wexler, and J. De Loecker (2014). Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy 122*(5), 1013–1063.

Atkeson, A. and A. Burstein (2010). Innovation, Firm Dynamics, and International Trade. *Journal of Political Economy 118*(3), 433–484.

Bai, Y., D. Lu, and X. Tian (2018). Do Financial Frictions Explain Chinese Firms' Saving and Misallocation? Working Paper 24436, National Bureau of Economic Research.

Baqaee, D. R. and E. Farhi (2020). Productivity and Misallocation in General Equilibrium. *The Quarterly Journal of Economics 135*(1), 105–163.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2013). Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review 103*(1), 305–334.

Bayer, C., A. Mecikovsky, and M. Meier (2018). Misallocation: Markups and Technology. *CEPR Discussion Papers* (12727).

Bento, P. and D. Restuccia (2017). Misallocation, Establishment Size, and Productivity. *American Economic Journal: Macroeconomics 9*(3), 267–303.

Bils, M., P. J. Klenow, and C. Ruane (2021). Misallocation or Mismeasurement? *Journal of Monetary Economics 124*, S39–S56.

Boar, C., D. Gorea, and V. Midrigan (2022). Why Are Returns to Private Business Wealth So Dispersed? Working Paper 29705, National Bureau of Economic Research.

Buera, F. J., J. P. Kaboski, and Y. Shin (2011). Finance and Development: A Tale of Two Sectors. *American Economic Review 101*(5), 1964–2002.

Busso, M., L. Madrigal, and C. Pagés (2013). Productivity and resource misallocation in Latin America. *The B.E. Journal of Macroeconomics 13*(1), 903–932.

Cavalcanti, T. V., J. P. Kaboski, B. S. Martins, and C. Santos (2021). Dispersion in Financing Costs and Development. Working Paper 28635, National Bureau of Economic Research.

Cirera, X., R. Fattal-Jaef, and H. Maemir (2020). Taxing the Good? Distortions, Misallocation, and Productivity in Sub-Saharan Africa. *The World Bank Economic Review 34*(1), 75–100.

Clementi, G. L. and B. Palazzo (2016). Entry, Exit, Firm Dynamics, and Aggregate Fluctuations. *American Economic Journal: Macroeconomics 8*(3), 1–41.

David, J. M., H. A. Hopenhayn, and V. Venkateswaran (2016). Information, Misallocation, and Aggregate Productivity. *The Quarterly Journal of Economics 131*(2), 943–1005.

David, J. M. and V. Venkateswaran (2019). The Sources of Capital Misallocation. *American Economic Review 109*(7), 2531–2567.

De Loecker, J. and F. Warzynski (2012). Markups and Firm-Level Export Status. *American Economic Review 102*(6), 2437–2471.

Eslava, M., J. Haltiwanger, and N. Urdaneta (2023). The Size and Life-Cycle Growth of Plants: The Role of Productivity, Demand, and Wedges. *The Review of Economic Studies*, rdad029.

Feng, Y. (2025). Firm life-cycle learning and misallocation. *The Review of Economics and Statistics*, 1–49.

Foster, L., J. Haltiwanger, and C. Syverson (2016). The Slow Growth of New Plants: Learning about Demand? *Economica 83*(329), 91–129.

Gao, H. and L. Zhang (2024). Uncover "Gazelles": The Macroeconomic Implications of Uncertainty and Learning for Entrepreneurship. *Working Paper*.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2013). Misallocation and financial market frictions: Some direct evidence from the dispersion in borrowing costs. *Review of Economic Dynamics 16*(1), 159–176. Special issue: Misallocation and Productivity.

Gollin, D. and C. Udry (2020). Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture. *Journal of Political Economy 129*(1), 1–80.

Gorodnichenko, Y., E. G. Mendoza, and L. L. Tesar (2012). The Finnish Great Depression: From Russia with Love. *American Economic Review 102*(4), 1619–44.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of u.s. workers reveal about lifecycle earnings dynamics? *Econometrica 89*(5), 2303–2339.

Haltiwanger, J., R. Kulick, and C. Syverson (2018). Misallocation Measures: The Distortion That Ate the Residual. Working Paper 24199, National Bureau of Economic Research.

Hopenhayn, H. and R. Rogerson (1993). Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *Journal of Political Economy 101*(5), 915–938.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Hsieh, C.-T. and P. J. Klenow (2014). The Life Cycle of Plants in India and Mexico. *The Quarterly Journal of Economics 129*(3), 1035–1084.

Jovanovic, B. (1982). Selection and the Evolution of Industry. *Econometrica 50*(3), 649–670.

Klenow, P. and A. Rodríguez-Clare (1997). *The Neoclassical Revival in Growth Economics: Has It Gone Too Far?*, pp. 73–114. MIT Press.

Ljungqvist, L. and T. J. Sargent (2018). *Recursive Macroeconomic Theory* (4 ed.). The MIT Press.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*(6), 1695–1725.

Midrigan, V. and D. Y. Xu (2014). Finance and Misallocation: Evidence from Plant-Level Data. *American Economic Review 104*(2), 422–458.

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

Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics 11*(4), 707–720.

Rotemberg, M. and T. K. White (2021). Plant-to-Table(s and Figures): Processed Manufacturing Data and Measured Misallocation. *Working Paper*.

Sterk, V., P. Sedláček, and B. Pugsley (2021). The Nature of Firm Growth. *American Economic Review 111*(2), 547–579.

Tian, C. (2022). Learning and firm dynamics in a stochastic equilibrium. *Journal of Economic Theory 203*, 105486.

