Maintenance Volatility, Firm Productivity, and the User Cost of Capital*

Poorya Kabir† and Eugene Tan‡

July 19, 2024

Abstract

We document that firms with higher volatility in maintenance rates exhibit lower performance in terms of lower revenue, productivity, and smaller capital stocks, but also have higher average revenue products of capital. We rationalize this result using a parsimonious model that incorporates endogenous maintenance decisions that are subject to maintenance cost shocks. Maintenance is necessary for output, but decreasing returns to maintenance imply that volatile maintenance services lead to lower average productivity and higher average user costs. We calibrate the model to firm-level data from India, and show that eliminating maintenance volatility can increase GDP by 4.1%. Lastly, we provide suggestive evidence higher maintenance volatility is associated with tighter financial constraints and firms that operate older capital stocks, providing new avenue for research into new channels to improve firm productivity.

Keywords: Productivity, Maintenance, Investment

*We would like to thank Andrea Lanteri for insightful discussions.
†National University of Singapore: poorya.kabir@nus.edu.sg
‡University of Toronto: eugene.tan@rotman.utoronto.ca

1 Introduction

Machines cannot be used optimally if they are not properly maintained. This suggests that maintenance decisions have a direct impact on firm performance. Perhaps reflecting the important role of maintenance, maintenance expenses by businesses are large. Maintenance expenses are about 6% of GDP, similar to spending on education and about 28% of spending on new physical capital (McGrattan and Schmitz, 1999).[1] Likewise, earlier research on best management practices (e.g., Bloom et al. (2012)) have documented an important role for maintenance. Yet, while there is a vast literature studying the role of investment dynamics on firm growth, there is sparse empirical evidence connecting maintenance decisions to firm productivity. Likewise, there is no theoretical framework connecting maintenance decisions to firm outcomes and aggregate productivity. Our article aims to cover both gaps in this literature.

Our main hypothesis is that firms which exhibit higher *volatility* in maintenance would exhibit lower long-term average performance, in the form of lower revenue and productivity. Our argument is based on the logic that while maintenance is necessary for maintaining the quality of machinery, maintenance benefits suffer decreasing returns. As a result, a constant and stable maintenance schedule is preferable to a volatile one that deviates from the schedule. A simple example is the case of car maintenance. A typical manufacturer's suggested maintenance schedule is bi-annual maintenance and servicing over the lifetime of the vehicle. Conversely, the same total amount of spending compacted into a single year is unlikely to deliver the same benefits.

With that in mind, we provide evidence that, controlling for a battery of confounding variables, firms with higher maintenance volatility have lower value added, lower measured total factor productivity, and have smaller capital stocks. For instance, we find that a one unit increase in the standard deviation of volatility of maintenance is associated with a 1.67% decrease in average value added. Then, we present a model of endogenous maintenance decisions to match this empirical fact. By calibrating the model to the data, we find from our baseline model that about 3.89% of the total dispersion in measured firm productivity (i.e., the firm-level Solow residual) is attributable to dispersion in maintenance. In turn, we find that a complete elimination of maintenance volatility can improve aggregate GDP by almost 4.1%, and lead to aggregate TFP gains of around 0.8%, suggesting that addressing the fundamental source of maintenance volatility could be an important channel for improving aggregate outcomes in India.

[1] In India, our country of analysis, the average maintenance spending is around 6.3% of the capital stock, comparable to average investment spending of around 24.6% of the capital stock.

Our data source for analysis comes from firm-level panel accounting data for companies in India. The data includes information on capital maintenance expenses and standard balance sheet and income statement items. This allows us to examine the relationship between firm performance and capital maintenance. We use this data to establish two reduced-form relationships that support our hypothesis.

First, we show, following our hypothesis, that maintenance is associated with better firm performance, but that it exhibits decreasing returns. Using a simple OLS analysis, we find that a 1% increase in the amount of maintenance spending per unit of capital (which we term "maintenance rate") is associated with a 0.23% increase in measured firm total factor productivity (TFP). Similarly, using an IV approach that addresses endogeneity concerns in maintenance increases this to between 0.43% to 0.47%, depending on the IV.

Second, we find that increases in the within-firm volatility of (log) maintenance rates are associated with decreases in firm performance. For instance, we find that holding fixed average maintenance spending, a one unit increase in the volatility of maintenance rates imply a 1.56% decrease in sales, 0.53% decrease in measured firm TFP, and 1.76% decrease in capital stock; all results consistent with our hypothesis that volatility in maintenance detrimental to firm performance. Interestingly, we also observe that increases in volatility is associated with *higher* average revenue product of capital (ARPK); suggesting that firms which face higher maintenance volatility perceive a higher user or shadow cost of capital.

To quantify the importance of maintenance volatility for firm performance, we next turn to a structural model of endogenous maintenance and investment decisions. Our model is a standard neoclassical investment model with productivity shocks and capital adjustment costs. To that, we augment the model with maintenance costs shocks that shift the marginal cost of maintenance, and extend the production function to allow for maintenance as a necessary input. Importantly, the firm "Solow residual"–the typical measure of firm TFP–is now a function of both exogenous productivity shocks and endogenous maintenance inputs. Furthermore, we allow for firm "fixed-effects" in both long-run productivity and maintenance volatility, as well as arbitrary correlation between the two. The parsimonious nature of the model allows us to match salient moments of the data, without taking an explicit stand on the source of maintenance volatility.

From the calibration, we find that maintenance volatility alone explains about 2.88% of the variation in long-run measured firm TFP, with the residual correlation between productivity and maintenance volatility resulting from exogenous correlation; that is, firms with low maintenance volatility also happen to be firms that have higher exogenous productivity. Finally, we find that maintenance volatility has a large impact on the dispersion

of ARPK, amplifying ARPK dispersion by almost one-third relative to an economy without maintenance shocks. This arises because maintenance itself is a part of the user cost of capital, and larger volatility in maintenance naturally amplify the dispersion of ARPK.

We next analyze, using our structural model, the economic impact of maintenance volatility on firm and aggregate outcomes. Our headline counterfactual analysis simply mutes maintenance volatility entirely, thus reducing our model to the standard investment model. We find that, absent maintenance shocks, firm values and TFP improve by about 3.8% and 2.1% respectively. Similarly, consistent with our empirical results, we find that muting maintenance shocks lead to a 5.5% decline in average firm ARPK, implying that maintenance volatility has a substantial impact on the aggregate user cost of capital.

At the aggregate level, we find that muting maintenance shocks lead to a 0.8% increase in aggregate TFP and 4.1% increase in GDP. Further decomposition reveals that these gains at the aggregate level comes through two unique channels: A direct channel coming through improvements in firm-level productivity, as we already discussed, but also an indirect channel through a decline in the dispersion of ARPK. This latter result occurs because a large fraction of the dispersion in ARPK arises due to shocks to the total user cost of capital, which includes maintenance. By muting maintenance shocks, we also indirectly mute its impact on ARPK dispersion, leading improvements to aggregate productivity. This result is reminiscent of the literature on capital misallocation.

A natural concern regarding our analysis is that some degree of maintenance expenses are deterministic and associated with increasing capital age. For instance, an older car generally needs a larger amount of maintenance expenses to achieve the same efficiency as a newer car. Therefore, our baseline analysis might overstate the "cost" of maintenance volatility to the economy, since some fraction of this volatility is intrinsic and cannot be resolved.

We address this issue by extending our baseline model to allow for exogenous aging of the capital stock, where maintenance costs rises exogenously as the capital ages, and firms make endogenous decisions as to when they upgrade their capital stock. When we re-calibrate the model to account for capital age, we find that maintenance shocks exhibit close to the same volatility, but exhibit a larger degree of decreasing returns. As a result, we find that muting maintenance volatility brings about larger improvements to firm performance and aggregate outcomes. For instance, firm values and TFP are around 4.7% and 2.1% higher, while GDP and aggregate TFP are around 4.8% and 1.8% higher.

Our analyses so far have taken maintenance volatility to be a model primitive, which leads to an open question as to what drives differences in volatility, and how firm managers or policymakers could even address this issue. To provide further insight, we

present evidence of two economic forces that could drive maintenance volatility.

First, we show evidence that firms that face tighter financial constraints exhibit higher maintenance volatility. For instance, we observe that firms with high leverage exhibit about 0.026 higher standard deviation of maintenance rates. Our hypothesis here is that firms trade-off between maintaining the capital stock today against delaying maintenance in favor of making payments towards other costs of production. As a result, firms that are more financially constrained are also more likely to delay maintenance until they can no longer do so. We verify this hypothesis using our data, showing that firms which experience tighter financial constraints reduce maintenance spending in the current period, and delay spending into the future. For instance, we find that a 1% increase in the debt-to-asset ratio reduces maintenance rates by around 3.0% in the current period, but increases spending by 3.6% three years down the road.

Second, we show indirect evidence that maintenance volatility itself might be a result of the vintage of the capital stock (as opposed to just the level of spending, as in our extended model). We hypothesize that new capital might be associated with less unexpected maintenance requirements. For instance, following our earlier analogy, a new car is less likely to require unexpected maintenance relative to a used car. We test this hypothesis using the insights of Cooper et al. (1999), who argue that lumpy investment is associated with capital upgrading. Indeed, we find that following an investment "spike" (Doms and Dunne (1998)), maintenance volatility decreases by around 0.053.

The rest of the article is structured as follows. Following this section, we present our literature review. Section 2 presents our data for analysis, and section 3 presents our headline reduced-form results. Section 4 presents our model, while section 5 presents the calibration approach. Section 6 presents our quantitative results, including our extended model that includes capital age. Finally, section 7 presents our additional empirical analyses, and section 8 concludes.

A Related Literature

Our article is related to three broad strands of the macroeconomics and corporate finance literature on firm and aggregate productivity.

First, our article is related to recent evidence by Sterk et al. (2021) showing that ex-ante "permanent differences" across firms is the main driver of the firm size distribution. However, it remains an open question as to the source of these permanent differences, with most extant work in macroeconomics focusing on firm innovation and intangible capital (such as customer capital) accumulation (e.g., Gourio and Rudanko (2014); Ace-

moglu et al. (2018); Argente et al. (2024)), while research in corporate finance focusing on some dimension of founder characteristics (e.g., Smith et al. (2019); Becker and Hvide (2021); Guzman and Li (2023)). Our article is the first to focus on maintenance spending as a driver of the firm-size distribution, showing that maintenance volatility alone has a direct and quantitatively meaningful impact on long-run firm performance, and in turn a quantitatively important impact on aggregate productivity. Furthermore, unlike the preceding research that emphasizes *productivity* alone as a driver of the long-run firm size distribution, our article also emphasizes that maintenance volatility additionally shapes the firm-size distribution through a secondary channel via the user cost of capital.

Second, our article also provides evidence connecting financial frictions and capital vintages to maintenance volatility, and which, through the lens of our model, impact firm productivity. While financial frictions and capital vintages alone have been a large focus of the extant research on firm and aggregate productivity, the focus of the extant research emphasizes a direct connection between the aforementioned mechanisms and productivity. For instance, articles such as those by Greenwood et al. (2010), Gorodnichenko and Schnitzer (2013), and Ottonello and Winberry (2024) focus on financial frictions and their impact on innovation at the firm-level, while research by Midrigan and Xu (2014) and Gopinath et al. (2017) emphasize the role of financial frictions in lowering aggregate allocative efficiency. Likewise, articles such as those by Caunedo and Keller (2021) focus on the role of capital-embodied technological process. Our article provides evidence regarding potential *additional* benefits to reducing financial frictions and improving new capital adoption, on top of the standard benefits that has been studied by the current literature.

Finally, while not a focus of our article, our findings are also related to the literature that emphasizes the role of "correlated distortions" in driving capital misallocation. Specifically, extant research has uncovered a pervasive fact across countries that firm-level ARPK is highly auto-correlated and positively correlated with productivity, and that such "correlated distortions" appear to be the pre-dominant driver of capital misallocation (e.g., Bartelsman et al. (2013); Restuccia and Rogerson (2008); David and Venkateswaran (2019)). Our mechanism, whereby the user cost of capital fluctuates with productivity due to endogenous maintenance decisions, provides one potential micro-foundation for correlated distortions, and adds to the small but growing literature (e.g., David et al. (2022)) that attempts to uncover the fundamental sources of correlated distortions.

2 Data and Measurement

In this section, we present our data source, measurement strategy, and report descriptive statistics.

A Data

We use data from Prowess, a database collected by the Center for Monitoring Indian Economy (CMIE). The dataset covers an unbalanced panel of firms from 1989 to 2019 and includes the typical income statement and balance sheet data available in other firm-level datasets such as COMPUSTAT. This dataset is uniquely suited for this paper because it includes information on the repair and maintenance expenses for a large sample of firms over time.

For our analysis, we keep firms with the fiscal month ending on March 31 (more than 95% of firm-year observations) to ensure sample consistency. We also require firms to have positive and non-missing capital, sales, assets, and maintenance. To reduce the effect of outliers in our analysis, we winsorize all variables at a 1% level. Throughout the paper, we use the term "industry" to refer to 6-digit NIC codes unless otherwise stated.[2] To have variation to estimate firm-level maintenance volatility, our sample of analysis only includes firms with more than 10 years of non-missing data in Prowess. Finally, we use GDP deflator data from the World Bank database to convert variables to 2019 values and use the average daily exchange rate of INR to USD in year 2019 (USD = 70.51 INR) to convert INR to USD. Next, we discuss how we measure our variables of interest.

B Measurement

Our analysis in the following sections will require measuring firm TFP, capital stock, investment, and maintenance. We briefly summarize our measurement strategy in this section and present the relevant summary statistics in Table 1.

