Capital Budgeting, Uncertainty, and Misallocation*

Ben Charoenwong†

Yosuke Kimura

Alan Kwan

Eugene Tan

May 5, 2021

Abstract

We use data on firms' expectations and planned capital expenditures to show that capital budgets (i) capture information beyond simple expectations of future profitability, (ii) are partially flexible, but (iii) are costly to deviate from. To explain these facts, we develop an investment model with capital budgeting, in which firms endogenously learn about firm fundamentals and make partially flexible investment plans. Our calibrated model shows that managers actively use both strategies, leading to substantial amelioration of misallocation arising from uncertainty. We show, through a decomposition exercise, that this arises primarily because high-productivity firms actively allocate a larger fraction of their expenses to better planning. In particular, a recalibrated model with homogenous learning predicts substantially larger misallocation. Our paper highlights the importance of accounting for firm heterogeneity in learning when studying capital misallocation under uncertainty.

*The authors thank Laurent Cavenaile, Murat Celik, Russell Cooper, Francois Gourio, Cosmin Ilut, Abdou Ndiaye, Omar Rachedi, Jincheng Tong, Choongryul Yang, participants at the Virtual Macroeconomic Seminar, and brownbag participants at the University of Toronto and National University of Singapore. The content in this article has been reviewed by the Ministry of Finance of Japan to ensure the anonymity of survey respondents. Any views or opinions expressed herein are not representative of those of the Policy Research Institute, Ministry of Finance, or the Cabinet Office of Japan. Ben Charoenwong acknowledges financial support from the NUS Start-Up Grant No. R-315-000-119-133 and the Singapore Ministry of Education AcRF Tier 1 Research Grant No. R-315-000-122-115. This paper was previously circulated as "Investment Plans, Uncertainty, and Misallocation." Any remaining errors are our own.

†Ben Charoenwong, Email: ben.charoenwong@nus.edu.sg, National University of Singapore Business School, 15 Kent Ridge Dr #07-69, Singapore 119245; Yosuke Kimura, Email: kimura.y.bq@m.titech.ac.jp, Tokyo Institute of Technology, 2-12-1-W9-54 Ookayama, Meguro-ku, Tokyo, 152-8552 Japan; Alan Kwan, Email: apkwan@hku.hk, The University of Hong Kong, K.K. Leung Building 923, University Dr, Lung Fu Shan, Hong Kong; Eugene Tan, Email: Eugene.Tan@rotman.utoronto.ca, Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, Ontario M5S 3E6.

1 Introduction

Corporate managers typically begin a fiscal year with capital budgets, making capital expenditure plans under substantial uncertainty regarding the realization of sales and costs in that year (Harris et al. 1982). As the year progresses, they may deviate from their plans by adjusting actual investments in response to additional information. Consequently, the quality of a firm's investment plan and the flexibility to adjust to real-time changes become jointly important in determining the effect of uncertainty on the profitability of a firm and, on an aggregate level, capital allocation across firms. While a large body of research has studied the role of imperfect information in firm decision-making, the primary focus has been on directly connecting sales or profitability uncertainty to capital expenditure decisions, bypassing the investment planning channel. In part, a difficulty in studying the role of investment planning and flexibility in mitigating the impact of uncertainty on firm bottom lines arises from the lack of access to a dataset with a sufficiently rich set of observations on sales forecasts, capital budgeting, and the underlying balance sheet of the firm.

To that end, in this paper, we use a unique setting in Japan that permits the construction of a nationally representative firm-level panel with headline financial statement variables and rich expectations and planned spending data, along with their ex post realizations, to examine the degree to which firms use two levers of corporate policy to mitigate uncertainty: (i) costly capital budgeting and (ii) the flexibility to adjust investments in response to real-time shocks. Then, we connect these strategies of uncertainty mitigation to a broader macroeconomic literature on dynamic capital misallocation. We present two sets of results. First, we establish three stylized facts that connect capital budgeting and actual investment choices. We show that (i) investment plans are a strong input into actual investment above and beyond a manager's expectations on profitability; (ii) realized investments deviate from plans in response to ex post shocks to profitability, but plans retain forecasting power *even after* accounting for these shocks; and (iii) deviating from plans appears to be costly, with firm profitability declining in the following year when managers deviate more from their initial investment plans through the fiscal year, even when controlling for both ex ante and ex post characteristics such as expected and realized TFP and firm size. These facts suggest that investment plans are not simple conduits through which expectations about sales or productivity feed into investments. Moreover, they are costly to adjust on the fly, but are not completely immutable.