Yang, M.-J. (2021). Micro-level Misallocation and Selection. *American Economic Journal: Macroeconomics 13*(4), 341–368.

Appendix A. Additional Descriptive Statistics

This appendix provides basic descriptive statistics for the variables used in our accounting exercise and in the calibration of the model. After applying our filtering criteria defined in Section 2 we have 741915 firm-year observations. Table 12 gives the number of firms for each year.

Table 12: Number of firms in the sample by year

Year	Number of firms
1995	31 377.00
1996	33 489.00
1997	37 154.00
1998	38 465.00
1999	38 807.00
2000	39 656.00
2001	40 153.00
2002	40 858.00
2003	40 897.00
2004	41 469.00
2005	41 286.00
2006	42 470.00
2007	44 706.00
2008	45 905.00
2009	45 276.00
2010	45 885.00
2011	47 060.00
2012	47 002.00

The key statistics are presented in Table 13. The upper panel reports the means and variances for capital, k (measured by total assets); value added, py; sales, py_{sales}; materials, $p_m m$; profits, π; employment compensation, wl; and full-time-equivalent (FTE) employment, l. All monetary variables are expressed in logs of 1,000 euros, whereas employment is reported in logs only. To evaluate growth rates, we also report log differences for each variable. The leftmost columns present these statistics for the full sample, regardless of whether firm age can be defined, while the remaining columns provide the corresponding statistics conditional on firm age. The category "age n.i." refers to firms for which age cannot be identified. The lower panel of Table 13 reports the number of firm-year observations and exit rates for each age bracket, along with the overall entry rate.

Table 13 illustrates life-cycle growth in inputs and outputs. Young surviving firms grow rapidly on average, with substantial variation in their growth profiles. As firms age, growth slows, and the dispersion in growth rates narrows. The table also shows considerable dispersion in the levels of inputs and outputs.

Table 13: Descriptive statistics of key variables in logs (no winsorization) and first differences, overall and by firm age group

	Overall		Age 0–2		Age 3–5		Age 5–10		Age 10–15		Age 15+		Age n.i.	
Variable	Mean	Variance	Mean	Variance	Mean	Variance	Mean	Variance	Mean	Variance	Mean	Variance	Mean	Variance
k	12.61	2.74	12.01	2.36	12.41	2.34	12.77	2.46	13.07	2.62	13.55	3.50	12.47	2.78
Δk	0.06	0.18	0.12	0.38	0.08	0.20	0.05	0.14	0.03	0.10	0.01	0.07	0.04	0.25
py	12.21	2.01	11.74	1.75	12.09	1.74	12.36	1.84	12.58	1.96	12.92	2.55	12.00	2.02
Δpy	0.04	0.31	0.17	0.48	0.06	0.31	0.02	0.25	0.00	0.24	−0.01	0.25	0.02	0.44
$(py)_{\text{sales}}$	13.28	2.19	12.83	1.84	13.16	1.91	13.42	2.05	13.63	2.20	13.99	2.86	13.12	2.23
$\Delta(py)_{\text{sales}}$	0.05	0.18	0.16	0.28	0.06	0.19	0.03	0.14	0.01	0.13	0.01	0.13	0.03	0.26
ms	12.67	2.71	12.22	2.37	12.54	2.44	12.81	2.56	13.02	2.68	13.39	3.23	12.53	2.87
Δms	0.05	0.26	0.16	0.42	0.06	0.27	0.03	0.21	0.01	0.18	0.02	0.16	0.03	0.36
π	10.13	3.28	9.76	3.04	9.99	3.00	10.23	3.13	10.49	3.34	10.75	4.04	10.02	3.32
$\Delta\pi$	0.03	1.71	0.02	2.14	0.07	1.70	0.03	1.59	0.02	1.56	0.02	1.66	0.02	1.85
wl	11.90	1.95	11.43	1.75	11.75	1.72	12.03	1.76	12.29	1.82	12.69	2.35	11.71	1.96
Δwl	0.07	0.21	0.22	0.37	0.08	0.24	0.05	0.15	0.03	0.12	0.01	0.11	0.05	0.30
l	1.65	1.35	1.37	1.04	1.58	1.18	1.73	1.35	1.84	1.48	2.11	1.96	1.53	1.30
Δl	0.02	0.10	0.13	0.16	0.02	0.10	0.00	0.07	−0.01	0.07	−0.02	0.07	0.00	0.13
n	741 915.00		152 638.00		145 469.00		176 728.00		119 999.00		51 698.00		95 383.00	
Exit rate	0.03		0.05		0.03		0.02		0.02		0.02		0.00	
Entry rate	0.05													

Notes: All monetary variables are expressed in logs of 1,000 euros, whereas employment is reported in logs without scaling.

Appendix B. Misallocation for Different Industries

This appendix explores the static misallocation and its decomposition for different industries. In addition, we also compare exporting and non-exporting firms.

Table 14 reproduces the static decomposition of Table 1, but focuses on one industry at a time. The main observation from the table is that, excluding electricity, gas and water supply, the results are remarkably stable across industries. Ex post misallocation levels are close to the level of pooled data (though slightly higher in wholesale and retail trade). The same is also true for the components of ex post misallocation. Overall, it seems that uncertainty and ex ante misallocation are more or less equally important, both accounting for around 40% of the total variation. The relative importance of covariance term is typically around 20%.