B.1 Capital, Investment, ARPK, and TFP

We use data on sales and capital inputs to recover firm-level total factor productivity (TFP), average revenue products of capital (ARPK), and investment. Value added (y) is the difference between sales and raw materials cost. Capital (k) is measured as the book value of "Plant, Property, and Equipment" (i.e., PPE). Log ARPK is the difference in log

[2] Our results are robust to alternative industry definitions.

value-added and log capital. For firm TFP (ω), we use

$$\log \omega = (1 - (1 - \alpha)\nu) \left(\log y - \frac{\alpha\nu}{1 - (1 - \alpha)\nu} \log k \right) \tag{1}$$

and impose that $\alpha = 1/3$ and $\nu = \frac{3}{4}$ following David and Venkateswaran (2019). As discussed in their article, this formula for computing TFP is consistent with a firm that operates a Cobb-Douglas production function with capital intensity parameter α and decreasing returns to scale ν, and faces a static decision problem in hiring labor.[3]

Finally, we impute investment using a perpetual inventory method with a depreciation rate of 10%.[4]

B.2 Maintenance

Repair and maintenance expenses are reported in Prowess's income statements in line item "repair and maintenance". For brevity, we refer to this item as maintenance expenses and define maintenance rates as maintenance expenses divided by the capital stock ($p\tilde{m} = \frac{pm}{k}$), where m is the quantity of maintenance services and p is the price of maintenance services, and hence pm is the expenditure on maintenance services. Our focus on maintenance rates (rather than levels) follows McGrattan and Schmitz (1999) and Caunedo and Keller (2021), for which the benefit of maintenance scales per unit of capital, rather than simply the absolute level of spending.

Maintenance expenditures in accounting are categorized into two types: expensed and capitalized. Expenditures expensed are recognized in the income statement in the period in which they are incurred. This approach is commonly applied to routine maintenance activities that neither enhance the asset's value nor prolong its useful life. Examples of such activities include regular servicing, minor repairs, and replacing small parts.

In contrast, capitalizing maintenance costs involves treating these expenses as additions to the value of an asset, consequently reflecting them under the Property, Plant, and Equipment (PPE). This method is typically employed for maintenance activities that substantially increase the asset's value, improve its performance, or extend its useful life. Activities falling under this category include major overhauls, significant upgrades, and enhancements that boost the asset's efficiency or capacity.

[3]The authors argue that this formula is preferable to standard production function estimation approaches when labor decisions are distorted, and labor distortions are not explicitly accounted for in the structural model, as is our case. Regardless, in Appendix Table OA.2, we present our headline results using TFP computed using estimated production functions (Levinsohn and Petrin, 2003; Ackerberg et al., 2015) and find similar results.

[4]The median accounting depreciation rate in our sample is 9.79%.

Our analysis focuses on the "repair and maintenance" line item from the income statement. This item accounts for maintenance expenditures when they are expensed. It is most relevant to our study as it aligns with our conceptualization of maintenance.

C Descriptive Summary Statistics

Table 1 reports the summary statistics for our sample. The first and second rows report statistics for the capital stock and value-added. The typical firm is relatively large, with median capital stock and value-added of US$5.1 million and US$9.5 million, respectively.

Rows three and four report gross investment and investment rates. Gross investment is defined as $i_t = PPE_t - (1-\delta)PPE_{t-1}$, where we choose $\delta = 10\%$ for all firms. Like other firm-level datasets, investment features high dispersion and a large excess kurtosis in the right tail. The next two rows show summary statistics for TFP and ARPK.

Summary statistics for maintenance are reported in the seventh to ninth rows of Table 1. As we discussed in the introduction, maintenance spending is large; on average, capital maintenance rates are about 6.3%, comparable to capital depreciation rates which are typically assumed to be between 6% to 10%. This suggests that maintenance, as a literal user cost of capital, is on the same magnitude as conventional estimates for capital user cost. More importantly, maintenance also exhibits large variability, suggesting substantial volatility in capital user costs.

3 Evidence

In this section, we present evidence that (i) maintenance volatility differences across firms are associated with differing firm performance, and (ii) that firm productivity improves with maintenance spending, but the benefits are concave.

A Maintenance Volatility and Firm Performance

To construct the maintenance volatility ($\sigma_i(\log(p\tilde{m}))$ and mean maintenance ($\mu_i(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ regressing it on its lagged value and industry-by-year fixed effects, then calculate the volatility and mean of the residuals across all observations for a firm. In constructing the residuals, we control for the lagged value of $\log(p\tilde{m})$ to align our empirical work with Section 4's model of the maintenance volatility process, which we assume is an AR(1). Residualizing maintenance relative to its lagged value has the added advantage that we account for a potential trend in maintenance. For instance,

a firm that only periodically upgrades its capital stock might see a natural increase in maintenance spending as maintenance requirements increases with the age of the capital stock. Our analysis focuses on deviations from this presumably "optimal" maintenance schedule. As such, residualizing with an AR(1) process allows us to interpret the dispersion of the residuals as "shocks" that lead to a departure from this schedule.[5] In contrast, for logged values of value added, ω, k, and ARPK, we simply residualize each variable using industry-by-year fixed effects and then take the mean of the residuals across all observations for a firm.

Table 2 shows the relationship between maintenance volatility and firm performance measures while controlling for average maintenance differences across firms. Column 1 shows a statistically significant and negative relationship between maintenance volatility and value added. The effect is economically significant, as well; a 1 standard deviation increase in maintenance volatility decreases value added by 28% of standard deviation. Column 2 shows that a 1 standard deviation increase in maintenance volatility decreases productivity by 22% standard deviation. Column (3) shows that a 1 standard deviation increase in maintenance volatility decreases average capital by 22% standard deviation.

Finally, we observe in Column (4) that an increase in maintenance volatility is associated with an increase in average firm-level ARPK. In particular, this is true even after controlling for average maintenance expenses. In other words, it appears that higher maintenance volatility increases the shadow cost of capital, suggesting that to some degree, the negative relationship associated with sales and capital is driven by both lower productivity and higher shadow cost of capital resulting from maintenance volatility. Figure OA.7 in the Appendix plots a binscatter plot for each of the outcome variables, further confirming that the findings are not driven by outliers

While Table 2 provides evidence consistent with our hypothesis, one potential concern with our interpretation is that the differences in firm outcomes, such as value added, could driven by differences in average and variance of productivity. In Appendix Table OA.1, we show that we find qualitatively similar results when controlling for differences in average and variance of productivity, suggesting that the findings can not be fully explained by differences in productivity average and variance.

In Figure OA.3 in the Appendix, we assess the sensitivity of our estimates to different assumptions about the industries. In Panel (A), we leave one NIC 6-digit industry out at a time and re-estimate regressions identical to the ones used in Table 2, and present the estimated βs in a boxplot for each outcome variable. We find that the estimated βs

[5]While we use AR(1) in our main specification, the main results are robust to using the AR(2) process, documented in Table OA.3.

vary very little, suggesting that the β estimated in Table 2 are not driven by one specific industry. In Panel (B), we estimate regressions for one NIC 6-digit industry at a time. We find that most of the single industry estimates have a similar sign to our main findings. This is despite the very narrow definition of industry, whereby the median industry in the sample used for estimation has a total of 14 firm-year observations.

B The effect of maintenance on productivity

One key assumption underlying our model is that higher maintenance increases firm-level productivity. To test this hypothesis, we use the following regression specification:

$$\log(\omega_{it}) = \alpha_i + \alpha_t + \beta \log(p_{i,t}\tilde{m}_{i,t}) + \varepsilon_{it}, \tag{2}$$

where α_i and α_t are firm and year fixed effects, β is the coefficient of interest and measures the effect of maintenance rate on firm productivity.

Column 1 of Table 3 shows β estimated using OLS. The positive and statistically significant coefficient estimate suggests that increased maintenance increases firm-level productivity. While suggestive, the OLS estimate suffers from the unobserved bias. For instance, if an increase in labor supply lowers the cost of labor, then firms might become more productive and also use more labor to maintain their capital, but the increase in productivity doesn't come from increased maintenance. Next, we describe how we address this challenge by using an IV strategy.

For constructing our IV, we utilize a "Hausman instrument" approach (Hausman et al., 1994; Hausman, 1996). Our instrument is the average of maintenance rate ($\log(p\tilde{m})$) in the firm's industry-year cell, excluding the firm's own maintenance rate, as well as potentially an expanded set of associated firms that are "similar" to the firm itself. The underlying assumption for this identification strategy is that maintenance rates across different firms within the same industry are correlated, but the omitted factors are not. We consider three "Hausman" IVs in our empirical analysis. Our first instrument excludes only the firm itself, while our second instrument excludes all other firms within the firm's district. Our last instrument excludes all other firms within the firm's state.[6] We consider a widening set of excluded firms because while a violation of the exclusion restriction can happen if a change in the omitted factors of other firms is also associated with the focal firm, this issue is less likely for firms that are geographically distant from the focal firm.

[6]Firms in our dataset are in 298 districts and 32 states.

For brevity, we only explain the details for "leave out focal firm's state". The rest follow the same logic. Let $\log(p_{i,j,t}\tilde{m}_{i,j,t})$ denote the maintenance rate of firm i in year t operating in industry j. First, we regress $\log(p_{i,j,t}\tilde{m}_{i,j,t})$ on firm and year FEs and take the residuals to construct the IV:

$$IV_{k,j,t} = \sum_{i \notin S_k} \log(p_{i,j,t}\tilde{m}_{i,j,t}) \tag{3}$$

where k indexes the focal firm, and S_k denotes the set of firms operating in the same state as the focal firm. The underlying assumption behind the instrument is that demand across firms operating in the same narrowly defined industry and year is correlated.

The relevance condition for the IV relies on firms in narrowly defined industries facing correlated maintenance in the same year, for instance, because of a common change in the price of spare parts. Table OA.4 in the Appendix provides the first-stage estimates and F-statistics, suggesting the relevance condition holds. The exclusion restriction for the IV assumes that the only channel through which changes in other firm's maintenance rate affect the focal firm's productivity is through changes in the focal firm's maintenance decisions. A violation of the exclusion restriction can happen if, for instance, an increase in labor supply can lead to an increase in the maintenance of other firms by lowering the cost of maintenance. The same increase in labor supply could make the focal firm more productive, not because of increased maintenance but simply because labor is cheaper. If this is the case, we expect the effect to weaken substantially as we increase the excluded set of firms since labor markets are local, which is not empirically supported in the data, as we explain next.

Columns (2) to (4) of Table 3 report our results for the three different IVs. We find estimates for β that are substantially larger than zero, ranging from 0.43 to 0.47. Furthermore, estimates are similar when excluding geographic units with differing coverage, suggesting that the exclusion restriction is likely to hold.

In Appendix Figure OA.6, we show that lagged maintenance leads to future fixed asset impairments. Prowess has a line item that documents asset impairments. We convert it to a dummy variable equal to one for positive asset impairments. In Panels (A) and (B), we show OLS and IV estimates of β, using a specification similar to 2, respectively. We find that lower maintenance leads to higher chances of asset impairment in the future, both in OLS and IV results.

4 Model

In this section, we present our model of firm investment with endogenous capital maintenance subject to maintenance shocks.

A Firms

We begin with the firms' decision problem, which is the focus of our article.

Time is discrete and infinite. Firms live forever, face perfect competition, and discount future profits with a real interest rate of r. Output is produced by the following decreasing return to scale production function

$$y_{it} = \left(1 - \exp\left(-\chi \tilde{m}_{it}\right)\right) a_i z_{it} \left(k_{it}^{\alpha} l_{it}^{1-\alpha}\right)^{\nu}. \tag{4}$$

i indexes a firm and t indexes time. a and z both refer to physical productivity of the firm; a is unique to the firm and invariant over time (i.e., a "fixed effect"), while z evolves stochastically following an AR(1) process in logs given by

$$\log z_{it} = \rho_z \log z_{i,t-1} + \sigma^z \epsilon_{it}, \tag{5}$$

with $\epsilon \sim N(0,1)$. k and l refers to capital and labor inputs respectively with $0 \leq \alpha \leq 1$ the capital intensity parameter and $0 \leq \nu < 1$ the degree of decreasing returns to scale. $\tilde{m}$ refer to "maintenance rate", which is defined as $\tilde{m} \equiv \frac{m}{k}$, where m refers to maintenance services in levels. $\chi \geq 0$ is a parameter that controls the returns to maintenance.

Our formulation of the production function implies that maintenance (i) is a necessary input into production, (ii) that output is strictly increasing in maintenance, and crucially, (iii) the marginal benefit of maintenance is decreasing and bounded at 1 in the limit. This formulation is related to earlier literature in capital maintenance (e.g., McGrattan and Schmitz (1999); Caunedo and Keller (2021); Kabir et al. (2024)) that model maintenance as an input that reduces the depreciation rate of capital.

Our departure from this literature is that we assume maintenance input has a direct immediate input on the "productivity" of capital, and thus affects current output. In contrast, earlier research has typically assumed that the impact on maintenance manifests in the next period as reduced capital stock (i.e., endogenous depreciation). In our case, this means the measured total factor productivity (TFP) of the firm, as measured by the

Solow residual of the production function and denoted by ω, is given by

$$\log \omega \equiv \log a + \log z + \log \left(1 - \exp\left(-\chi \tilde{m}_{it}\right)\right), \tag{6}$$

where we see that measured TFP is directly impacted by maintenance. Furthermore, because the benefit function of maintenance is concave, this implies that higher variations in $\tilde{m}$ will lower the *average* measured productivity of the firm. Moreover, the impact of maintenance variation is increasing in χ. This assumption on the benefit of maintenance endogenously lead to a negative relationship between the volatility of maintenance and measured TFP.