Next, to study the economic implications of our findings, we formulate a parsimonious investment model with costly endogenous learning and capital budgeting in which investment plans are flexible but costly to adjust, and calibrate our model to the data. Our baseline model follows Hopenhayn (1992) with only two additional key ingredients: costly information acquisition and costly deviation from investment plans. The first ingredient — costly information acquisition — means that firms start with uncertainty but may acquire better information (to make better investment plans) at a cost. In practice, this cost may take the form of conducting additional market research or hiring better managers with more forecasting ability.[1] The second ingredient — flexible but costly deviation from investment plans — means that firms want to reduce their uncertainty due to adjustment costs for deviations as the year unfolds. In practice, costs of adjusting plans may manifest in several ways, such as through financing or organizational frictions that involve a large collection of agents coordinating to deviate from initial plans.[2] These micro-founded frictions based on the empirical and theoretical accounting and corporate finance literature introduce a novel intertemporal tradeoff for firms between purchasing better information ex ante against planning to potentially deviate from investment plans ex post.

A key insight of our model is that firms with different characteristics tolerate different levels of uncertainty due to the cost to deviate from initial capital budgets. A higher level of uncertainty is costlier for more productive firms because the returns to learning increase in

[1] See Baik et al. (2011); Lee et al. (2012); Goodman et al. (2014).

[2] Studies such as Harris et al. (1982); Antle and Eppen (1985); Bernardo et al. (2004); and Malenko (2019) study the theoretical implications of intrafirm agency and information frictions for capital allocation within a firm, usually from the angle of studying incentive mechanisms for managers within a firm. However, there is little research on the macroeconomic implications of costly and noisy capital expenditure plans.

productivity. This mechanism gives rise to two unique predictions regarding TFP forecast errors and investment plan deviations, which we test and verify using our data: (1) the dispersion of TFP forecast errors decreases in initial firm productivity and (2) the dispersion of investment plan deviations increases in the absolute size of realized TFP shocks. Because our panel data contain information on forecasts/plans and realizations for both sales and investments as well as other financial statement variables, our data are uniquely suited for testing these predictions. Notably, due to the same underlying mechanism, our model also predicts that firm forecast errors are forecastable, even though we assume full rationality in our model, and all uncertainty is always resolved at the end of the model period.

Our calibrated model suggests that while firms actively adjust their capital expenditures on the fly, they generally prefer to seek better information rather than be forced to adjust ex post. The total resources spent on information acquisition is more than five times that of investment plan deviation. One outcome of this is that revenue is positively correlated with signal acquisition costs expended, but negatively correlated with investment plan adjustment costs. That said, as firms do actively use both levers of corporate policy, attributing capital misallocation to either channel alone could lead to overstatement of the effects of one or the other.

From a macroeconomic perspective, we find that the ability to mitigate uncertainty via investment planning and flexible budget adjustments substantially ameliorates capital misallocation arising from uncertainty. For instance, in our baseline model, we find that aggregate wages are about 0.4% lower than a reference model of costless investment planning and adjustments.[3] Notably, the bulk of the amelioration in capital misallocation arises from our model assumptions that allow managers to optimally choose how much information to acquire. In contrast, a counterfactual model that imposes homogeneous uncertainty (as is common in the literature on uncertainty and misallocation; e.g., Asker et al. 2014; David and Venkateswaran 2019) predicts substantially higher levels of misallocation; for instance,

[3]Such a reference model relates to the framework and timing used in most macroeconomic firm dynamics model; e.g., David and Venkateswaran (2019).

wages are about 2.5% lower relative to the same reference. In contrast, the imposition of immutable plans (while maintaining the endogenous learning assumption) leads to an economy in which wages are about 0.5% lower relative to the same reference model.

Finally, our main theoretical result, whereby higher productivity firms endogenously choose lower subjective uncertainty, also has implications for the business cycle. We show in an extended model with aggregate productivity shocks that our model endogenously generates (1) persistent and countercyclical subjective uncertainty over both idiosyncratic and aggregate conditions and (2) countercyclical capital misallocation, with the latter a natural consequence of the former. Critically, our theoretical results depend on the fact that investment plans are costly to adjust. While both facts have been established in the literature, we also use our data to directly test these predictions. As with our earlier results from the stationary model, the extended predictions are also matched in the data.

After discussing the related literature, the rest of our paper proceeds as follows: Section 2 provides empirical evidence that investment plans are important predictors of actual investment choices, even after accounting for other variables such as expected sales, Section 3 describes our baseline model, Section 4 discusses and tests our main model predictions, and Section 5 discusses key quantitative results through our calibrated model. Section 6 presents our extended model with aggregate risk, and finally, Section 7 concludes.

Related Literature

Our results contribute to several strands of the corporate finance literature on learning, uncertainty, and investment dynamics as well as macroeconomics.

First, our paper builds on the well-established insight in corporate finance that argues for the importance of capital budgeting. Closest to our research is Harris and Raviv (1996), who contend that unanticipated shocks to productivity can lead to inefficient investment, since it is difficult to adjust actual investment in real time. Similar to our paper, they argue using a small sample of firms, that due to inertia in adjusting investment, firms underinvest (overinvest) in response to positive (negative) shocks. Our empirical framework demonstrates that this is true for a broad representative cross-section of firms, and our model allows us to map these "poor" firm responses to macroeconomic outcomes.