Given these general patterns, electricity, gas and water supply (Industry E) stands out. The variance of $\ln(1-\tau)$ is more than 4 times higher than in other industries. Consistent with this, ex post misallocation is also at a substantially higher level. A distinctive feature of this industry is that it is much more heavily regulated than other industries. In addition, competition is limited at best: gas supply is in the hands of a nationwide monopoly, while water supply and electric grid are run by local monopolies. Moreover, local governments own the water supply companies and are also often major owners of electricity companies. These sector-specific features are likely behind the elevated level of misallocation. Another difference between Industry E and the rest of the economy is the negative covariance term. Uncertainty, on the other hand, is on a par with other industries. Together these facts imply that mots of the variation, almost 80%, in the HK wedge is accounted for by ex ante

Table 14: Ex post misallocation and its decomposition for different industries.

	$\text{Var}(\ln(1-\tau_{HK}))$	$\text{Var}(\ln(1-\tau))$	$\text{Var}(\ln(1-\varphi))$	$2\text{Cov}(\ln(1-\tau),\ln(1-\varphi))$
Construction (F)	0.174	0.074 [43%]	0.080 [46%]	0.019 [11%]
Manufacturing (D)	0.175	0.090 [51%]	0.080 [46%]	0.006 [3%]
Electricity, gas and water supply (E)	0.536	0.375 [70%]	0.109 [20%]	0.052 [10%]
Wholesale and retail trade (G)	0.265	0.104 [39%]	0.108 [41%]	0.053 [20%]
Hotels and restaurants (H)	0.207	0.088 [43%]	0.092 [44%]	0.028 [14%]
Transport, storage and communication (I)	0.188	0.097 [52%]	0.080 [43%]	0.011 [6%]

Notes: The relative contributions are given in brackets.

misallocation.

One of the characteristic features of the Finnish economy that sets it apart from large countries such as the US, China or India is the relatively high importance of exports. This might be relevant for us if wedges (ex ante wedges or prediction errors) or their joint variation differ systemically between exporters and non-exporters. For example, markups could differ between exporters or non-exporters (see De Loecker and Warzynski, 2012) or currency fluctuations might affect prediction errors differently. Our data allows us to observe firms' exports for years 1995-2007. However, when calculating our decomposition for exporters alone, we observe that this dimension does not affect our results. For pooled data over years 1995-2007 with exporters only, we get $\text{Var}(\ln(1-\tau_{HK})) = 0.217$, $\text{Var}(\ln(1-\tau)) = 0.101$, $\text{Var}(\ln(1-\varphi)) = 0.092$ and $2\text{Cov}(\ln(1-\tau), \ln(1-\varphi)) = 0.024$.

Appendix C. Year and Cohort Effects in Misallocation

This appendix reports the results of the baseline accounting exercise for selected years and cohorts.

In Table 15 we measure misallocation in different years. The dispersion of the residual wedge, our measure of ex ante misallocation, is again stable over the subsets of the data. Moreover, the values are close to the ones observed in the pooled data. The same is true for the measure of uncertainty and the covariance term with the exception of the year 1996.

Table 15: The variance of HK-style revenue wedge and its decomposition to uncertainty and residual wedge for different years.

	$\text{Var}(\ln(1-\tau_{HK}))$	$\text{Var}(\ln(1-\tau))$	$\text{Var}(\ln(1-\varphi))$	$2\text{Cov}(\ln(1-\tau),\ln(1-\varphi))$
Year 1996	0.313	0.089 [28%]	0.164 [52%]	0.060 [19%]
Year 2001	0.188	0.081 [43%]	0.078 [41%]	0.029 [15%]
Year 2006	0.197	0.097 [49%]	0.088 [45%]	0.013 [7%]
Year 2011	0.191	0.095 [50%]	0.079 [41%]	0.017 [9%]

Notes: The relative contributions are given in brackets.

In 1996, Finland was still recovering from the most severe recession of its peacetime history and the structural change triggered by the collapse of trade with the Soviet Union was still underway (see, e.g., Gorodnichenko et al, 2012). In light of this, an increased level of uncertainty is perhaps to be expected. [20]

Figure 7 presents the evolution of the decomposition for selected cohorts of firms. We have chosen these cohorts such that the starting years are in line with the years reported in Table 15. The figure illustrates that our main conclusion, decreasing trends in ex post misallocation and uncertainty, is present for all cohorts. It is worth noting, though, that the first year of the 1996 cohort stands out. From Table 15, we already know that this year ex post misallocation and uncertainty were exceptionally high. This is clearly true for startups as well. Interestingly, however, at their second year and thereafter, the 1996 cohort does not stand out when compared against other cohorts.

Appendix D. Accounting Exercise with Capital income

In this appendix, we redo our measurement of ex post misallocation and its components using capital income rather than wage-bill information. Compared to labor costs, which are measured at market prices and collected for taxation purposes, we only observe accounting-based measures of capital, which we combine with the assumption of a common rental rate. In particular, we use two measures of firms' capital stock: total assets (net of short-term

[20] We also observe similar high values for year 1995. From 1997 onwards the observed values are in line with 2006 and 2011.