We now turn to a discussion of the factors that affect a firm's decision to purchase maintenance, investment, and hire labor.

A.1 Maintenance Costs

We assume that for m units of maintenance services purchased for a capital stock of k, the total cost of maintenance is given by $p_{it}\tilde{m}k$ units of the output good. Furthermore, we assume that the relative price p_{it} is subject to random and persistent shocks. Specifically, we assume that it evolves according to the following AR(1) equation

$$\log p_{it} = (1 - \rho_m) \log \bar{p} + \rho_m \log p_{i,t-1} + \sigma_i^m u_{i,t}, \tag{7}$$

where $u \sim N(0,1)$. $\bar{p}$ is a parameter that controls the long-run maintenance cost of a firm, while ρ_m and σ_i^m determine the persistence and volatility respectively. Importantly, we allow for each firm i to face a different volatility in maintenance shocks, which motivates the i subscript for σ^m. Finally, we assume that $\tilde{m}_{it}$ is chosen after both p_{it} and z_{it} are observed.

A.2 Investment and Capital Accumulation

Capital stock at the firm level evolves according to the accumulation equation

$$k_{i,t+1} = (1 - \delta) k_{it} + i_{it}, \tag{8}$$

where i is gross investment and δ is the depreciation rate of capital. The cost of each unit of capital is normalized to 1 unit of output.

Investment is subject to a non-convex adjustment cost. Following Cooper et al. (1999), we assume that a firm with capital stock k conducting positive investment face a fixed

adjustment cost f that is proportional to its capital stock (i.e., the cost, conditioned on adjustment, is fk). However, for firms selling capital, we assume that they do not face any adjustment costs. As is standard, we assume that i is chosen with perfect information over z_{it} and p_{it}, but before $z_{i,t+1}$ and $p_{i,t+1}$ is observed.

A.3 Labor

As is standard, we assume that labor is also chosen after z_{it} and p_{it} is observed. Therefore, for a given wage rate w, the gross profit function (net of labor cost) can be written as

$$\pi\left(z_{it}, m_{it}, k_{it}\right) \equiv y_{it} - w_t l_{it} = A\left(w_t\right)\left(\left(1 - \exp\left(-\chi \tilde{m}_{it}\right)\right) a_i z_{it}\right)^{\Theta_z} \left(k_{it}\right)^{\Theta_k}, \tag{9}$$

where A is a function of wage w, $\Theta_z \equiv \frac{\nu}{1-(1-\alpha)\nu}$, and $\Theta_k \equiv \alpha\Theta_z$.

A.4 Joint distribution of firm-specific productivity and volatility

Our final assumption on the economic environment faced by the firm relates to their individual-specific productivity a_i and volatility σ_i^m. Here, we assume that they are drawn from a jointly log-normal process given by

$$\begin{pmatrix} \log a \\ \sigma^m \end{pmatrix} \sim N \begin{pmatrix} \varsigma_a^2 & \gamma\varsigma_a\varsigma_m \\ \gamma\varsigma_a\varsigma_m & \varsigma_m^2 \end{pmatrix} \tag{10}$$

where γ is the correlation coefficient, and ς_a and ς_m are the standard deviation of a and σ^m respectively.

B Households

We close out our discussion with the household sector, which is presented for the sake of completeness. An infinitely lived representative household ranks streams of consumption and labor effort according to the following utility function:

$$U_0 \equiv \sum_{t=0}^{\infty} \beta^t \log C_t, \tag{11}$$

where C_t is aggregate consumption and $\beta \in (0, 1)$ is the discount factor. We assume that the household labor supply is perfectly inelastic, and we normalize labor to 1.

C Bellman equations

We now summarize the firm's dynamic programming problem. Let $V(p, z, k; a, \sigma^m)$ denote the value function of a firm with productivity z, maintenance price p, physical capital stock k, and firm-specific long-run productivity a and maintenance price volatility σ^m. Let "prime" variables denote next-period variables. The Bellman equation is,

$$\begin{aligned} V(z, p, k; a, \sigma^m) = \max_{i,m} \pi(z, p, k; a, \sigma^m, w) - i - \mathcal{I}(k' > k) f k - p\tilde{m}k + ... \\ ... \frac{1}{1+r} \mathbb{E}\left[V\left(p', z', k'; a, \sigma^m\right) | z, p\right] \qquad s.t. \\ k' = i + (1-\delta) k, \\ \log z' = \rho_z \log z + \sigma^z \epsilon', \\ \log p' = (1-\rho_m) \log \bar{p} + \rho_m \log p + \sigma^m u', \end{aligned}$$

where $\mathcal{I}(k' > k)$ is an indicator function that evaluates to 1 if the argument is true. Note that because we focus on a stationary equilibrium, the risk-free rate is constant and equal to $\frac{1}{\beta} - 1$.

D Equilibrium

A stationary equilibrium in this economy is (i) a set of value and policy functions for the firm (V,$\tilde{m}$,k',l), (ii) a wage rate w, and (iii) a joint distribution over (a, σ^m, z, p, k) such that taking wages as given, (a) the value and policy functions solve the firm's optimization problem, (b) the output and labor markets clear ($\int y d\Lambda = C$, $\int l d\Lambda = 1$), and (c) the distribution Λ is invariant over time.

5 Calibration

We follow standard approach for calibration, by fixing a subset of parameters to prior literature and calibrating internally the remaining parameters. Specifically, we fix $\alpha = 1/3$, $\nu = 3/4$, $r = 0.04$ and $\delta = 0.10$ following standard assumptions (and consistent with our measurement strategy). The remaining parameters are calibrated to identifying moments which we discuss in more detail below. Finally, we conduct a simple decomposition exercise to quantify the contribution of heterogeneity in the volatility of maintenance to permanent heterogeneity in productivity.

A Internally calibrated parameters

Panel A of Table 4 summarizes the values of our internally calibrated parameters. Given the parsimonious nature of our model, most parameters are identified with standard moments. f, our fixed cost parameter, is used to target the fraction of observations with investment rates in excess of 20% (i.e., $|\frac{i}{k}| > 0$), following the approach of Cooper and Haltiwanger (2006). ρ_z, σ^z, ρ_m, and σ^m are used to jointly target the autocorrelation and standard deviation of the within-firm Solow residual and maintenance rates. To be precise, we fit (in logs) an AR(1) model with firm-fixed effects to the empirically measured ω and $p\tilde{m}$, and the AR(1) parameter and standard deviation of the residuals serve as the moments to target. This approach deviates slightly from the conventional approach of simply using these AR(1) parameters as model primitives. This deviation is necessary, since in our model, the stochastic properties of both productivity and maintenance rates are endogenous.

$\bar{p}$ and χ are used to target the average maintenance rate and the regression coefficient of a regression of log maintenance rate on the log average revenue product of capital, which is simply the ratio of output to capital (i.e., y/k). To see this logic, consider the first order condition for maintenance expenses, which is given by

$$p\tilde{m} = \frac{p}{\chi} \log \frac{\chi}{p} + \frac{p}{\chi} \log(A_x z_z^{\Theta} k^{\Theta_k - 1}). \tag{12}$$

Notice that holding all else constant, the average maintenance expenses is directly increasing in p, and thus increases in $\bar{p}$. Importantly, while maintenance *services* ($\tilde{m}$) are decreasing in $\bar{p}$, maintenance *expenses* (the empirically observed variable) is increasing as the average price of maintenance increases. In contrast, increases in χ, holding all else constant, *decreases* maintenance expenses, as it leads to a decrease in maintenance services but does not directly affect the price of maintenance.

Next, building on the above equation, a further rewriting obtains

$$\frac{d \log p\tilde{m}}{d \log \frac{y}{k}} = \frac{1}{\tilde{m}} \left(\frac{1}{\chi + p \left(\frac{y}{k}\right)^{-1}} \right). \tag{13}$$

Therefore, from the above equation, holding all else constant, an increase in χ always leads to a decrease in the average elasticity between maintenance and ARPK. This is because of the direct impact χ has in the equation, but also because increasing χ leads to a decrease in ARPK due to lower user costs of capital. As a result, the elasticity falls. In contrast, we see that changing p has ambiguous effects on the elasticity. Whereas the di-

rect impact of p is that increases in p also lead to also decreases in the average elasticity between maintenance and ARPK, increases in p also lead to decreases in ARPK. As a result, there is no one-to-one "direct" mapping between p and the elasticity, unlike χ. This differential impact allows us to use these two moments to separately identify these two parameters. Panels A and B of Figure 4 presents a comparative statics exercise around our calibrated parameters to provide visual illustration to this identification strategy. Here, we see clearly that increases in χ lead to a decrease in average m/k and beta, along with decrease in ARPK. Conversely, increases in p lead to an increase in m/k and beta, along with an increase in ARPK.

Finally, we set ς_a and ς_m to replicate the standard deviation of $\mu_i(\log(\omega))$ and $\sigma_i(\log(p\tilde{m}))$ respectively; and set γ to match the regression coefficient of a regression of $\mu_i(\log(\omega))$ on $\sigma_i(\log(p\tilde{m}))$. Recall that $\mu_i(\log(\omega))$ and $\sigma_i(\log(p\tilde{m}))$ refer to the within-firm average $\log\omega$ and standard deviation in the residuals of log maintenance rates, following our data construction in generating Table 2. Importantly, we note that while our model endogenously generates a negative correlation between $\mu_i(\log(\omega))$ and $\sigma_i(\log(p\tilde{m}))$, this relationship is purely dependent on the other parameters of the model that are calibrated to other moments. As a result, the residual variation is used to identify γ.

B Model Fit and External Validation

Panel B of Table 4 summarizes our model fit, showing a close fit to the data. We also report three untargeted moments, namely the regression coefficients from Table 2. Our model also predicts a close fit to the data. Importantly, note that the positive relationship between maintenance volatility and ARPK is also replicated. In other words, as in the data, our model predicts that maintenance volatility is associated not only with lower productivity, but also a higher shadow cost of capital.

A main emphasis, and departure of our model from standard framework, is the introduction of maintenance shocks. We therefore further evaluate our model fit regarding the *dynamic* properties of maintenance spending. Specifically, one issue at hand is that the autocorrelation and volatility alone are not sufficient statistics to fully characterize a dynamic process. For instance, as noted by the vast literature on capital adjustment costs (e.g., Cooper and Haltiwanger (2006)), simply calibrating an external shock process to replicate the autocorrelation and volatility of investment would miss the fact that investment is lumpy. A natural question is whether our parsimonious model, without any frictions to maintenance spending, sufficiently captures the dynamics of maintenance.

To evaluate our model fit in this dimension, we borrow the idea of lumpy investment

(Doms and Dunne (1998)) and apply it to maintenance. To operationalize this, in the data, for each firm, we first identify the year with the highest maintenance rate in the entire sample and keep the observations surrounding the event ($t = -4$ to $t = 5$). Second, to ensure that differences in firm decomposition do not drive the estimation, we further require that all observations be non-missing for the sample period around the event. Then, we estimate the following equation

$$y_{it} = \sum_{s=-3}^{s=5} \beta_s \mathbb{1}_s + \alpha_i + \varepsilon_{it}. \tag{14}$$

We include firm-fixed effects, and the omitted group is t = -4. Hence, the point estimates are the change in the variable of interest relative to t = -4. We then repeat the same procedure for a simulated panel of firms in our model.

Figure 1 plots the estimated coefficient, where the solid blue line plots our model's predicted path for maintenance, and the dashed red line with crosses plots the corresponding estimates the data. Note that while the dynamic path of maintenance was not an explicit target for our calibration exercise, it closely replicates the behavior of maintenance around maintenance spikes in the data.

C How much does maintenance contribute to dispersion in firm TFP?

How much does maintenance contribute to differences in measured TFP across firms? In this subsection, we examine this question by focusing only on the statistical contribution of maintenance to variations in firm productivity. Specifically, we utilize the fact that firm TFP can be directly decomposed into two components, namely $\log(1 - \exp(-\chi\tilde{m}))$ and $\log a_i z$, to inform us of the contribution of maintenance to differences in TFP across firms.

In Column (1) of Table 5, we report the total decrease in the variance of average firm productivity when all variation in maintenance is completely muted.[7] We find that, in totality, maintenance contributes to around 3.89% of the variation in productivity.

A natural issue with the first decomposition is that maintenance varies in response to *both* variations to productivity (az) and to marginal costs (p). As a result, across firms, firms with higher permanent component of productivity (a) will also exhibit higher average maintenance. As a result, the dispersion of TFP is amplified by maintenance even absent any shocks, and the first decomposition overstates the contribution of differences in maintenance volatility across firms.

[7] Let σ_b and σ_c represent the standard deviation of productivity in the baseline and counterfactual respectively. We report $100 \times (1 - \frac{\sigma_c^2}{\sigma_b^2})$.

To address this, we next re-evaluate the model by muting all maintenance shocks, that is, by imposing $\sigma_m = 0$ and $\gamma = 0$. As a result, the only driver of maintenance variation is now responses to productivity variation. Column (2) of Table 5 reports our results, where we find that maintenance shocks alone contributes to 2.88% of the total variance in productivity.

Through our calibration exercise, we also observe a strong negative exogenous correlation between maintenance volatility and firm productivity (namely, $\gamma \approx 0.3$). Therefore, a natural question is how much of the "explained" variance by maintenance shocks is simply due to this exogenous correlation, as opposed to the endogenous mechanism we emphasize. To answer this, we re-evaluate the model by only imposing $\gamma = 0$, but keeping the other parameters at their baseline value. Column (3) of Table 5 reports our result, showing that the exogenous correlation alone contributes to only 0.66% of the total variance. Therefore, while the correlation coefficient appears to be large, it does not have a meaningful impact on the total dispersion in firm productivity.