Second, our paper connects to the broad literature on information rigidity and rational inattention (e.g., Sims 2003; Reis 2006; Van Nieuwerburgh and Veldkamp 2009; Coibion and Gorodnichenko 2015; Benhabib et al. 2016; Ilut and Valchev 2020). As is the literature, we assume firms are rationally inattentive due to costly information acquisition. However, we focus on how costly information acquisition interacts with partially flexible investment plans and connect this channel to a broader study on dynamic capital misallocation. Our primary focus is on studying ex ante improvements in the quality of capital budgeting in mitigating uncertainty, especially when firms are constrained in how they can respond in real-time to contemporaneous shocks. Importantly, we directly verify our mechanisms using a unique novel data set of firm expectations, in which we observe both expected and realized sales and investment as well as the underlying balance sheet of nationally representative firms.[4]

In this vein, our paper is also related to the literature on behavioral corporate managers that studies how systematic forecast errors affect firm investment decisions (e.g., Gennaioli et al. 2015; Bordalo et al. 2018, 2020; Ma et al. 2020). In this literature, behavioral corporate managers implement capital expenditure policies based on biased expectations, while investment plan formation and adjustments are typically not studied. In contrast, our model is closer to that of standard models of rational inattention, where managers are fully rational, and we emphasize the role of capital budget formation and flexibility in ameliorating the impact of uncertainty. In this broader context, while not a main focus of our paper, our model also generates forecastable sales forecast errors even though managers are assumed to be fully rational.

Third, our findings also relate to the literature on uncertainty as a source of capital misallocation. Recent studies have argued that the combined effect of investment under

[4]In contrast, prior research (e.g., Bachmann et al. 2013; Bachmann and Elstner 2015; Bachmann et al. 2017) use survey data that are less comprehensive and/or only include qualitative forecasts regarding sales or investment.

uncertainty, and capital physical adjustment frictions, can lead to substantial capital misallocation (e.g., David et al. (2016) impose an implicit time-to-build element, while Asker et al. (2014) and David and Venkateswaran (2019) combine time-to-build and physical adjustment costs). Unlike in the literature, uncertainty in our paper generates misallocation because firms need to plan capital budgets under incomplete information and face costly real-time investment plan adjustments. Therefore, our focus is on adjustment costs that arise from nontangible costs such as coordination or opportunity costs that arise due to deviations from initial budgets. We find that uncertainty plays a small role in generating capital misallocation, similar to the finding of David and Venkateswaran (2019). However, we differ in the underlying economic mechanism. In our paper, the diminished role of uncertainty for misallocation arises because of endogenous learning and, importantly, because high-productivity firms optimally incur the bulk of the cost of learning. In contrast, David and Venkateswaran (2019) assume that a fraction of uncertainty is resolved costlessly ahead of time. In this context, we show that while capital misallocation arising from uncertainty is low, this entails a tradeoff: expending more resources to buy better information (at least, from the firm's perspective). Notably, while our framework focuses on a different economic environment, we believe that our underlying economic insight should apply to the broader literature as well.

Our framework also relates to the broader literature that studies aggregate fluctuations in uncertainty. Similar to recent work (e.g., Benhabib et al. 2016; Ilut and Saijo 2020), we propose endogenous procyclical learning as an explanation for countercyclical uncertainty, in contrast to the broader literature in which countercyclical uncertainty is exogenously imposed (e.g., Bloom 2009; Bachmann and Bayer 2014; Bachmann and Elstner 2015; Bloom et al. 2018; Senga 2018). Unlike the recent literature, we focus on a different mechanism whereby countercyclical uncertainty arises because the returns to learning are increasing in firm productivity. In contrast, countercyclicality in Benhabib et al. (2016) arises from a general-equilibrium effect of complementarity in signal acquisition, while in Ilut and Saijo

(2020), fluctuations in uncertainty arises from passive information accumulation.[5] Importantly, we empirically verify our main mechanism using our data by showing that firms with higher productivity have smaller forecast errors over fundamentals. Our empirical finding is similar to that of Tanaka et al. (2020), who find, using a similar but distinct dataset in Japan, that higher productivity firms have lower uncertainty in the form of lower absolute forecast errors over aggregate variables. They hypothesize that this result is due to better management ability, which implies a model in which marginal costs to acquire information decrease in productivity. In contrast, we show that given standard assumptions, the returns to being "correct" are naturally higher for firms with higher productivity, even when the marginal cost of acquiring information is homogeneous across firms.

Finally, our paper is related to the literature on investment dynamics, beginning with Kydland and Prescott (1982), that studies the role of time-to-build and time-to-plan. Our paper's focus is similar to that of Christiano and Todd (1996); Bar-Ilan and Strange (1996); and Kuehn (2011) in studying the implications for firm and aggregate outcomes when planned and realized investments do not necessarily coincide. And although we do not study asset prices, our findings and model are both related to and consistent with empirical findings in Lamont (2000) and modeling assumptions in Li et al. (2020) and Li and Wang (2020), who document the asset-pricing implications of investment plan frictions.