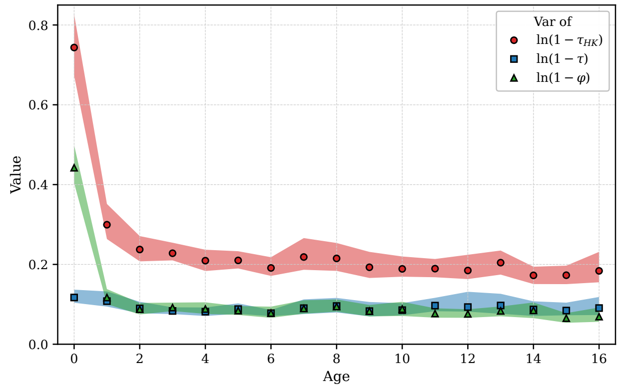

(a) 1996

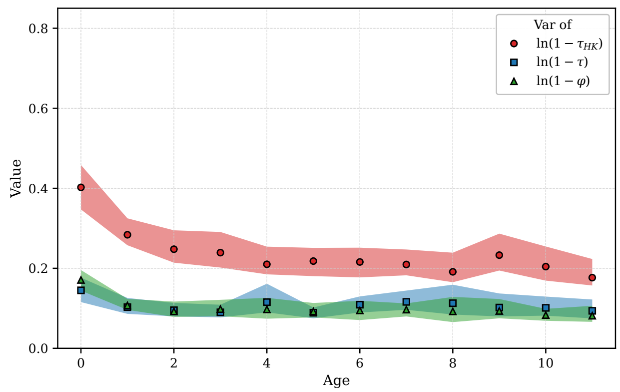

(b) 2001

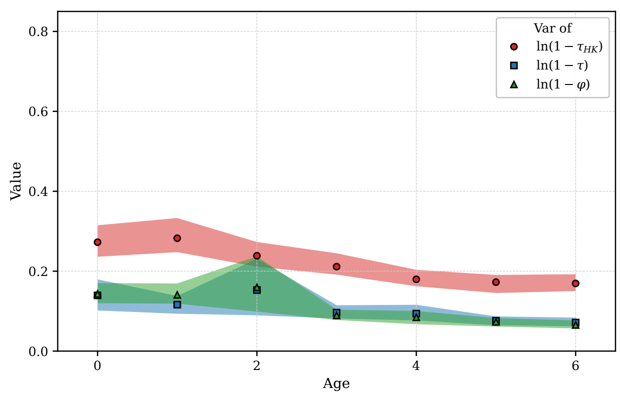

(c) 2006

Figure 7: Life cycle of uncertainty and misallocation for selected cohorts.

Notes: Shaded areas represent 95% confidence intervals from a block bootstrap with individual firms as blocks.

debt) and, following HK, fixed capital (machinery, equipment, and structures).

We report the results of the misallocation accounting exercise using capital income in Table 16. The level of ex post misallocation is sensitive to the choice of capital measure. With total assets, the variance of the HK revenue wedge is 0.91, while with fixed capital

	$\mathrm{Var}(\ln(1-\tau_{HK}))$	$\mathrm{Var}(\ln(1-\tau))$	$\mathrm{Var}(\ln(1-\varphi))$	$2\mathrm{Cov}(\ln(1-\tau), \ln(1-\varphi))$
Total assets	0.914	0.746 [82%]	0.121 [13%]	0.048 [5%]
Fixed capital	1.805	0.944 [52%]	0.760 [42%]	0.101 [6%]

Table 16: Variance decomposition for capital cases

it nearly doubles to 1.81. Both estimates are substantially higher than our baseline labor-based results, with levels that, if taken at face value, would imply very large TFP gains from reallocation, well in excess of those typically reported for advanced economies.

Turning to the decomposition, the two capital measures yield qualitatively different pictures. With total assets, the level of uncertainty is on par with the baseline results, but ex ante misallocation is substantially elevated both in level and in its relative contribution. It accounts for 82% of the total variation, while uncertainty contributes 13%. With fixed capital, the decomposition is much more balanced, with ex ante misallocation accounting for 52% and uncertainty for 42%, in line with our baseline labor-based findings. In this case, however, both uncertainty and ex ante misallocation are substantially higher in levels than in the baseline. The covariance term remains small under both measures.

The sensitivity of the results to the choice of input measure reflects the variances and covariances of the underlying ratios. Ex post misallocation is pinned down by the variance of the ratio of value added to input costs, whether $\ln(py/wl)$ or $\ln(py/Rk)$. The variance of this ratio depends on the dispersion of both components and their covariance. In our data, the variance of log capital income is substantially larger than the variance of log labor costs under both of our capital measures. This is the primary reason why capital-based ex post misallocation exceeds its labor-based counterpart. Moreover, the variance of log capital income is itself highly sensitive to how capital is measured, being considerably larger for fixed capital than for total assets. This sensitivity translates directly into substantial differences in measured ex post misallocation across the two capital specifications.

A similar logic applies to the uncertainty component, which is primarily identified from the profits-to-input-costs ratio. Here, however, the covariance channel also plays a role: profits co-move more strongly with total assets than with fixed capital, compressing measured uncertainty under the former. This is consistent with total assets including current assets, such as receivables and inventories, that mechanically track current-period profitability. Together with the differences in the dispersion of capital income, this covariance channel accounts for the shift in the decomposition across the two capital measures.

Appendix E. Labor distortions

Like most of the literature, we have not allowed for the use of labor to be affected by an output wedge and an input wedge jointly. Instead, we have focused on the decomposition between the prediction error and the output wedge. In this appendix, we generalize our accounting exercise to also include labor distortions that directly affect profits. To do so we combine our baseline approach which uses information in profits to our alternative approach that relies on the material use.