A natural question is whether our statistical decomposition matches our data. To that end, we conduct a similar decomposition exercise using the data as follows. Recalling our exercise in Table 2, we now conduct a sequence of permutation exercises, in the context of a Shapely-Shorrocks decomposition (Shorrocks (2013)), to compute the marginal contribution of maintenance volatility to the total R^2 in the regression model.[8] Column (4) of Table 5 reports our results: On average, maintenance volatility accounts for around 4.34% of the variance. This result is remarkably close to that in Column (2), the closest model counterpart to the data, suggesting that our model is a good reflection of the actual relationship between maintenance volatility and firm productivity.

While our statistical decomposition suggests that maintenance has a meaningful contribution to variation in firm productivity, it remains an open question as to whether this contribution has economic significance. To that end, we now turn to the next section to evaluate this problem using our structural model.

6 Quantitative Results

We use our calibrated model to conduct three counterfactual exercises to evaluate the impact of maintenance volatility on firm performance and aggregate outcomes. Two counterfactuals are motivated by our statistical decomposition from Section 5. In this case, we separately mute the volatility of maintenance to assess the total impact of mainte-

[8] Recall that we consider the standard deviation of TFP and average log maintenance rates as controls in our baseline regression model.

nance volatility, then we only mute the exogenous correlation of maintenance volatility and productivity to assess the role of this dimension of firm productivity.

In practice, only a subset of the economic gains derived from muting all maintenance volatility might be achievable, since some fraction of shocks could be an inherent nature of business and cannot be forecast (e.g., unexpected flood that damages half the equipment in a manufacturing plant). In Section 7, we provide some reduced form evidence that some of these gains might be reasonably realized.

A Firm values and the shadow cost of capital

We first assess the impact of maintenance volatility on firm performance. Maintenance volatility affects firm performance in two related ways. First, it simply reduces the productivity of a firm, thus reducing profitability and firm sizes. Secondly, and unique to maintenance, is that maintenance uncertainty increases the effective shadow cost of capital, which raises the marginal (and thus average) revenue product of capital (ARPK). To see this, consider the case where there are no adjustment costs. As such, the capital decision of a firm is given by

$$k' = \left(\frac{\mathbb{E}\left[(\omega')_z^{\Theta} \,|p, z \right]}{r + \delta + \mathbb{E}\left[p' \tilde{m}' | p, z \right]} \right)^{\frac{1}{1-\alpha_V \Theta_z}}. \tag{15}$$

Because $p\tilde{m}$ is convex in maintenance, increases in the variance of maintenance implies a higher effective user cost of capital, which reduces capital demand. In turn, this lowers firm revenues and profit above and beyond the first order effect of lowered productivity. Importantly, this also leads to *higher* ARPK for firms that face higher volatility in maintenance shocks, which gives rise to the relationship between maintenance volatility and ARPK that our model predicted in Table 4.

Panel A of Table 6 reports statistics of the distribution of firm values under the counterfactuals, relative to the baseline at each point in the distribution. We see substantial gains when volatility of maintenance is muted. Firms on average see an improvement in firm value by about 3.8%, while the median firm sees an improvement in firm value of around 4.0%. In contrast, muting the exogenous negative correlation between firm productivity and maintenance volatility substantially *reduces* firm values. Average firm values decline by 4.0%, while the median firm value decreases by 4.1%.

Why does muting the exogenous correlation reduce firm values? In standard parametrizations of the production function, as in our article, output is convex in firm productivity.

As a result, the marginal benefit of increasing productivity for a high-productivity firm is higher than that for a low-productivity firm. Recall now that in our model, higher volatility in maintenance reduces firm productivity, due to concavity in the benefit function of maintenance. As a result, for identical marginal distributions of maintenance volatility across firms, an economy where high-productivity firms face exogenously lower volatility would see higher average firm values than an alternative where the distribution of firm productivity and maintenance volatility was uncorrelated. In other words, the fact that productivity and maintenance volatility are estimated to be exogenously negatively correlated imply some degree of positive sorting in the data. This helps to mitigate, to some degree, the negative effects of maintenance volatility.

We can see this effect play out in Panels B and C, which separates out the impact of maintenance volatility on measured firm productivity and the shadow cost of capital. For the shadow cost of capital, we compute the log difference between ARPK and the user cost of capital, motivated by the wedge accounting literature (e.g., Hsieh and Klenow (2009)). Looking to the first two rows, we see that for both cases where maintenance volatility is lower, measured productivity is higher and the shadow cost is lower. For instance, average productivity is 2.1% higher when all volatility is muted, and around 0.9% higher when volatility is partially muted; conversely, ARPK is around 6.23% and 5.4% lower, respectively. In other words, lowering volatility improves firm performance by simultaneously increasing productivity and reducing the implicit cost of investment. Conversely, in the case where we impose a zero correlation between productivity and maintenance volatility, we see that average productivity is lower by around 0.2%, while average ARPK is also slightly lower by around 0.07%.

B Aggregate Outcomes

We now assess the impact of maintenance volatility on aggregate economic performance, by reporting the impact of maintenance volatility on aggregate GDP, capital, TFP, and the dispersion of firm-level ARPK. We use ARPK as a proxy for the allocative efficiency of capital (Hsieh and Klenow (2009)).

Table 7 summarizes our results. When we mute maintenance volatility (first row), we observe that GDP increases by almost 4.1%. This improvement in GDP is driven by parallel increases in the aggregate capital stock (around 4.9%) and aggregate TFP (0.8%). In the last column, we observe that the increase in aggregate TFP is not simply driven by increases in average firm productivity, but is in fact accompanied by a fall in capital "misallocation" (the dispersion of ARPK).

Why does the dispersion of ARPK fall? In our model, ARPK dispersion arises due to both uncertainty in productivity (i.e., time-to-build, as in Asker et al. (2014)), and shocks to the user cost of capital via shocks to maintenance costs. From an aggregate perspective, these shocks to capital user costs is isomorphic in their impact on aggregate productivity as the reduced-form capital "wedges" identified in the misallocation literature (e.g., David and Venkateswaran (2019)). As such, in the counterfactual where these shocks are absent, we observe a corresponding decline in the dispersion of ARPK, and in turn, an increase in measured aggregate TFP.

In contrast, when we mute the correlation between productivity and maintenance volatility (last row), we find that GDP and the aggregate capital stock would have been lower by around 3.6% and 3.1% respectively, if firm productivity and maintenance volatility were exogenously uncorrelated. In turn, aggregate TFP is lower by around 1.1%. Therefore, similar to our analysis at the firm-level, we find that the exogenously negative correlation generates positive sorting that negates some of the negative impact of maintenance volatility. However, we observe that reducing the correlation does not meaningfully affect capital "misallocation", which reflects the fact that the distribution of maintenance cost shocks–and thus user cost shocks–have not changed in this counterfactual.

C Sensitivity Analysis: Model with Endogenous Capital Vintages

Our analysis so far assumes that all the volatility in maintenance is due to shocks to either productivity or the marginal costs of maintenance. However, in reality, some variation in maintenance could simply be due to mechanical aging of the capital stock. For instance, the cost of maintenance could be rising for an older car. This means that our counterfactual might be over-stating the possible gains, since there is nothing that can be done about this dimension of variation.

We now extend our model to incorporate this dimension. Specifically, we extend our model to include a "vintage" effect in capital, where the relative price of maintenance now follows the following process

$$\log p_{it} = (1-\rho_m)\log \bar{p} + (1-\rho_m)\log \mu_j + \rho_m \log p_{i,t-1} + \sigma_i^m u_{it}, \tag{16}$$

where the main difference from our baseline is the inclusion of the variable μ_j. Here, $j = 1, 2, ...$ indexes capital age, with $j = 1$ being the newest vintage. We assume that μ_j is increasing in j, where specifically, μ evolves as

$$\log \mu_j = \log \mu_{j-1} + \log \theta, \tag{17}$$

where $\theta \geq 0$ is a constant and $\mu_1 = 1$. This therefore nests our baseline model, for which $\log \theta = 0$.

Finally, following Cooper et al. (1999), we assume that every positive investment is associated with a complete replacement of the capital vintage. Therefore, whenever a positive investment is done, the firm resets its capital vintage to $j = 1$. In the interest of space, Appendix B presents a complete description of the model, along with a formal definition of the firm's decision problem.

With this extended framework, we now have one more parameter to calibrate, θ. To calibrate θ, we draw on evidence from contemporaneous research by Kabir et al. (2024). The authors observe that for the same set of firms in India, maintenance co-evolve around investment spikes in a way that is not easy to resolve with standard models absent capital vintages. Specifically, they find that maintenance rates follow an increasing pattern in the five years leading up to an investment spike, peaks in the same year as the investment spike, and immediately mean-reverts and stays around the new lower level for the next five years. Panel A of Figure 5 replicate their figure for reference.

Conversely, in our extended model, this co-movement can be replicated through capital aging. Specifically, because investment spikes in our model are primarily generated through capital aging, that is, when the capital user cost is too high, firms are motivated to replace their capital stock, this investment spike also occurs simultaneously with maintenance. In our case, firms in the inaction region "wait" for the capital to eventually reach a stage of excessively high maintenance before they replace the capital stock. With this logic in mind, we set θ to target two moments: (i) the difference in maintenance rate between $t = -5$ and the event, and (ii) the *average* annual growth rate in maintenance rates during the five year window preceding the spike. Furthermore, we restrict the maximum number of vintages in the economy to five (for computational reasons). Panel B of Figure 5 plots our model's result with regards to the event study, and shows a close fit to the data.

In contrast, Figure OA.5 in the appendix, which plots the same event study in our baseline model, further show that the inclusion of the age dimension is key; absent the vintage effect, maintenance always hits a maximum *before* the investment spike. The reason this happens is because investment spikes are also partly triggered by an unexpectedly large reversal in maintenance cost. Specifically, when maintenance cost is high, investment is lower. Whereas when maintenance cost is low, since maintenance shocks are persistent, investment spikes when firms unexpectedly see a large fall in maintenance shocks.

Finally, Table OA.5 in the appendix presents our overall model fit. while Tables OA.6

and OA.7 present our decomposition and analysis mirroring the earlier discussion. Notice that the results largely mirrors the baseline. For instance, GDP is about 4.8% higher when maintenance volatility is muted, similar to the 4.1% in our baseline.

D Summary of Results

To summarize, we uncover that while maintenance volatility accounts only for a small fraction of the total statistical variation in firm productivity, it has a potentially large impact on both firm values and aggregate outcomes. Contrary to our initial hypothesis, however, we also uncover that the strong negative co-variation between productivity and maintenance volatility is predominantly driven by exogenous co-variation, not endogenous reasons. As a result, this generates positive sorting that partly mitigates the losses from maintenance volatility.

A natural subsequent question is to what extent these differences in volatility of maintenance across firms is an inherent "nature" of the firm and thus immutable, making these gains not realizable; or whether these differences in volatility is a result of firm decisions. In the final section, we examine two channels through which firms can reduce volatility in maintenance, showing that these gains could potentially be realized.

7 Microfoundations for differences in maintenance volatility

So far, we have taken it as given that firms differ in their fundamental maintenance volatility. In this section, we hypothesize two potential drivers of maintenance volatility and provide empirical evidence supporting our hypothesis.

A Financial constraints

A factor that could contribute to higher maintenance volatility is financial constraints. Financially constrained firms could delay maintenance expenses. This delay, however, could be potentially costly further down the road, and firms not maintaining their capital properly might be forced to compensate later on by paying high maintenance expenses if poor maintenance leads to breakdown and asset write-off.[9]

[9]Figure OA.6 in the Appendix provides evidence that firms with lower maintenance rates are more likely to have a fixed asset write-off in the future.

To measure maintenance volatility ($\sigma^{t,t-3}(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility over the last four years (years t to t-3).

We use five proxies for financial constraints to test the relationship between financial constraints and maintenance volatility. We use the logged value of total assets and sales to proxy for firm size (Hadlock and Pierce, 2010). We use a dividend dummy, equal to one if the firm paid dividends in that year and zero otherwise (Fazzari et al., 1987; Whited and Wu, 2006). Firm leverage, defined as the total debt divided by total assets, and high leverage dummy is a dummy variable equal to one for the highest quartile of leverage (Giroud and Mueller, 2017). We use debt coverage ratio, defined as EBITDA/Debt, and Low Coverage Ratio Dummy equals one for the lowest quartile of debt coverage ratio. To use identical samples across columns, we require non-missing values for all variables. To ensure no overlap between observations used to calculate maintenance volatility and the proxies for financial constraints, we lag proxies for financial constraints by four years.

Table 8 shows the relationship between maintenance volatility and various proxies for financial constraints. Column (1) shows that firms with higher total assets have lower maintenance volatility. Column (2) shows that firms with higher sales show lower maintenance volatility, as well.[10] Columns (3), (4), and (5) provide further evidence suggesting that financially constrained firms have higher maintenance volatility.