2 Firm Forecasts and Investment Plans

In the following subsections, we first discuss our main data sources and the definition and construction of the key variables we will use in our analysis. Then, we present three stylized facts regarding firm investment plans, which we use to argue that (i) investment plans capture more information than just the expectations of future profitability; (ii) investment plans are only partially flexible; and (iii) investment plans are costly to deviate from. We then use these facts as motivation for our theoretical model in the next section.

[5]In particular, Benhabib et al. (2016) assume that firm productivity is i.i.d., which implies that the initial conditions of firms do not matter with respect to information acquisition.

2.1 Data Sources

Our main data sources are the Business Outlook Survey (BOS) and the Annual and Quarterly Financial Statements Statistics of Corporations by Industry in Japan. The BOS contains firm-level expectations, forecasts, and spending plan data and their realizations, while the Annual (Quarterly) Financial Statistics survey provides detailed year-end (quarter-end) financial statistics such as cash holdings, debt, total employment, cost of labor, and other financial statement variables for the fiscal period. Although the two sets of surveys are distinct, both are administered by the Ministry of Finance (MoF) and follow the same sampling procedure. Appendix D contains a more detailed description of the sampling structure and explains why the sampling procedure effectively permits the construction on a nationally representative and complete panel data of large firms, both publicly listed and private, in Japan.

We merge these two data sets together for our sample, which comprises around 6,000 firms a year, from fiscal year (FY) 2005 (April 2005 to March 2006) to FY 2016 (April 2016 to March 2017). The firms in our sample account for around 60% of total employment in Japan. Table I reports the broad general characteristics of our merged sample. Panel A reports the balance sheet variables as reported in the Annual Financial Statistics survey, and Panel B reports spending plan data from the BOS as well as constructed key variables such as TFP, which we define below.

[Table I Here]

2.2 Definition and Construction of Key Variables

In our sample, we observe a firm's forecasted sales, profits, and investment spending plans for the full year, conditional on information available to a firm as of the first fiscal quarter.[6] For sales and profits, we refer to the differences between actual and realized variables as

[6]Firms are surveyed around six weeks into the first fiscal quarter regarding their expectations for the full fiscal year. Appendix D contains more detailed descriptions of the sampling timing.

"*shocks*." In contrast, for investment plans, we refer to the differences between realized and planned values as "*deviations from plan*." Our interpretation is that differences in actual and predicted sales or profits are largely driven by circumstances outside of a firm's control, and thus can be interpreted as innovations to a firm's information set.[7] In contrast, investment spending is under a firm's control, and thus any differences between realized and planned values are *choices* made by a firm. Such deviations reflect both a firm's reactions to real-time innovations to its information set and the degree of flexibility of a firm to adjust capital expenses relative to its planned expenses. This distinction between innovations to the information set and degree of investment plan flexibility will be the core feature of our analyses and model. Appendix Section A.2 shows that the observed capital spending plans appear to be informative for realized spending, generating results quantitatively similar to or even larger than those in Gennaioli et al. (2015). These results suggest that BOS corporate plans and forecasts are good predictors of actual realizations, and hence economically relevant.

2.2.1 TFP, Expected TFP, and TFP Shocks

The firm-level "fundamental" we consider is its revenue total factor productivity (henceforth TFP), as is commonly used in the literature. For computing realized TFP, we follow the measurement strategy detailed in Asker et al. (2014). We follow the standard convention by assuming that firms operate Cobb-Douglas physical production functions with capital (labor) intensity α $(1-\alpha)$, and face isoelastic demand curves with elasticity η. Given values for α and η, we can back out TFP z using the identity

$$z = \frac{(py)^{\frac{\eta-1}{\eta}}}{k^{\alpha} l^{1-\alpha}},$$

where py is the firm's total value added for the fiscal year, k is the firm's physical capital stock at the end of the last fiscal year, and l is the number of full-time equivalent labor hired as of the end of the current fiscal year. A full description is deferred to Appendix A.3 in the interest of space.

[7]This is consistent with the assumptions made in Bachmann and Elstner (2015).