With this extension, the HK wedge identified from the first-order condition of labor is a combination of output distortion $(1-\tau^y_{t,s,i})$, labor distortion $(1+\tau^n_{t,s,i})$ and prediction error $(1-\varphi_{t,s,i})$. To separate these wedges, we fallow our alternative approach and assume that a firm chooses its materials after uncertainty has been revealed. Moreover, we also assume that the firm's use of materials is not distorted by an input wedge. Now we can identify the output distortion from the first-order condition of materials. Next, the ex post profit condition allows us to pin down the prediction error. When the prediction error and the output wedge are uncorrelated, a violation of our identification assumptions would underestimate the uncertainty and overestimate the misallocation arising from the output wedges. Finally, the first-order condition with respect to labor enable us to quantify the labor wedge.

The firm's problem is the following:

$$\max_{k_{t,s,i},n_{t,s,i}} \left\{ \mathbb{E}\left[\max_{m_{t,s,i}} (1-\tau^y_{t,s,i}) p_{t,s,i} y_{t,s,i} - p^m_{t,s,i} m_{t,s,i} \right] - (1+\tau^n) w_{t,s,i} n_{t,s,i} - R_t k_{t,s,i} \right\},$$

where $p_{t,s,i} = P_t \left(\frac{y_{t,s,i}}{Y_t}\right)^{-1/\sigma}$ and $y_{t,s,i} = z_{t,s,i} k^{\alpha_s}_{t,s,i} n^{\xi_s-\alpha_s}_{t,s,i} m^{1-\xi_s}_{t,s,i}$. Let us define

$$\alpha_{1,s} \equiv \left(1-\frac{1}{\sigma}\right)\alpha_s, \qquad \alpha_{2,s} \equiv \left(1-\frac{1}{\sigma}\right)(\xi_s-\alpha_s), \qquad \alpha_{3,s} \equiv \left(1-\frac{1}{\sigma}\right)(1-\xi_s).$$

Because materials are chosen after productivity is observed, the materials first-order condition identifies the output wedge:

$$(1-\tau^y_{t,s,i})\alpha_{3,s} p_{t,s,i} y_{t,s,i} = p^m_{t,s,i} m_{t,s,i},$$

so that

$$1-\tau^y_{t,s,i}=\frac{\sigma}{\sigma-1}\frac{1}{(1-\xi_s)\frac{p_{t,s,i}y_{t,s,i}}{p^m_{t,s,i}m_{t,s,i}}}. \tag{E.1}$$

The equation is exactly the same as in the main text when we use our alternative identification strategy.

The labor and capital first-order conditions are

$$\begin{aligned}(1+\tau^n_{t,s,i})w_{t,s,i}n_{t,s,i}&=\alpha_{2,s}(1-\tau^y_{t,s,i})(1-\varphi_{t,s,i})p_{t,s,i}y_{t,s,i},\\ R_tk_{t,s,i}&=\alpha_{1,s}(1-\tau^y_{t,s,i})(1-\varphi_{t,s,i})p_{t,s,i}y_{t,s,i}.\end{aligned}$$

Observed profits satisfy

$$\pi_{t,s,i}=p_{t,s,i}y_{t,s,i}-p^m_{t,s,i}m_{t,s,i}-(1+\tau^n_{t,s,i})w_{t,s,i}n_{t,s,i}-R_tk_{t,s,i}.$$

Substituting the three optimality conditions into this expression gives

$$\pi_{t,s,i}=\left[1-(1-\tau^y_{t,s,i})\left(\alpha_{3,s}+(\alpha_{1,s}+\alpha_{2,s})(1-\varphi_{t,s,i})\right)\right]p_{t,s,i}y_{t,s,i}.$$

Using $\alpha_{1,s}+\alpha_{2,s}=\left(1-\frac{1}{\sigma}\right)\xi_s$ and $\alpha_{3,s}=\left(1-\frac{1}{\sigma}\right)(1-\xi_s)$, we get

$$1-\varphi_{t,s,i}=\frac{1-\xi_s}{\xi_s}\frac{p_{t,s,i}y_{t,s,i}-\pi_{t,s,i}-p^m_{t,s,i}m_{t,s,i}}{p^m_{t,s,i}m_{t,s,i}}. \tag{E.2}$$

Finally, combining (E.1) and (E.2) with the labor first-order condition gives an expression for the input wedge:

$$1+\tau^n=\frac{\xi-\alpha}{\xi}\frac{py-\pi-p^mm}{wn}. \tag{E.3}$$

Equations (E.1), (E.2), and (E.3) identify the unobservable output wedge, the prediction error, and the observable input wedge, respectively. To keep our results comparable with the main text, we report the results using a combined ex ante wedge wedge

$$1-\tau_{t,s,i} \equiv \frac{1-\tau^{y}_{t,s,i}}{1+\tau^{n}_{t,s,i}} = \frac{\sigma}{\sigma-1}\frac{w_{t,s,i}n_{t,s,i}}{p_{t,s,i}y_{t,s,i}}\frac{\xi_s\, p^{m}_{t,s,i}m_{t,s,i}}{(\xi_s-\alpha_s)(1-\xi_s)(p_{t,s,i}y_{t,s,i}-\pi_{t,s,i}-p^{m}_{t,s,i}m_{t,s,i})}. \tag{E.4}$$

The empirical measures used and the calibration of the elasticity parameters are as in Section 2.4, where the materials were first introduced. Table 17 gives the results of the extended accounting exercise.