We further evaluate our hypothesis that the higher volatility in maintenance is driven by delays in maintenance as firms attempt to shift expenses to the future when faced with tighter borrowing constraints. To that end, we estimate the following equation,

$$\log p\tilde{m}/k_{it} = \alpha_{i,t} + \sum_{j=0}^{2} \beta_j \log x_{i,t-j} + \sum_{j=1}^{2} \gamma_j \log\left(p\tilde{m}/k\right)_{i,t-j} + \zeta \log \omega_{it} + \epsilon_{it}, \qquad (18)$$

where $\alpha_{i,t}$ denotes firm and year fixed effects and x_{it} refers to a specific measure of financial constraint. Our estimates of interest are the sequence of β_j; in particular, we would anticipate $\beta_0 < 0$ while $\beta_2 > 0$. That is, firms that face tighter constraints currently will delay maintenance into the future ($\beta_2 > 0$) by cutting maintenance spending today ($\beta_0 < 0$). Note that for this specification, we additionally control for two lags of maintenance spending, capturing our earlier concern that some degree of maintenance spending is directly associated with the capital vintage itself. Furthermore, we also control for contemporaneous measured productivity (ω) to account for the fact that, in our

[10] This is also consistent with the findings in Table 2 where we find that larger firms have less volatile maintenance rates. Nevertheless, the variation used to estimate that regression is across firms, whereas here, we are using within-firm variation to estimate maintenance volatility.

model, maintenance is also complementary to productivity. This address the concern that our measures of financial constraint is associated with productivity itself.

Table 9 below reports our results, where we use three measures of financial constraints, namely, the debt to asset ratio, the asset to dividend ratio, and debt to EBITDA ratio. In all cases, we observe $\beta_0 < 0$ and $\beta_2 > 0$, with all estimates statistically significant at conventional levels.

B Capital upgrading

A second factor that we consider is the role of capital vintages. Specifically, it is possible that newer (and thus higher quality) capital breaks less often, and hence maintenance requirements become more predictable (and thus less volatile). Note that this hypothesis is different from our extended model, which only allows for maintenance *cost* to rise with age. In contrast, this hypothesis implies that maintenance volatility itself also rises with capital age.

We test this assumption by leveraging the observation that new investments are typically associated with increases in capital quality, such as purchasing new machinery (Kabir and Mansouri, 2021). Our empirical strategy borrows from Cooper et al. (1999) and our earlier extended model, postulating that investment "spikes" are associated with capital upgrading. To that end, we study the evolution of maintenance *volatility* around the year of peak investment. We use the regression specification in equation 14 of Section 6, where the main difference is that the event ($t = 0$) is the year with peak gross investment. Panel A in the Appendix OA.1 shows that investment is lumpy, consistent with earlier literature (Doms and Dunne, 1998). Panel (B) shows that maintenance rate raises before the peak investment rate, spikes at the peak, and then drops after the event; a result we already discussed earlier.

In Figure 3, we plot maintenance volatility around the peak investment event. First, we observe that from ($t = -3$ to $t = 0$) maintenance volatility starts to increase. From ($t = 1$ to $t = 4$) maintenance volatility increases mechanically because of the sharp change in maintenance rate documented in Panel (A) of Figure OA.1. This sharp change in maintenance volatility lasts for four years. More importantly, we also observe that in periods 5 and afterward, which don't include both pre- and post-observations, maintenance volatility goes down, suggesting that increased capital quality lowers maintenance volatility. In Table OA.8, we use regression to compare the differences in maintenance volatility before and after the event (excluding periods 1-4, to remove the influence of the mechanical increase in maintenance volatility), and find that after the peak investment maintenance

volatility drops.

8 Conclusion

Capital maintenance, just as capital investment, is a necessary part of firm operations. However, unlike capital investment, there are very few studies, if any, on the impact of capital maintenance decisions on firm productivity or macroeconomic outcomes.

Our article shows, using indirect inference through a structural model where maintenance volatility is a model primitive, that volatility in maintenance has negative impact on firm performance and the aggregate economy, because the benefits to maintenance exhibits decreasing returns. Furthermore, we provide additional evidence that such volatility might be endogenous decisions related to firm financing and adoption of lower quality capital, providing further avenue for improvement by policymakers or firm managers hoping to raise aggregate productivity or firm performance. Altogether, our study provides a first look at the important role of maintenance in firm and aggregate productivity, thus suggesting an important avenue for future research into the drivers of firm productivity.

9 Figures and Tables

A Tables

Table 1: **Summary Statistics**

	Mean	Std. Dev.	p25	Median	p75
capital (k)	31.895	93.970	1.653	5.133	18.389
value added (y)	47.367	119.289	3.125	9.554	31.119
gross investment $(i = k' - (1-\delta)k)$	3.922	13.679	-0.021	0.193	1.667
gross investment rate $(\frac{i}{k})$	0.246	0.776	-0.012	0.052	0.200
log measured productivity $(\log \omega)$	0.975	0.802	0.510	0.997	1.483
ARPK $(\frac{y}{k})$	3.824	7.918	0.950	1.817	3.711
maintenance expenses (pm)	0.963	2.616	0.048	0.165	0.594
maintenance rate $(p\frac{m}{k} = p\tilde{m})$	0.063	0.091	0.016	0.035	0.073
maintenance to value added $(\frac{pm}{y})$	0.029	0.035	0.009	0.019	0.036
maintenance volatility $(\sigma^{t,t-3}(\log(p\tilde{m})))$	0.474	0.370	0.232	0.379	0.599
Number of firm-year observations	94,853				

This table presents the summary statistics for the main sample from 1989 to 2019. Capital (k) is measured as the value of PPE. Value added (y) is the difference between sales and cost of material input. Investment (i) is the gross investment and is measured as $i_t = PPE_t - (1-\delta) \times PPE_{t-1}$, where $\delta = 10\%$. Maintenance (m) is the repair and maintenance item from the income statement. $\log \omega$ is measured logarithm of TFP and is defined as a Solow residual $\log \omega = \frac{1}{\Theta_z}(\log py - \Theta_k \log k)$, and $\log ARPK = \log y - \log k$. To measure maintenance volatility $(\sigma^{t,t-3}(\log(p\tilde{m})))$, we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility over the last four years (years t to t-3). Where relevant, the values are reported in millions of US dollars in 2019.

Table 2: **Maintenance volatility and firm outcomes**

Dependent variable	$\mu_i(\log(\text{value added}))$	$\mu_i(\log(\omega))$	$\mu_i(\log(k))$	$\mu_i(\log(\text{ARPK}))$
	(1)	(2)	(3)	(4)
$\sigma_i(\log(p\tilde{m}))$	-1.67***	-0.53***	-1.76***	0.10**
	(0.078)	(0.033)	(0.077)	(0.040)
$\mu_i(\log(p\tilde{m}))$	-0.73***	0.44***	-2.82***	2.07***
	(0.11)	(0.046)	(0.11)	(0.057)
Observations	5,277	5,277	5,277	5,277

This table reports the relationship between maintenance volatility and firm performance measures, i.e., value added, productivity (ω), capital (k), and ARPK. Each observation represents one firm. To construct the firm-level maintenance volatility over time ($\sigma_i(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility of the residuals over time for each firm. To calculate firm-level maintenance average over time ($\mu_i(\log(p\tilde{m}))$, we take the mean of the residuals over time for each firm. To calculate $\mu_i(\log(\text{value added}))$, $\mu_i(\log(\omega))$, $\mu_i(\log(k))$, and $\mu_i(\log(\text{ARPK}))$ we residualize each variable using industry-by-year fixed effects and then take the mean of the residuals over time for each firm. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3: **The effect of maintenance on productivity**

Dependent variable	$\log \omega_{i,t}$			
	(1)	(2)	(3)	(4)
$\log(p_{i,t}\tilde{m}_{i,t})$	0.23***	0.43***	0.47***	0.47***
	(0.0041)	(0.032)	(0.040)	(0.043)
Fixed Effects				
Firm	✓	✓	✓	✓
Year	✓	✓	✓	✓
IV		Leave out focal firm	Leave out focal firm's district	Leave out focal firm's state
Observations	94,853	94,853	94,853	94,853

This table reports the effect of maintenance ($\log(p_{i,t}\tilde{m}_{i,t})$) on productivity ($\log(\omega_{i,t})$) using equation 2. Column (1) shows the OLS estimate, while Columns (2) to (4) provide IV estimates using the $\log(p\tilde{m})$ of firms in the same industry-year cell, leaving out the focal firm, focal firm's district, and focal firm's state, respectively. Standard errors are clustered at the firm level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: **Calibration and Parameters**

(a) Parameters

ρ_z	σ_z	$\bar{p}$	ρ_m	σ_m	χ	f	ς_a	ς_m	γ
0.8047	0.2150	0.0997	0.730	0.744	6.610	0.0314	0.674	0.367	-0.217

(b) Model fit

Moments	**Data**	**Model**
Targeted		
Autocorr. of $\log \omega$	0.678	0.664
Vol. of $\log \omega$	0.270	0.267
Avg. maintenance rate	0.064	0.064
Autocorr. of log maintenance rate	0.621	0.656
Vol. of log maintenance rate	0.537	0.619
Fraction of firms reporting $\lvert i/k \rvert > 0.2$	0.235	0.232
$\beta_{\log p\tilde{m},\log arpk}$	0.638	0.574
Std. dev. of μ_i $(\log \omega)$	0.687	0.706
Std dev. of σ_i $(\log p\tilde{m})$	0.256	0.245
$\beta_{\mu_i(\log \omega),\sigma_i(\log p\tilde{m})}$	-0.611	-0.610
Not targeted		
$\beta_{\mu_i(\log \text{value added}),\sigma_i(\log p\tilde{m})}$	-1.67	-1.26
$\beta_{\mu_i(\log k),\sigma_i(\log p\tilde{m})}$	-1.76	-1.17
$\beta_{\mu_i(\log ARPK),\sigma_i(\log p\tilde{m})}$	0.10	0.065

This table reports internally calibrated parameters and the statistics we use for calibration. Panel A reports the calibrated parameters, while Panel B reports our model fit. $\beta_{\log p\tilde{m},\log arpk}$ refers to the regression coefficient of a regression of log maintenance on log ARPK, controlling for firm and industry-by-year fixed effects. $\beta_{\mu_i(\log \omega),\sigma_i(\log p\tilde{m})}$ refers to the regression coefficient of$\mu_i \log \omega$ on $\sigma_i \log p\tilde{m}$. "Targeted" refers to moments explicitly targeted in our calibration routine, while "untargeted" refers to moments that were not explicitly targted by our calibration routine.

Table 5: **Statistical decomposition of contribution of maintenance volatility to total variance in measured firm TFP**

Total	$\sigma_m = \gamma = 0$	$\gamma = 0$	**Marginal R^2 data**
(1)	(2)	(3)	(4)
3.89%	2.88%	0.66%	4.34%

Table 6: **Impact of maintenance volatility on firm outcomes**

	Mean	**25th**	**50th**	**75th**
	Panel A: Change in firm values			
0 shocks	3.82	4.1	3.96	3.79
0 correlation	-4.04	-4.21	-4.12	-3.96
	Panel B: Change in ω			
0 shocks	1.38	1.41	1.39	1.36
0 correlation	-0.2	-0.2	-0.16	-0.22
	Panel C: Change in ARPK			
0 shocks	-6.21	-3.94	-5.35	-8.06
0 correlation	-0.07	-0.27	0	0.24

This table reports firm-level outcomes relative to the baseline. All values are reported as percentage increase relative to baseline.

Table 7: **Impact of maintenance volatility on aggregate outcomes**

	Capital	**GDP**	**TFP**	**Dispersion of ARPK**
0 shocks	4.9	4.1	0.83	-32%
0 correlation	-3.1	-3.63	-1.1	0%

This table reports aggregate outcomes relative to the baseline. All numbers are reported as percentage increase relative to baseline. Dispersion of APRK is reported as a percentage change relative to the variance of log ARPK in the baseline.

Table 8: **Maintenance volatility and financial constraints**

Dependent variable	$\sigma^{t,t-3}(p\tilde{m})$				
	(1)	(2)	(3)	(4)	(5)
$\log(\text{assets}_{i,t-4})$	-0.050***				
	(0.0049)				
$\log(\text{sales}_{i,t-4})$		-0.062***			
		(0.0043)			
$\text{Dividend Dummy}_{i,t-4}$			-0.027***		
			(0.0051)		
$\text{High Leverage Dummy}_{i,t-4}$				0.026***	
				(0.0057)	
$\text{Low Coverage Ratio Dummy}_{i,t-4}$					0.032***
					(0.0049)
Fixed Effects					
Firm	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓
Observations	63,534	63,534	63,534	63,534	63,534

This table reports the relationship between maintenance volatility and various proxies for financial constraints. To measure maintenance volatility ($\sigma^{t,t-3}(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility over the last four years (years t to t-3). All proxies for financial constraints are measured in $t-4$. Dividend dummy equals one if the firm paid (positive) dividends in that year. High Leverage Dummy equals one for the highest quartile of logged value of total debt divided by assets. Low Coverage Ratio Dummy equals one for the lowest quartile of coverage ratio (EBITDA/Debt). Standard errors are clustered at the firm level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 9: **Dynamics of maintenance spending and financial constraints**

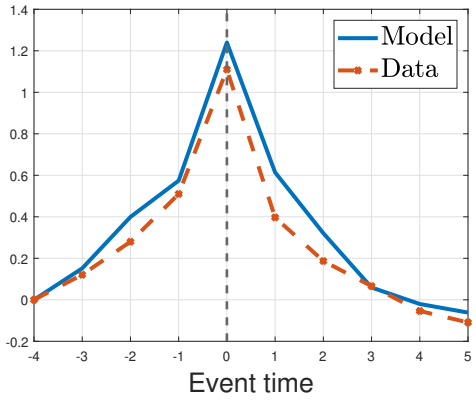
Dependent variable	$\log(p\tilde{m})$		
	(1)	(2)	(3)
β_1	-0.029***	-0.009**	-0.029***
	(0.004)	(0.004)	(0.003)
β_2	-0.024***	-0.012***	-0.029***
	(0.005)	(0.004)	(0.003)
β_3	0.022***	0.010***	0.018***
	(0.005)	(0.004)	(0.003)
Fixed Effects			
Firm	✓	✓	✓
Year	✓	✓	✓
Observations	64,793	18,452	56,234