We also compute an approximate measure of the firm's expected TFP using data on expected sales and balance sheet variables and use it to infer the TFP shock the firm faces. This is unlike prior research, which typically uses the fitted value of an AR(1) regression as a proxy for a firm's expected TFP and the residuals as the TFP shock a firm faces. Specifically, we define expected TFP z^e as

$$z^e \equiv \frac{(py^e)^{\frac{\eta-1}{\eta}}}{k^\alpha l^{1-\alpha}},$$

where py^e is the firm's *expected* total value added for the fiscal year (defined as expected sales minus realized costs of goods sold), and k and l are the same variables as before. We will define a TFP "shock" $\Delta \log z$ as

$$\Delta \log z \equiv \log z - \log z^e,$$

and we will interpret the cross-sectional dispersion of $\Delta \log z$ as the average amount of uncertainty across firms at a particular point of time.[8,9]

2.2.2 Investment Plans and Deviations from Investment Plans

In the BOS, firms are surveyed regarding their planned and actual investment-related spending, including information over three broad categories of spending, namely (i) physical property and equipment (PP&E), (ii) land, and (iii) software. For most of our analysis, we will focus on PP&E as our definition of "capital." We denote i as actual investment, i^p as planned investment, and $\Delta i \equiv i - i^p$ as "deviation from investment plans." Moreover, for all of our analysis, we use investment *rates* rather than investment *levels*, where we normalize

[8] Readers might observe that our measure of expected TFP mismeasures the true expected TFP, and consequently also mismeasures the true TFP shock. In Appendix A.3.3, we demonstrate that given the timing convention and parametric assumptions of our model in Section 3, these measurement issues do not affect our qualitative results, and are also unlikely to be quantitatively large.

[9] Our measure of expected value added also requires an assumption that firms have perfect foresight regarding intermediate factor input costs for the year, as we do not observe forecasted intermediate input cost. For robustness, we therefore verify that all of our reported empirical results using any expected TFP or its derivatives are also robust to using the as-reported expected sales or profits (i.e., replacing TFP shocks with sales shocks or profits shocks).

investment by the initial period capital stock. Correspondingly, investment-plan deviations are defined as

$$\Delta \frac{i}{k} \equiv \frac{i - i^p}{k}.$$

One issue we face is that the BOS focuses on capital expenses. This means that for firms that plan on doing disinvestment and/or end up doing disinvestment at the end of the year, they will report $i^p = 0$ and/or $i = 0$. This manner of reporting will bias our measure of Δi toward 0, which means that we will underestimate investment flexibility. In Appendix A.3.4, we discuss this issue in more detail. In practice, we do not believe that this bias meaningfully affects our empirical results.

2.3 Three Facts on Investment Plans

We now present three stylized facts regarding firm investment plans.

2.3.1 Fact 1. Investment Plans Contain Incremental Information Not Captured in Expected Future Performance Alone

To study the importance of investment plans relative to other common alternative explanatory variables (in particular, expected profitability), we run regressions of the form:

$$\frac{i_{i,t}}{k_{i,t}} = \alpha_{j(i),t} + \beta \frac{i^p_{i,t}}{k_{i,t}} + \gamma x^{t+1}_{i,t} + X'_{i,t-1}\Gamma + \varepsilon_{i,t}, \tag{1}$$

where i indexes a firm, $j(i)$ indexes the industry firm i is in, and t is a fiscal year; $X_{i,t-1}$ are control variables including log total assets, cash to total assets, and the long-term book leverage ratio; and $\alpha_{j(i),t}$ represents industry-by-year fixed effects. The regression specification compares firms in the same industry and year, in which the coefficients of interest are β, which captures the relation between investment plans and realized investments, and γ which studies the relation between investment i_t and an explanatory expected future period $t+1$ performance variable $x^{t+1}_{i,t}$. For our performance variable, we consider as alternatives expected sales relative to capital, expected value-added relative to capital, and expected

TFP. We also run the same regressions without investment plans (i.e., dropping $i^p_{i,t}$) to show the benchmark result of the performance variable alone. This set of specifications allows us to essentially interpret the regressions as a horse race between investment plans and $x^{t+1}_{i,t}$ in determining actual investment.

[Table II Here]

We report our results in Panel B of Table II, in which Columns 1, 3, and 5 report regressions on log expected sales divided by capital, log expected value added divided by capital, and log expected TFP per the definition in Section 2.2.1, while Columns 2, 4, and 6 report regressions that include investment plans as a regressor. A sharp and consistent result emerges: Investment plans are statistically significantly related to actual capital expenses even after controlling for these conventional variables. Moreover, the presence of investment plans *attenuate* the importance of these expected performance variables by over two-thirds and improve R^2 more than 4-fold. For instance, in comparing Columns 5 and 6 in Panel B, we find that the coefficients on expected TFP decrease by over 75% and R^2 increases more than 5-fold from 0.080 to 0.473 when adding investment plans as a regressor.

Our results are surprising: Conventional firm dynamics models would predict that investment plans will have no predictive power for actual investment once we control for expected TFP, since any correlation between actual and planned investment is driven by the firm's expectations of productivity. Moreover, our results are not simply due to our measure of expected TFP, as similar results obtain when simply considering alternative measures that require no structural functional form assumptions.[10] Therefore, investment plans appear to statistically and economically significantly explain realized investments, beyond acting as a conduit of expected sales, expected value added, or expected TFP.