Table 17: Decomposition of ex post misallocation under the extended identification with labor distortions

Variable	Value	Share
$\text{Var}(\ln(1-\tau_{HK}))$	0.320	
$\text{Var}(\ln(1-\varphi))$	1.092	3.41
$\text{Var}(\ln(\frac{1-\tau_y}{1+\tau_n}))$	0.608	1.90
$2\text{Cov}(\ln(1-\varphi), \ln(\frac{1-\tau_y}{1+\tau_n}))$	-1.380	-4.31

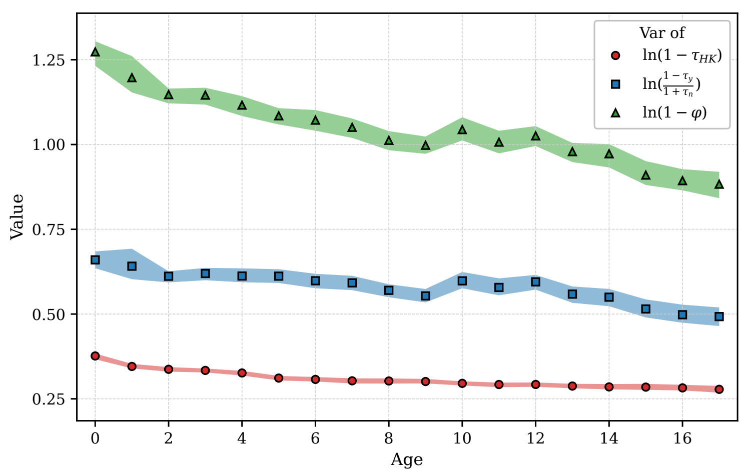

Figure 8: Age profiles under the extended identification with labor distortions

Notes: Shaded areas represent 95% confidence intervals from a block bootstrap with individual firms as blocks.

The results are broadly in line with the alternative approach without labor wedges discussed in Section 2.4: the extended approach yields an elevated level of uncertainty and a strongly negative covariance term. The main difference is the substantial increase in ex ante misallocation, which now captures the contribution of labor wedges in addition to output wedges. Figure 8 shows that the age-dependent patterns are preserved: uncertainty declines sharply over the life cycle while ex ante misallocation remains relatively stable.

Appendix F. Correlated Distortion

In this appendix, we provide the details of the model with correlated distortions discussed in Section 5. We begin by describing the incumbent's problem, and then turn to the calibration and the static decomposition.

F.I. Incumbent's Problem

As in the baseline model, firm productivity depends on a transitory component and a persistent component, which follows an AR(1) process (see Eqs. (15) and (16) in the main text). Moreover, the firm cannot decompose these signals.

The ex ante wedge takes the form

$$\ln(1-\tau) = \tau_p + \tau_z,$$

where τ_p is an observable, normally distributed permanent wedge drawn at entry, and

$$\tau_z = \tau_{tr} + \iota z_{p,-},$$

with $\tau_{tr} \sim N(0, \sigma^2_{\tau_{tr}})$ a transitory wedge and ι linking the wedge to productivity. To prevent full inference of past productivity from the current wedge, the firm observes only τ_z, not its components.

This structure gives rise to sequential Kalman filtering. At the beginning of a period, the firm holds a belief $m_{-|z_-}$ about its persistent productivity in the previous period, based on information up to the realization of z_-. Conditional on continuing, the firm observes the current ex ante wedge and updates this belief according to

$$m_{-|\tau,z_-} = m_{-|z_-} + K_{-|\tau}(\tau_z - \iota m_{-|z_-})$$

with Kalman gain

$$K_{-|\tau} = \frac{\iota \Sigma_{-|z_-}}{\iota^2 \Sigma_{-|z_-} + \sigma^2_{\tau_{tr}}}$$

and updated variance

$$\Sigma_{-|\tau,z_-} = (1 - \iota K_{-|\tau})\Sigma_{-|z_-}$$

Next, the firm uses this information to form a prediction of its current-period persistent productivity and the associated variance of the prediction error:

$$m_{|\tau,z_-} = \rho m_{-|\tau,z_-}$$

$$\Sigma_{|\tau,z_-} = \rho^2 \Sigma_{-|\tau,z_-} + \sigma^2_{\epsilon_p}$$

However, for employment decisions, the firm needs a prediction of its total productivity,

$$z = z_t + z_p = m_{|\tau,z_-} + z_t + (z_p - m_{|\tau,z_-})$$

not just its permanent component. Conditional on $m_{-|\tau,z_-}$ and $\Sigma_{-|\tau,z_-}$, z is normally distributed with

$$\mathrm{E}[z|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}] = \rho m_{-|\tau,z_-} \tag{F.1}$$

$$\mathrm{Var}[z|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}] = \rho^2 \Sigma_{-|\tau,z_-} + \sigma^2_{\epsilon_p} + \sigma^2_{z_{tr}} \tag{F.2}$$

Finally, after production, the firm observes z and updates its belief about the current permanent productivity component, z^p, according to

$$m_{|z} = m_{|\tau,z_-} + K_z(z - m_{|\tau,z_-}),$$

$$K_z = \frac{\Sigma_{|\tau,z_-}}{\Sigma_{|\tau,z_-} + \sigma^2_{z_{tr}}},$$

$$\Sigma_{|z} = (1 - K_z)\Sigma_{|\tau,z_-}.$$

Given these equations, the expected value of $m_{|z}$ before observing current productivity is distributed as

$$\mathrm{E}[m_{|z}|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}] \sim N\Big(\rho m_{-|\tau,z_-}, K_z^2(\rho^2 \Sigma_{-|\tau,z_-} + \sigma^2_{\epsilon_p} + \sigma^2_{z_{tr}})\Big) \tag{F.3}$$

We use this to compute the expected firm value at the beginning of the next period, given information up to the current ex ante wedge.