This table reports the results of equation 18. Column 1 uses debt-to-assets as the independent variable, Column 2 uses assets-to-dividend, and Column 3 uses debt-to-EBITDA. Coefficients for the control variables have been suppressed in the interest of space. Standard errors are clustered at the firm level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

B Figures

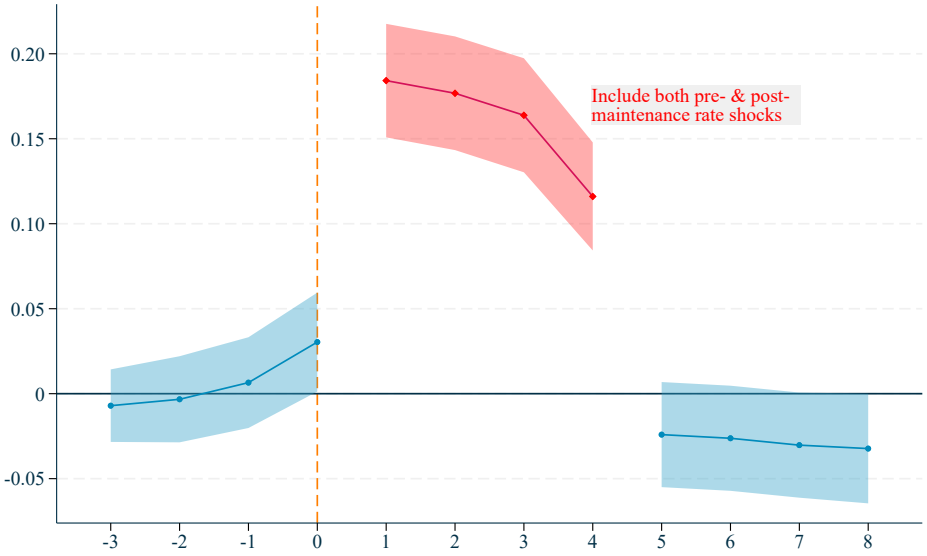
Figure 1: **Maintenance rate around the peak maintenance rate**

Figure 2: Log maintenance rate



The solid blue line plots the point estimate for the model, and the dashed red line with crosses plot the point estimates and 95% confidence interval, for log maintenance rate ($\log(m/k)$) relative to t = -4. Event time = 0 is the period with the peak maintenance rate for a firm.

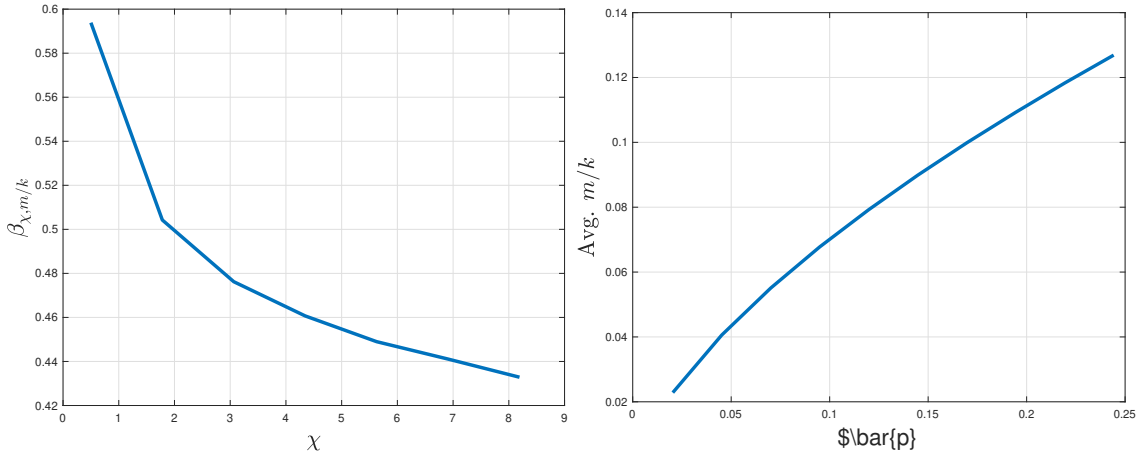
Figure 3: **Maintenance volatility around the peak net investment rate**



This figure plots the point estimates and 95% confidence interval for maintenance volatility relative to $t = -4$. Event time = 0 is the period with the peak investment rate for a firm. To measure maintenance volatility ($\sigma^{t-3,t}(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ by regressing it on its lagged value and industry-by-year fixed effects, then calculate the volatility over the prior years four years (years $t-3$ to t). Maintenance volatility is high in years 1 to 4 because it includes the sharp change in maintenance rate from $t = 0$ to $t = 1$.

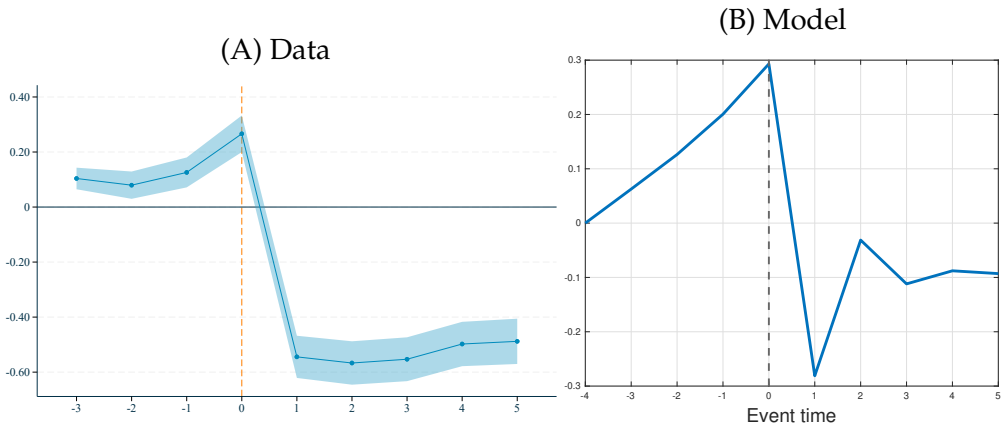
Figure 4: **Comparative statics.**

(A) β vs χ (B) Avg. maintenance vs $\bar{p}$



This figure plots a comparative statics exercise.

Figure 5: **Maintenance dynamics around investment spikes.**



Panel A plots the dynamics of maintenance around an investment spike in the data. Panel B plots our model's predicted maintenance dynamics around an investment spike.

References

ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): "Innovation, Reallocation, and Growth," *American Economic Review*, 108, 3450–91.

ACKERBERG, D. A., K. CAVES, AND G. FRAZER (2015): "Identification properties of recent production function estimators," *Econometrica*, 83, 2411–2451.

ARGENTE, D., D. FITZGERALD, S. MOREIRA, AND A. PRIOLO (2024): "How Do Entrants Build Market Share? The Role of Demand Frictions," *American Economic Review: Insights*.

ASKER, J., A. COLLARD-WEXLER, AND J. D. LOECKER (2014): "Dynamic Inputs and Resource (Mis) Allocation," *Journal of Political Economy*, 122, 1013–1063.

BARTELSMAN, E., J. HALTIWANGER, AND S. SCARPETTA (2013): "Cross-country differences in productivity: The role of allocation and selection," *American Economic Review*, 103, 305–334.

BECKER, S. O. AND H. K. HVIDE (2021): "Entrepreneur Death and Startup Performance," *Review of Finance*, 26, 163–185.

BLOOM, N., B. EIFERT, A. MAHAJAN, D. MCKENZIE, AND J. ROBERTS (2012): " Does Management Matter? Evidence from India," *The Quarterly Journal of Economics*, 128, 1–51.

CAUNEDO, J. AND E. KELLER (2021): "Capital obsolescence and agricultural productivity," *Quarterly Journal of Economics*, 136, 505–561.

COOPER, R., J. HALTIWANGER, AND L. POWER (1999): "Machine replacement and the business cycle: Lumps and bumps," *American Economic Review*, 89, 921–946.

COOPER, R. W. AND J. C. HALTIWANGER (2006): "On the Nature of Capital Adjustment Costs," *The Review of Economic Studies*, 73, 611–633.

DAVID, J. M., L. SCHMID, AND D. ZEKE (2022): "Risk-adjusted capital allocation and misallocation," *Journal of Financial Economics*, 145, 684–705.

DAVID, J. M. AND V. VENKATESWARAN (2019): "The sources of capital misallocation," *American Economic Review*, 109, 2531–2567.

DOMS, M. AND T. DUNNE (1998): "Capital adjustment patterns in manufacturing plants," *Review of Economic Dynamics*, 1, 409–429.

FAZZARI, S., R. G. HUBBARD, AND B. C. PETERSEN (1987): "Financing constraints and corporate investment," .

GIROUD, X. AND H. M. MUELLER (2017): "Firm leverage, consumer demand, and employment losses during the great recession," *The Quarterly Journal of Economics*, 132, 271–316.

GOPINATH, G., KALEMLI-ÖZCAN, L. KARABARBOUNIS, AND C. VILLEGAS-SANCHEZ (2017): "Capital Allocation and Productivity in South Europe*," *The Quarterly Journal of Economics*, 132, 1915–1967.

GORODNICHENKO, Y. AND M. SCHNITZER (2013): "Financial Constraints And Innovation: Why Poor Countries Don't Catch Up," *Journal of the European Economic Association*, 11, 1115–1152.

GOURIO, F. AND L. RUDANKO (2014): "Customer Capital," *The Review of Economic Studies*, 81, 1102–1136.

GREENWOOD, J., J. M. SANCHEZ, AND C. WANG (2010): "Financing Development: The Role of Information Costs," *American Economic Review*, 100, 1875–91.

GUZMAN, J. AND A. LI (2023): "Measuring Founding Strategy," *Management Science*, 69, 101–118.

HADLOCK, C. J. AND J. R. PIERCE (2010): "New evidence on measuring financial constraints: Moving beyond the KZ index," *The review of financial studies*, 23, 1909–1940.

HAUSMAN, J., G. LEONARD, AND J. D. ZONA (1994): "Competitive analysis with differenciated products," *Annales d'Economie et de Statistique*, 159–180.

HAUSMAN, J. A. (1996): "Valuation of New Goods under Perfect and Imperfect Competition," in *The Economics of New Goods*, ed. by T. F. Bresnahan and R. J. Gordon, Chicago: National Bureau of Economic Research, 209–247.

HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 127, 1403–1448.

KABIR, P. AND S. M. MANSOURI (2021): "Capital quality, productivity, and financial constraints: Evidence from India," *Productivity, and Financial Constraints: Evidence from India (November 5, 2021)*.

KABIR, P., E. TAN, AND I. VARDISHVILI (2024): "Quantifying the Allocative Efficiency of Capital: The Role of Capital Utilization," Working Paper.

LEVINSOHN, J. AND A. PETRIN (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, 70, 317–341.

MCGRATTAN, E. AND J. SCHMITZ (1999): "Maintenance and repair: Too big to ignore," *Federal Reserve Bank of Minneapolis Quarterly Review*, 23, 2–13.

MIDRIGAN, V. AND D. Y. XU (2014): "Finance and misallocation: Evidence from plant-level data," *American Economic Review*, 104, 422–458.

OTTONELLO, P. AND T. WINBERRY (2024): "Capital, Ideas, and the Costs of Financial Frictions," *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

RESTUCCIA, D. AND R. ROGERSON (2008): "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic Dynamics*, 11, 707–720.

SHORROCKS, A. F. (2013): "Decomposition procedures for distributional analysis: a unified framework based on the Shapley value," *The Journal of Economic Inequality*, 11, 99–126.

SMITH, M., D. YAGAN, O. ZIDAR, AND E. ZWICK (2019): "Capitalists in the Twenty-First Century*," *The Quarterly Journal of Economics*, 134, 1675–1745.

STERK, V., P. SEDLÁČEK, AND B. PUGSLEY (2021): "The Nature of Firm Growth," *American Economic Review*, 111, 547–79.

WHITED, T. M. AND G. WU (2006): "Financial constraints risk," *The review of financial studies*, 19, 531–559.

Appendix

This appendix contains supplementary material, tables, and figures.