These findings are also consistent with investment plans being partially flexible. On one hand, plans are informative of actual capital investment, since the coefficient is statistically

[10] Our regressions also control for firm liquidity and asset levels, which are important determinants of firm investment decisions, as emphasized by the large firm dynamics literature. This means that our investment plans variable is not simply picking up the effects of omitted variables.

different than zero. On the other hand, if plans were completely irreversible, we would expect an estimate very close to unity given that realized investments cannot deviate from plans. In the next subsection, we test the flexibility of investment plans more directly.

2.3.2 Fact 2. Investment Plans Are Partially Flexible

To test the flexibility of investment plans, we consider a regression specification of the form

$$\Delta \frac{i_{i,t}}{k_{i,t}} = \alpha_{j(i),t} + \beta \Delta \log z_{i,t} + X'_{i,t-1}\Gamma + \eta_{i,t}, \tag{2}$$

where the indices $i, j(i)$, and t, as well as the set of control variables $X'_{i,t-1}$, follow Equation 1. The coefficient of interest β tells us how much deviation from planned investment is possible in response to TFP shocks. Standard errors are clustered by firm.

Table III reports the empirical results to document the flexibility of investment plans, dropping the i, t subscript when there is no confusion. Columns 1, 3, and 5 study how investment deviations $\frac{\Delta i_{i,t}}{k_{i,t}}$ respond to performance shocks with respect to sales shocks, value-added shocks, and TFP shocks, respectively. We find that firms deviate from investment plans in response to all three measures of performance shocks in conventional ways. For instance, compared with firms in the same industry, those that receive a more positive TFP shock ramp up investment relative to their plans, and those that receive more negative TFP shocks scale down relative to their plans. Taken together, this suggests that deviations from investment plans are possible.

[Table III Here]

A concern with our specification is that our measure of investment plan deviations Δi partially captures the endogeneity of an investment plan itself to TFP shocks—which effectively introduce a reverse causality problem—rather than the variation in realized investment responses themselves.[11] This counterfactual is plausible if firms have some private informa-

[11] For instance, influential work by Jaimovich and Rebelo (2009); and Schmitt-Grohé and Uribe (2012) concludes that "news shocks" — whereby firms have advanced information on future innovations—are important drivers of investment decisions.

tion that is not reflected in our expected TFP measure. To alleviate this concern, Columns 2, 4, and 6 in Table III consider a slightly different regression specification that uses realized investment $\frac{i}{k}$ as a dependent variable and investment plans $\frac{i^p}{k}$ as a control to more directly study changes in actual investment responses. We show that actual investment indeed responds to TFP shocks and investment plans remain statistically and economically significant. Our findings are robust to studying both sales and value-added shocks as the performance shock rather than TFP. Therefore, investment plan deviations do appear to capture intra-year deviations from initial budgets.

We now return to addressing two earlier concerns. First, we noted in Section 2.2.1 that our TFP shocks measure is biased. However, the fact that we mismeasure TFP shocks is not a major threat to our empirical results, given the robustness of the qualitative correlations to using sales or value-added shocks, similarly defined. Since these definitions require fewer assumptions than TFP, we believe our empirical results are not driven solely by the bias in the TFP shock measurement. Second, we noted in Section 2.2.2 that our measure of investment deviation is biased toward 0, which in turn implies an underestimation of investment flexibility. As our results show, even with this bias in place, actual investment is still responsive to TFP shocks, which implies that the bias might not be that severe. Therefore, we believe that our two measures, while imperfect, are not mechanically biasing our empirical results toward our conclusions.[12] Regardless, to directly account for these empirical concerns for our counterfactual analyses, we will address both concerns simultaneously when calibrating our model in Section 5.

2.3.3 Fact 3. Deviations from Capital Budgets Are Costly

Finally, we examine whether investment plans are only partially flexible, in part because it is costly to deviate from the capital budget. To show this relation, we study whether a firm's future profitability relates to its capital budget deviations using a regression of the following

[12] Both stylized facts are robust to accounting for firm fixed effects (which will also account for heterogeneity in firm-level $\mathbb{V}$). However, we prefer the specification with no firm fixed effects to maintain a close connection to our model.

form:

$$y_{i,t+1} = \alpha_{j(i),t} + \beta \log\left(1 + |\Delta \frac{i_{i,t}}{k_{i,t}}|\right) + \gamma i_{i,t} + \phi \log z_{i,t} + \psi(\Delta \log z_{i,t-1}) + X'_{i,t-1}\Gamma + \varepsilon_{i,t+1}, \tag{3}$$

where the indices $i, j(i)$, and t, as well as the set of control variables $X'_{i,t-1}$, follow Equation 1. The outcome variable $y_{i,t+1}$ is a firm's future gross profit margins, defined as the ratio of ordinary profits to sales. Our empirical specification studies the impact of investment plan deviations from planned investment on firm's bottom line, holding fixed all other relevant firm characteristics. We effectively compare two otherwise identical firms in the same year in the same industry with similar fundamentals (including realized investment rates), for which the only difference is their initial investment plans (and hence, with the same actual investment choices, different levels of deviation). By controlling for actual realized TFP, we also control for actual outcomes, relating any changes in profitability to plan deviations.