With this information structure in place, we can set up the Bellman equation for an incumbent firm at the beginning of the period as

$$V(m_{-|z_-}, a, n_-, \tau_p) = \max\Big\{\mathrm{E}_{\tau_z}\big[W(m_{-|\tau,z_-}, a, n_-, \tau_p, \tau_z)\big],\ -2\lambda n_-\Big\}. \tag{F.4}$$

The relevant state variables are the same as in Section 4, except that we now use the belief regarding the previous period's persistent productivity instead of the current period's. However, the value of staying at the stage when employment is chosen, $W(\cdot)$, depends on the still unknown wedge τ_z.

To compute the expectation, note that conditional on $m_{-|z_-}$ and $\Sigma_{-|z_-}$, τ_z is normally distributed with

$$\mathrm{E}[\tau_z | m_{-|z_-}, \Sigma_{-|z_-}] = \iota m_{-|z_-},$$

and variance

$$\mathrm{Var}[\tau_z | m_{-|z_-}, \Sigma_{-|z_-}] = \sigma^2_{\tau tr} + \iota^2 \Sigma_{-|z_-}.$$

The value of staying is given by

$$\begin{aligned} W(m_{-|\tau,z_-}, a, n_-, \tau_p, \tau_z) = \max_n\ & C^{\frac{1}{\sigma}} P e^{\tau_p+\tau_z} \mathrm{E}_{z|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}}\big[(e^z)^{\frac{\sigma-1}{\sigma}}\big] n^{\frac{\sigma-1}{\sigma}} \\ & - n - c_f - \lambda\Big(\frac{n-n_-}{\bar{n}}\Big)^2 \bar{n} \\ & + \beta(1-\gamma)\,\mathrm{E}_{m_{|z}|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}}\big[V(m_{|z}, a+1, n, \tau_p)\big]. \end{aligned} \tag{F.5}$$

The first expectation, $\mathrm{E}_{z|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}}[\cdot]$, is computed using the fact that z, conditional on information available up to observing the ex ante wedge, is normally distributed with parameters given by equations (F.1) and (F.2). The second expectation, $\mathrm{E}_{m_{|z}|m_{-|\tau,z_-}, \Sigma_{-|\tau,z_-}}[\cdot]$, uses the distribution specified in equation (F.3).

F.II. Calibration

To discipline the calibration of the model with more flexible ex ante wedges, we apply a two-stage calibration strategy. In the first stage, we fix the parameters governing the productivity process, and thus the information friction, using information directly related to the evolution of measured productivities and the relationship between productivity and firm exits. In the second stage, we calibrate the remaining parameters by targeting the same moments as in Section 4.

To proceed, we first measure firms' productivity. Following HK, we evaluate it as

$$z_{t,s,i} \propto \frac{(p_{t,s,i}y_{t,s,i})^{\frac{\sigma}{\sigma-1}}}{n_{t,s,i}^{1-\alpha_s}k_{t,s,i}^{\alpha_s}}, \tag{F.6}$$

where, similarly to Section 2, $\sigma = 3$, value added is used as a measure $p_{t,s,i}y_{t,s,i}$, and labor shares at the 3-digit industry level are used to pin down α_s.

Without endogenous exits, we could identify the parameters governing the evolution of productivity directly from the autocovariances of lags 0, 1, and 2,

$$\gamma_0 = \frac{\sigma_\varepsilon^2}{1-\rho^2} + \sigma_{z_{tr}}^2, \quad \gamma_1 = \rho\frac{\sigma_\varepsilon^2}{1-\rho^2}, \quad \gamma_2 = \rho^2\frac{\sigma_\varepsilon^2}{1-\rho^2},$$

where γ_h denotes the autocovariance of the productivity process at lag h. However, the fact that firms with lower productivity are more likely to exit complicates the calibration.

To address the selection issue, we follow Guvenen et al. (2021) and assume that firm exits follow a logistic probability model. Specifically, each firm faces an exit shock with probability $\nu_{a,z}$, which depends on the firm's age, a, and productivity, z, through a logistic function:

$$\nu_{a,z} = \frac{e^{\xi_{a,z}}}{1+e^{\xi_{a,z}}}, \quad \xi_{a,z} = b_0 + b_1 a + b_2 z + b_3 az.$$

We estimate the parameters of the productivity process (ρ, σ_ε^2, and $\sigma_{z_{tr}}^2$) jointly with those governing firm exits (b_0, b_1, b_2, and b_3) using the method of simulated moments. Given a set of parameter values, we simulate a cohort of firms and compare the model-implied moments to their empirical counterparts computed from the data. The target moments include the relevant autocovariances, exit rates across different age brackets, and the correlation between exits and productivity conditional on age. The empirical targets are calculated using only the three cohorts of firms that can be tracked up to age 15 (those entering between 1995 and 1997).

The upper panel of Table 18 reports the preferred parameterization obtained using the method of simulated moments ("SMM"). For comparison, we also report the parameter values for the productivity process implied by the empirical autocovariance function when firm exits are ignored ("Direct I" and "Direct II"). For Direct I, we use the whole pooled data, while for Direct II, similarly to our SMM, we only use cohorts that started in 1995–1997. The difference between column Direct II and SMM illustrates the impact of exits on the estimated parameters. The lower panel of the table presents the targeted moments alongside their model counterparts.