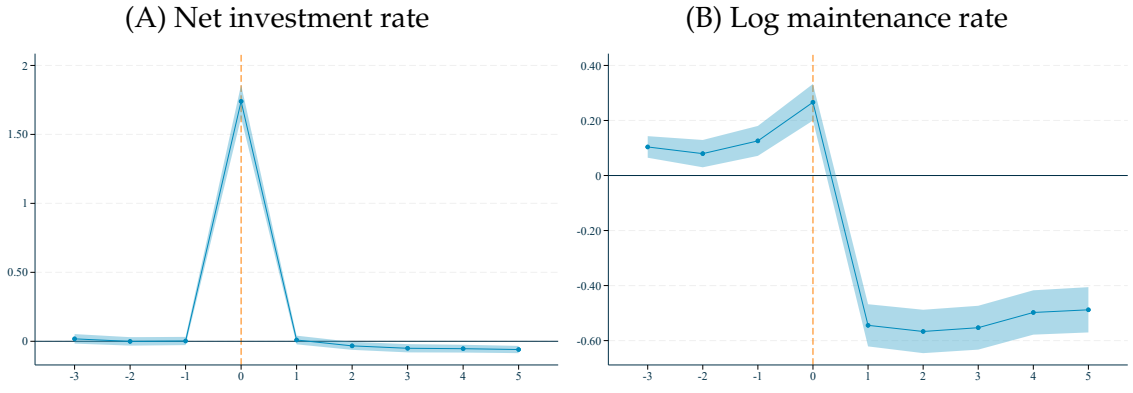
A Extended Empirical Analyses

Table OA.1: **Robustness: Maintenance volatility and firm outcomes with more controls**

Dependent variable	$\mu_i(\log(\text{value added}))$	$\mu_i(\log(k))$	$\mu_i(\log(\text{ARPK}))$
	(1)	(2)	(3)
$\sigma_i(\log(p\tilde{m}))$	-0.54***	-1.09***	0.55***
	(0.032)	(0.065)	(0.033)
$\mu_i(\log(p\tilde{m}))$	-1.69***	-3.41***	1.70***
	(0.043)	(0.088)	(0.044)
$\sigma_i(\log(\omega))$	0.16***	0.36***	-0.17***
	(0.038)	(0.078)	(0.039)
$\mu_i(\log(\omega))$	2.22***	1.44***	0.78***
	(0.013)	(0.026)	(0.013)
Observations	5,277	5,277	5,277

This table tests the robustness of Table 2 to controls. The table shows the relationship between maintenance volatility and firm performance measures, i.e., value added, capital (k), and ARPK. Each observation represents one firm. To construct the firm-level maintenance volatility over time ($\sigma_i(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility of the residuals over time for each firm. To calculate firm-level maintenance average over time ($\mu_i(\log(p\tilde{m}))$), we take the mean of the residuals over time for each firm. To calculate $\mu_i(\log(\text{value added}))$, $\mu_i(\log(\omega))$, $\mu_i(\log(k))$, and $\mu_i(\log(\text{ARPK}))$ we residualize each variable using industry-by-year fixed effects and then take the mean of the residuals over time for each firm. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure OA.1: **Net investment and maintenance rate around the peak net investment rate**



Panel (A) and (B) plot the point estimates and 95% confidence interval for net investment (i/k) and log maintenance rate ($\log(p\tilde{m})$) relative to t = -4. Event time = 0 is the period with the peak investment rate for a firm.

Table OA.2: **Robustness: Maintenance volatility and various productivity measures**

Dependent variable	$\mu_i(\log(\omega))$	
Estimation method	Levinsohn and Petrin (2003)	Ackerberg et al. (2015)
	(1)	(2)
$\sigma_i(\log(p\tilde{m}))$	-0.25***	-0.24***
	(0.033)	(0.032)
$\mu_i(\log(p\tilde{m}))$	0.26***	0.48***
	(0.046)	(0.046)
Observations	5,277	5,277

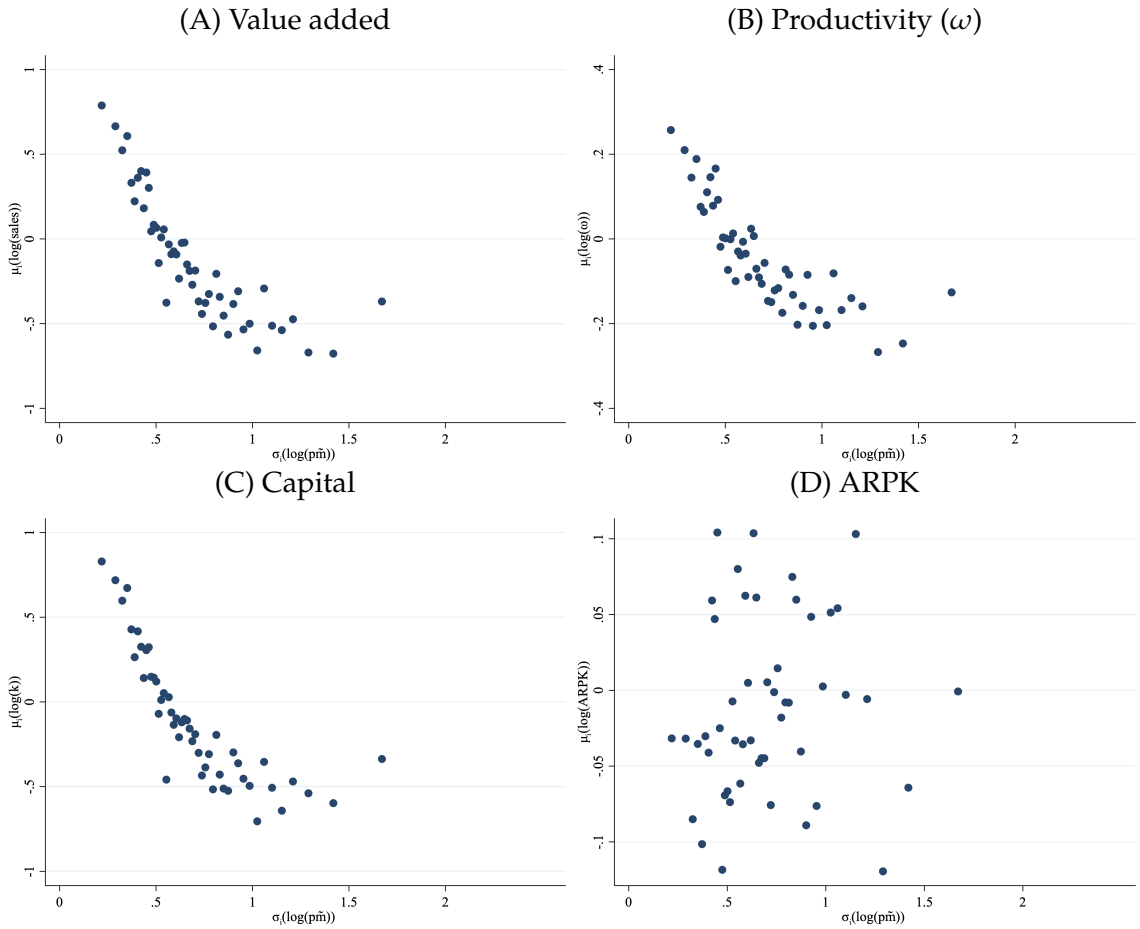
This table tests the robustness of the relationship between maintenance volatility and productivity measured using Levinsohn and Petrin (2003) and Ackerberg et al. (2015). Each observation represents one firm. To construct the firm-level maintenance volatility over time ($\sigma_i(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ regressing it on its lagged value and industry-by-year fixed effects, then calculate the volatility of the residuals over time for each firm. To calculate firm-level maintenance average over time ($\mu_i(\log(p\tilde{m}))$), we take the mean of the residuals over time for each firm. To calculate $\mu_i(\log(\omega))$, we residualize each variable using industry-by-year fixed effects and then take the mean of the residuals over time for each firm. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.3: **Robustness: Maintenance volatility using AR(2) process and firm outcomes**

Dependent variable	$\mu_i(\log(\text{value added}))$	$\mu_i(\log(\omega))$	$\mu_i(\log(k))$	$\mu_i(\log(\text{ARPK}))$
	(1)	(2)	(3)	(4)
$\sigma_i(\log(p\tilde{m}))$	-1.61***	-0.53***	-1.61***	0.020
	(0.079)	(0.033)	(0.078)	(0.042)
$\mu_i(\log(p\tilde{m}))$	-0.68***	0.33***	-2.40***	1.70***
	(0.11)	(0.048)	(0.11)	(0.061)
Observations	5,275	5,275	5,275	5,275

This table tests the robustness Table 2 to using AR(2) process for constructing maintenance volatility. The table reports the relationship between maintenance volatility and firm performance measures, i.e., value added, productivity (ω), capital (k), and ARPK. Each observation represents one firm. To construct the firm-level maintenance volatility over time ($\sigma_i(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ by regressing it on its first and second lags and industry-by-year fixed effects, then calculate the volatility of the residuals over time for each firm. To calculate firm-level maintenance average over time ($\mu_i(\log(p\tilde{m}))$), we take the mean of the residuals over time for each firm. To calculate $\mu_i(\log(\text{value added}))$, $\mu_i(\log(\omega))$, $\mu_i(\log(k))$, and $\mu_i(\log(\text{ARPK}))$ we residualize each variable using industry-by-year fixed effects and then take the mean of the residuals over time for each firm. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

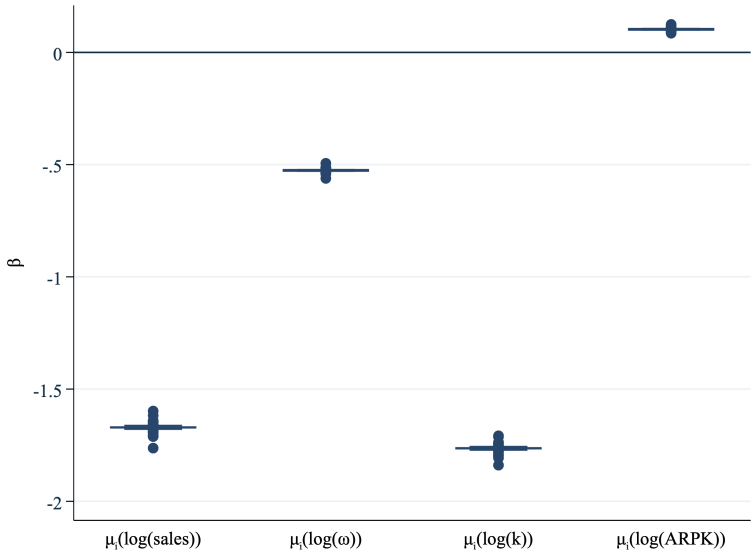
Figure OA.2: **Maintenance volatility and firm outcomes: Binscatter plots**



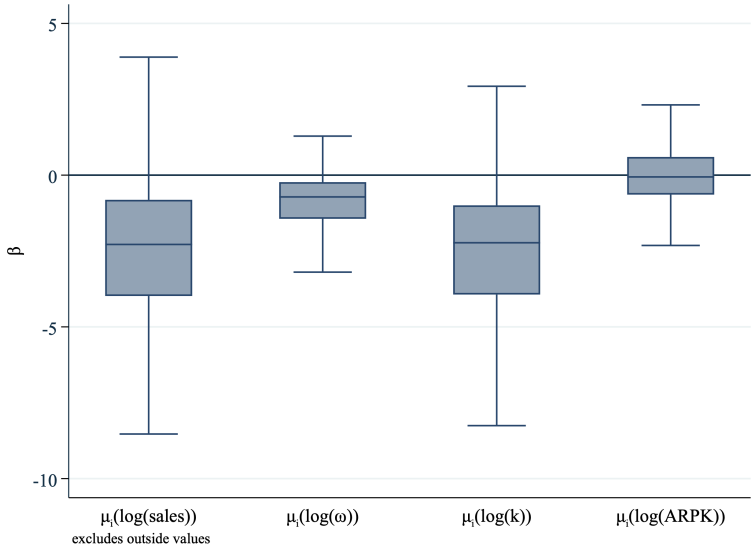
This figure is the binscatter plot of the relationship between maintenance volatility and firm performance measures, i.e., value added, productivity (ω), capital (k), and ARPK, with each panel corresponding to a column with identical regression specification in Table 2. To construct the firm-level maintenance volatility over time ($\sigma_i(\log(p\tilde{m}))$, we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility of the residuals over time for each firm. To calculate firm-level maintenance average over time ($\mu_i(\log(p\tilde{m}))$, we take the mean of the residuals over time for each firm. To calculate $\mu_i(\log(\text{value added}))$, $\mu_i(\log(\omega))$, $\mu_i(\log(k))$, and $\mu_i(\log(\text{ARPK}))$ we residualize each variable using industry-by-year fixed effects and then take the mean of the residuals over time for each firm. We control for the mean maintenance rate ($\mu_i(\log(p\tilde{m}))$) in all the plots.

Figure OA.3: **Maintenance volatility and firm outcomes: Industry exclusion and focus**

(A) Excluding industries individually



(B) Only one industry at a time



This figure visualizes the robustness of the coefficients reported in Table 2 by examining the impact of excluding and focusing on specific industries. We use an identical specification to Table 2 and the dependent variables are $\mu_i(\log(\text{value added}))$, $\mu_i(\log(\omega))$, $\mu_i(\log(k))$, and $\mu_i(\log(\text{ARPK}))$. Panel (A) presents the boxplots of the estimated coefficients when excluding one 6-digit NIC industry at a time from the analysis. Panel (B) presents the boxplots of the estimated coefficients when focusing on only one 6-digit NIC industry at a time, with outside values excluded. To construct the firm-level maintenance volatility over time ($\sigma_i(\log(p\tilde{m}))$, we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility of the residuals over time for each firm. To calculate firm-level maintenance average over time ($\mu_i(\log(p\tilde{m}))$, we take the mean of the residuals over time for each firm. To calculate $\mu_i(\log(\text{value added}))$, $\mu_i(\log(\omega))$, $\mu_i(\log(k))$, and $\mu_i(\log(\text{ARPK}))$ we residualize each variable using industry-by-year fixed effects and then take the mean of the residuals over time for each firm. We control for the mean maintenance rate ($\mu_i(\log(p\tilde{m}))$) in all the plots.