Table IV reports our results. Column 1 controls for realized TFP, Column 2 also controls for TFP shocks, and Column 3 includes additional firm-level controls. Regardless of the suite of control variables, the deviations in investment choices from plans — whether it is necessary to ramp up or scale down — reduces future profit margins. These results are economically meaningful. For example, Column 1 shows a one-percentage-increase in deviations leads to a 0.3-percentage-point decrease in gross margins. In our data, average sales are around US$ 670 million while the average profit margin is 6.70% of sales. Taken together, this therefore amounts to forgone profits of around US$ 2 million.[13] Importantly, in Column 3, we also control for firm total assets with quantitatively similar effects. The robustness of our point estimates suggests that our investment plan deviations measures are not simply picking up the effect of physical adjustment costs.

[Table IV Here]

[13]This back-of-the-envelope calculation assumes that investment plan deviations only affect profits rather than sales, which is consistent with our model in the next section.

3 Model

In this section, we present a parsimonious model of firm dynamics motivated by the stylized facts from Section 2. We consider a discrete time, infinite-horizon economy, populated by a representative household, a representative final goods firm, and heterogeneous intermediate goods firms. Our model features two key elements: (i) endogenous signal acquisition and (ii) costly deviation from investment plans. We will use this as our baseline economic environment to study the aggregate implications of partially irreversible investment plans on TFP forecast errors, investment plan deviations, and capital allocation.

3.1 Households

The household is infinitely lived, discounts time at rate β, and owns all of the firms in the economy. It inelastically supplies a fixed quantity of labor $N = 1$ and has preferences over consumption of a final aggregate consumption good. The household plays a limited role in our analysis, but is presented for completeness.

3.2 Final Goods Firm

There is a representative final goods firm that aggregates all the intermediate goods y_i. Aggregate output is a constant elasticity of substitution (CES) aggregate over a measure 1 of different varieties of good $Y_t = \left(\int_0^1 y_{i,t}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}$, where i is a generic variety and η is the elasticity of substitution across goods. The usual cost minimization problem for the final goods firm yields the standard demand schedule for each good i as $p_{i,t} = y_{i,t}^{-\frac{1}{\eta}} P_t Y_t^{\frac{1}{\eta}}$, where $P_t \equiv \left(\int_0^1 p_{i,t}^{1-\eta} di \right)^{\frac{1}{1-\eta}}$ is the usual CES price index, giving us $P_t Y_t = \int_0^1 p_{i,t} y_{i,t} di$. For the rest of the paper, we will set the final consumption good as the numeraire and normalize P to 1.

3.3 Intermediate Goods Firms

The economy is populated by a continuum of intermediate goods firms indexed by $i \in [0, 1]$. Firms are infinitely lived and are run by managers who discount future utility at a constant rate β. Managers derive utility over dividend flow but also have preferences over how actual investment deviates from their plans. In the interest of clarity, we delay further discussion of manager preferences until the rest of the economic environment has been described.

3.3.1 Production

Intermediate goods firms produce a differentiated good using the production function $y_{it} = z_{it}k_{it}^{\alpha}l_{it}^{1-\alpha}$, where z_{it} is idiosyncratic stochastic productivity, k_{it} is the beginning-of-period capital stock, l_{it} is labor hired by the firm, and $\alpha \in (0,1)$ is the capital share of the firm. We assume that $\log z_{it}$ follows an AR(1) process given by $\log z_{it} = \rho \log z_{i,t-1} + \epsilon_{it}$, where $\epsilon_{it} \sim N\left(0, \sigma_{\epsilon}^{2}\right)$. In addition, capital depreciates at a geometric rate δ. Finally, each unit of labor costs a wage w, and we will assume that l_{it} is chosen after z_{it} has been observed. Under these assumptions, without loss of generality, we can directly rewrite the firm's gross profit function net of labor cost as

$$\pi = \mathcal{A}\left(w, Y\right) z_{it}^{\Theta_z} k_{it}^{\Theta_k}, \tag{4}$$

where $\Theta_z \equiv \frac{\frac{\eta-1}{\eta}}{1-(1-\alpha)\frac{\eta-1}{\eta}}$, $\Theta_k \equiv \frac{\alpha\frac{\eta-1}{\eta}}{1-(1-\alpha)\frac{\eta-1}{\eta}}$, and $\mathcal{A}\left(w, Y\right) > 0$ is a function of the endogenous aggregate wage and output.