Table 18: Calibration of the Productivity Process

	Direct I	Direct II	SMM
Parameters:			
ρ	0.94	0.92	0.98
σ_{ε}	0.25	0.27	0.14
$\sigma_{z_{tr}}$	0.41	0.43	0.48
b_0 (const)			-2.74
b_1 (age)			-0.13
b_2 (z)			-0.42
b_3 ($z\times$ age)			-0.01
Target Moments for SMM (1995–1997 cohorts):			
		Model	Data
Var(z)		0.69	0.69
Cov(z_{t+1}; z_t)		0.45	0.46
Cov(z_{t+2}; z_t)		0.44	0.43
Exit, age 0–2		0.06	0.07
Exit, age 3–5		0.05	0.04
Exit, age 6–10		0.03	0.03
Exit, age 11–15		0.02	0.02
Corr(Exit, age 0–2; z_t)		-0.09	-0.08
Corr(Exit, age 3–5; z_t)		-0.08	-0.10
Corr(Exit, age 6–10; z_t)		-0.07	-0.07
Corr(Exit, age 11–15; z_t)		-0.06	-0.05

Notes: Direct I uses autocovariances of lags 0, 1, 2 calculated from the pooled data to estimate the parameters of the productivity process. Direct II uses the cohorts 1995-1997 for the same analysis. Finally, column SMM uses cohorts 1995-1997 and takes into account the selection by specifying a logistic function that determines probability of exiting the sample, simulates one cohort for 15 years and finds the parameter values that minimize the relative distance between simulated and empirical moments.

In the second stage, we calibrate the remaining model parameters by targeting the same growth rate moments, exit rates, employment patterns and size distribution as before. We consider three versions of the model. M1 corresponds to the baseline specification. M2 extends it by introducing temporary wedges but sets $\iota = 0$, while M3 further allows the correlation parameter ι to be estimated from the data. The empirical targets and corresponding model moments are reported in Table 19 and the associated parameter values are presented in Table 20.

Table 19: Target and model counterparts for alternative calibrations

Moment	Data	M1	M2	M3
By Age Group				
Mean growth rate, age 0–2	0.13	0.27	0.22	0.26
Mean growth rate, age 3–5	0.02	0.02	0.01	0.02
Mean growth rate, age 6–10	0.00	0.00	−0.01	0.00
Mean growth rate, age 11–15	−0.01	0.00	−0.01	0.00
Variance of growth rate, age 0–2	0.16	0.34	0.39	0.16
Variance of growth rate, age 3–5	0.10	0.05	0.06	0.03
Variance of growth rate, age 6–10	0.07	0.03	0.04	0.03
Variance of growth rate, age 11–15	0.07	0.03	0.04	0.03
Exit rate, age 0–2	0.05	0.05	0.04	0.04
Exit rate, age 3–5	0.03	0.03	0.03	0.02
Exit rate, age 6–10	0.02	0.02	0.02	0.02
Exit rate, age 11–15	0.02	0.02	0.02	0.02
Autocorrelation				
First-order autocorrelation of log employment	0.96	0.98	0.97	0.96
Second-order autocorrelation of log employment	0.94	0.96	0.94	0.93
By Employees Group				
Share of firms with $1 \leq$ employees < 5	0.57	0.53	0.57	0.58
Share of firms with $5 \leq$ employees < 10	0.20	0.16	0.20	0.23
Share of firms with $10 \leq$ employees < 20	0.11	0.11	0.11	0.11
Share of firms with $20 \leq$ employees < 50	0.07	0.11	0.08	0.07
Share of firms with employees ≥ 50	0.05	0.10	0.03	0.02

Notes: M1 refers to the original model under the alternative calibration strategy; M2 adds temporary wedges; and M3 additionally allows ex ante wedges to be correlated with productivity. Moments shown with a gray background are not used as calibration targets.

F.III. Static Misallocation for Alternative Models

Table 21 presents the decomposition of static misallocation based on the calibrated models M1, M2, and M3. As in our baseline model, M1 assigns roughly equal importance to variation in prediction errors and variation in ex ante wedges, consistent with the empirical results in Section 2. However, the levels of both variances are higher than in the data, which is consis-

Table 20: Parameter values for alternative calibrations

Description	Parameter	M1	M2	M3
Entry cost	c_e	126.936	44.284	41.963
Fixed cost	c_f	0.336	0.617	0.618
Tax-like wedges (permanent)	σ_{τ_p}	0.592	0.193	0.186
Tax-like wedges (transitory)	$\sigma_{\tau_{tr}}$		0.050	0.049
Adjustment costs	λ	0.261	0.190	0.190
Exogenous exit	γ	0.017	0.015	0.015
Shock correlation	ι			-0.102

Notes: M1 refers to the original model under the alternative calibration strategy; M2 adds temporary wedges; and M3 additionally allows ex ante wedges to be correlated with productivity.

Table 21: Ex post misallocation and its decomposition for alternative calibrations

Variable	M1	M2	M3
$\mathrm{Var}(\ln(1-\tau_{HK}))$	0.36	0.18	0.17
$\mathrm{Var}(\ln(1-\tau))$	0.22	0.04	0.05
$\mathrm{Var}(\ln(1-\varphi))$	0.16	0.16	0.13
$2\mathrm{Cov}(\ln(1-\tau), \ln(1-\varphi))$	-0.03	-0.02	-0.01

Notes: M1 refers to the original model under the alternative calibration strategy; M2 adds temporary wedges; and M3 additionally allows ex ante wedges to be correlated with productivity.

tent with the higher ex post misallocation in M1. For M2, by contrast, ex post misallocation is close to its empirical counterpart. However, the alternative calibration strategy appears to place greater emphasis on the uncertainty channel relative to the baseline calibration or the empirical decomposition.M3 is closest among the three alternative calibrations, although it still puts too much weight on uncertainty and too little on ex ante misallocation.