Table OA.4: **First stage for Table 3: The effect of maintenance on productivity**

Dependent variable	$\tilde{m}$		
	(1)	(2)	(3)
IV excluding focal firm	0.39***		
	(0.027)		
IV excluding focal firm's district		0.32***	
		(0.025)	
IV excluding focal firm's state			0.29***
			(0.024)
Fixed Effects			
Firm	✓	✓	✓
Year	✓	✓	✓
F-stat	218.12	155.54	139.48
Observations	94,853	94,853	94,853

This table reports the first-stage regressions for the effect of maintenance ($\log z$) on productivity ($\log u$) in Table 3. Columns (1) to (3) provide the first-stage estimates using the average maintenance rate of firms in the same industry-year cell, leaving out the focal firm, focal firm's district, and focal firm's state, respectively. Standard errors are clustered at the firm level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

B Extended Model with Capital Age

B.i Bellman Equations

B.ii Results

Table OA.5: **Calibration and Parameters: Extended model with capital vintage**

(a) Parameters

ρ_z	σ_z	$\bar{p}$	ρ_m	σ_m	χ	f	θ	ς_a	ς_m	γ
0.7954	0.2198	0.105	0.731	0.750	8.85	0.0373	0.4084	0.674	0.380	-0.220

(b) Model fit

Moments	**Data**	**Model**
Targeted		
Autocorr. of $\log \omega$	0.678	0.594
Vol. of $\log \omega$	0.270	0.454
Avg. maintenance rate	0.064	0.067
Autocorr. of log maintenance rate	0.621	0.466
Vol. of log maintenance rate	0.537	0.783
Fraction of firms reporting $\lvert i/k \rvert > 0.2$	0.235	0.465
$\beta_{\log p\tilde{m},\log arpk}$	0.638	0.570
Diff. between maintenance peak and trough	0.02	0.02
Avg. annual difference in pre-period	0.004	0.004
Std. dev. of μ_i $(\log \omega)$	0.687	1.05
Std dev. of σ_i $(\log p\tilde{m})$	0.256	0.241
$\beta_{\mu_i(\log \omega),\sigma_i(\log p\tilde{m})}$	-0.611	-0.880
Not targeted		
$\beta_{\mu_i(\log \text{value added}),\sigma_i(\log p\tilde{m})}$	-1.67	-1.34
$\beta_{\mu_i(\log k),\sigma_i(\log p\tilde{m})}$	-1.76	-1.24
$\beta_{\mu_i(\log ARPK),\sigma_i(\log p\tilde{m})}$	0.10	0.065

This table reports internally calibrated parameters and the statistics we use for calibration. Panel A reports the calibrated parameters, while Panel B reports our model fit. $\beta_{\log p\tilde{m},\log arpk}$ refers to the regression coefficient of a regression of log maintenance on log ARPK, controlling for firm and industry-by-year fixed effects. $\beta_{\mu_i(\log \omega),\sigma_i(\log p\tilde{m})}$ refers to the regression coefficient of$\mu_i \log \omega$ on $\sigma_i \log p\tilde{m}$. "Targeted" refers to moments explicitly targeted in our calibration routine, while "untargeted" refers to moments that were not explicitly targted by our calibration routine.

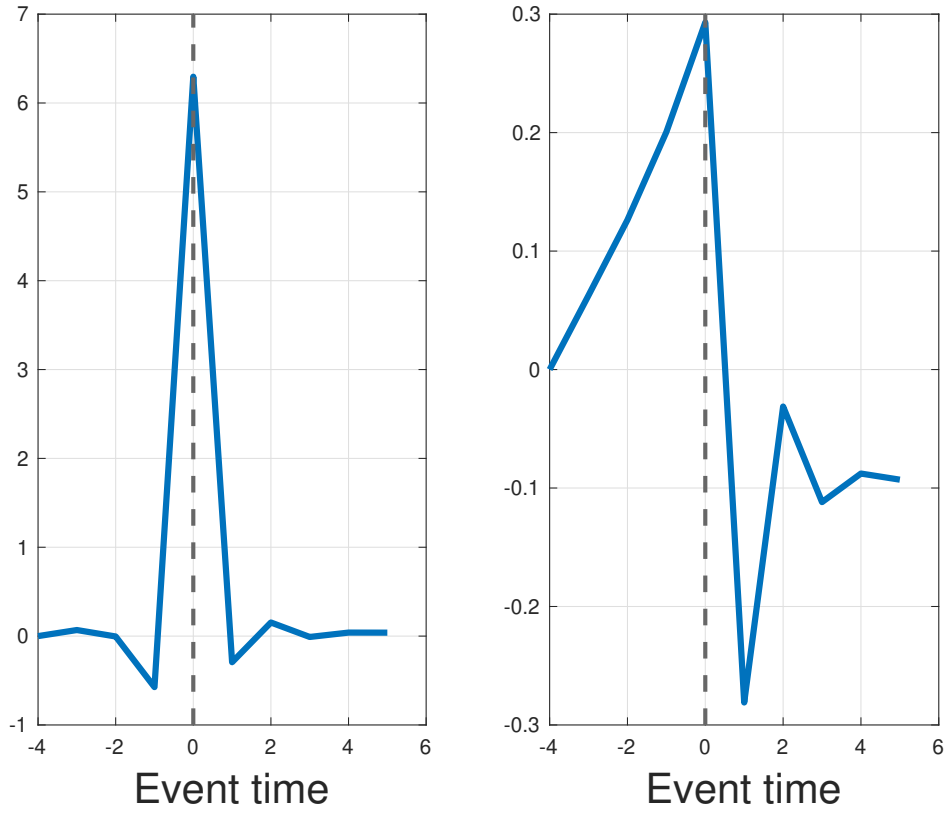
Table OA.6: **Impact of maintenance volatility on firm outcomes**

	Δ in firm values		Δ in ω		Δ in ARPK	
	Mean	**Median**	**Mean**	**Median**	**Mean**	**Median**
0 shocks	4.72	4.66	3.19	3.10	-4.81	-4.85
0 correlation	-2.91	-3.08	-2.91	-3.08	0.72	0.50

Table OA.7: **Impact of maintenance volatility on aggregate outcomes**

	Capital	**GDP**	**TFP**	**Dispersion of ARPK**
0 shocks	6.0	4.8	1.8	-32%
0 correlation	-2.1	-2.7	-1.7	0%

Figure OA.4: **Event study with capital age.**



C Additional evidence for microfoundation

Figure OA.6 plots our results showing that a reduction in maintenance spending leads to greater asset impairment in the following years.

Figure OA.5: **Event study without capital age.**

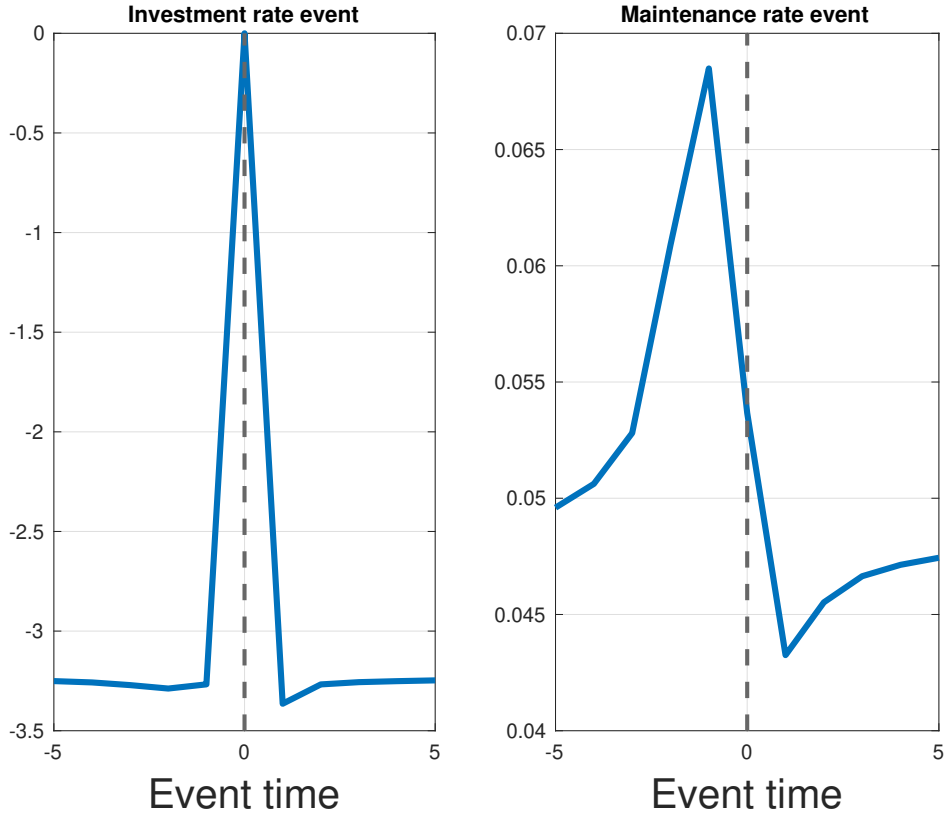
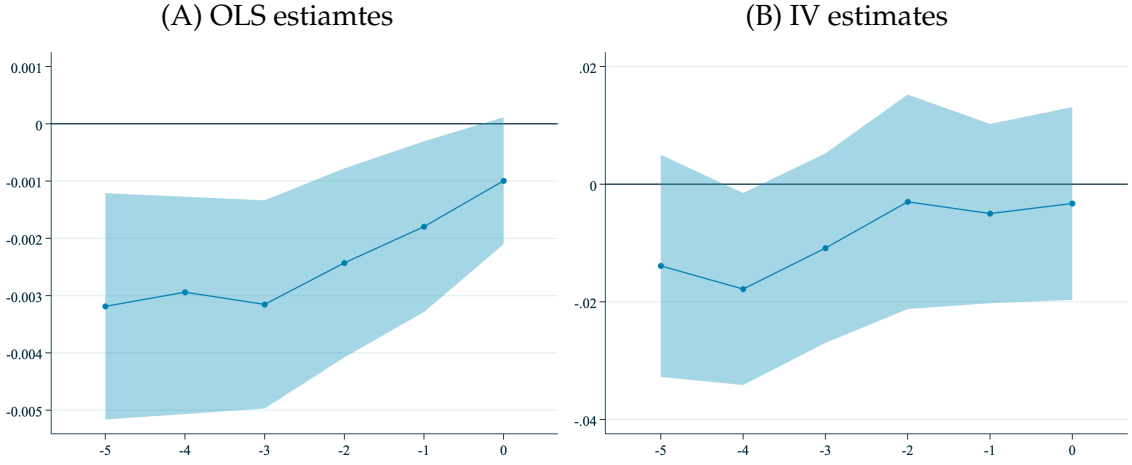


Figure OA.6: **Lagged maintenance and fixed asset impairment**



This figure presents the coefficients and 95% confidence intervals from regressions examining the relationship between fixed asset impairment (indicated by a dummy variable) and lagged maintenance rates. Panel (A) shows the results using OLS estimates, while Panel (B) displays the results from IV estimates. The horizontal axis represents the lag period (in years) of the maintenance rate, with zero indicating the current year and negative values indicating previous years.

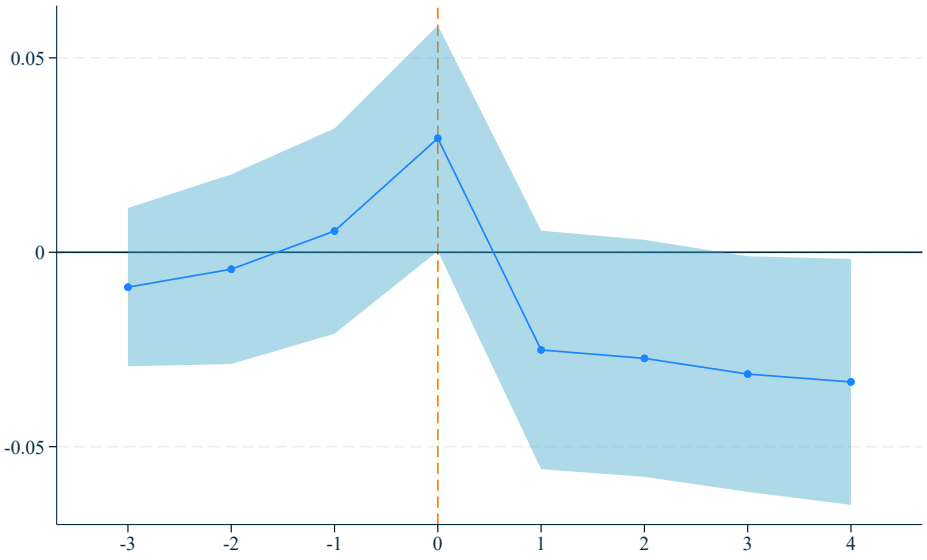
Figure OA.8 plots maintenance volatility around an investment spike, where we use a shifted window to account for the mechanical effects discussed in the main text.

Table OA.8: **Maintenance volatility around the event**

Dependent variable	$\sigma^{t-3,t}(\tilde{m})$	
	(1)	(2)
Post	-0.035***	-0.053***
	(0.012)	(0.016)
Fixed Effects		
Firm	✓	✓
Year	—	✓
Observations	6,143	6,143

This table presents the change in maintenance volatility after the event ($t >= 5$) relative to before the event ($t <= 0$). Event time = 0 is the period with the peak investment rate for a firm. We exclude $1 <= t <= 4$ from this analysis due to a the sharp change in , Figure OA.1 shows the. To measure maintenance volatility ($\sigma^{t,t-3}(\log(p\tilde{m}))$), we first residualize $\log(p\tilde{m})$ using its lagged value and industry-by-year fixed effects, then calculate the volatility over the prior years four years (years t to t-3). Maintenance volatility due to a sharp drop in maintenance rates.

Figure OA.7: **Maintenance volatility around the event**



This figure plots the point estimates and 95% confidence interval for maintenance volatility relative to t = -4. Event time = 0 is the period with the peak investment rate for a firm. To measure maintenance volatility ($\sigma^{t+1,t+4}(\tilde{m})$), we first residualize log(m/k) using its lagged value and industry-by-year fixed effects, then calculate the volatility over the years four years (years t+1 to t+4). To measure maintenance volatility, we use $\sigma^{t-3,t}$ for before the event ($t <= 0$) and $\sigma^{t+1,t+4}$ for after the event ($t > 0$) so that measurement of maintenance volatility is not contaminated by the sharp decline in maintenance rate at event time t = 1. Further details go to equation.