3.3.2 Capital Budgeting

An important element of the model is that, as a model primitive, we will assume firm managers have to make investment *plans* $k_{i,t+1}^{p}$ prior to making an actual investment to achieve the next-period capital stock $k_{i,t+1}$. While we do not take a stand on exactly why firms have to make investment plans, both the corporate finance literature and our empirical results show that firm managers in fact do care strongly about making correct plans, and that any need for ex post adjustments is costly.[14] We assume that given some plans and actual investment, the firm managers face the following cost function $\phi(k_{i,t+1}^{p}, k_{i,t+1})$ for any

[14] In practice, capital budgeting is a core process in the allocation of capital within a firm, and any deleterious effects of poor capital budgeting could arise for a variety of reasons (e.g., Mao 1970; Myers 1974; Schall et al. 1978; Arnold and Hatzopoulos 2000; Ryan and Ryan 2002), which is outside the scope of this paper.

deviation from the plan:

$$\phi\left(k_{i,t+1}^{p}, k_{i,t+1}\right) = \frac{\chi}{2}\left(\frac{k_{i,t+1}}{k_{i,t+1}^{p}} - 1\right)^{2} k_{i,t+1}^{p}, \tag{5}$$

where $\chi \geq 0$ denotes the severity of the cost function and $\chi = 0$ implies that plans are irrelevant. We discuss the importance of this assumption through the lens of the information structure and model timing below.[15]

3.3.3 Information and Timing

Why do firm managers not simply set $k_{i,t+1}^{p} = k_{i,t+1}$? In a standard firm dynamics model such as Hopenhayn (1992), firms perfectly observe contemporaneous productivity z_{it}, so all planned investments equate to realized investments. However, as discussed in the empirical section, firms generally do not perfectly observe their current productivity when they are making plans. Instead, at the beginning of a period, they forecast the sales (and hence TFP) for the year and make operating expenditure plans accordingly based on currently available information from the past period's productivity. Then, given the plans, they are able to partially deviate as the year progresses and more information is revealed. To account for this fact, we will need to adopt a different timing assumption. To keep the model as simple as possible but still have this intra-year feature, we split a time period into "day" and "night" subperiods.

Day In a given period, during the day, the manager only has information on the previous period's productivity $z_{i,t-1}$ (which was realized at the previous period at night) but does not observe ϵ_{it}. However, the firm manager is able to improve on her information by acquiring signals with precision $1/\sigma$. For some choice of σ, the manager will receive some signals s_{it}

[15] This assumption on costly plan deviation can be microfounded by operational inefficiencies such as those modeled by Harris and Raviv (1996) and Malenko (2019), which feature information asymmetry within a firm and costly verification (auditing) of spending within a firm. In both models, the constrained optimal capital budgeting rule is to allocate a planned amount for capital expenditure, then incur costly verification for realized spending that deviates from plans. In this framework, costs are due to information frictions in the decentralized organization, such as auditing or various meetings.

about current productivity, given by

$$s_{it} = u_{it} + \epsilon_{it},$$

where $u_{it} \sim N(0, \varsigma^2)$ and $\varsigma^2 \equiv \left(\frac{1}{\sigma^2} + \frac{1}{\sigma_u^2}\right)^{-1}$. The parameter σ_u^2 is the intrinsic upperbound to the uncertainty regarding current productivity, while σ^2 is the amount of uncertainty the manager chooses to reduce relative to this upperbound.[16] In order to improve on her information about the current period productivity, the firm manager is able to acquire better signals at a cost. Specifically, for some choice of σ, the manager bears a utility cost

$$\mathcal{C}(\sigma) = \xi \times \left(\frac{1}{\mathbb{V}} - \left(\frac{1}{\sigma_u^2} + \frac{1}{\sigma_\epsilon^2}\right)\right),$$

where $\mathbb{V} \equiv \left(\frac{1}{\varsigma^2} + \frac{1}{\sigma_\epsilon^2}\right)^{-1}$. This assumption implies that the marginal cost of improving signal precision is constant and homogeneous across firms.

After observing the signal, the manager is able to update her belief over the posterior distribution of $z_{i,t}$ using Bayesian updating. Specifically, her belief follows

$$\log \tilde{z}_{it} \sim N(\log \hat{z}_{it}, \mathbb{V}),$$

where $\log \hat{z}_{it} \equiv \rho \log z_{i,t-1} + \frac{\sigma_\epsilon^2}{\varsigma^2 + \sigma_\epsilon^2} s_{it}$ is the expected current period productivity and $\mathbb{V}$ is the posterior variance of productivity (equivalently, $\mathbb{V}^{-1}$ gives us the precision of the signal) from above. In contrast, a standard timing whereby $\sigma = 0$ for all firms results in a posterior distribution of $\log z_{it}$ that is degenerate (all firms know exactly what their current productivity is). The manager then makes an investment plan, taking into account the adjustment cost function $\phi(k_{i,t+1}^p, k_{i,t+1})$ in Equation 5.

[16] Notice that with our specification, $\varsigma \in (0, \sigma_u)$, where $\varsigma \to 0$ when $\sigma \to 0$ and $\varsigma \to \sigma_u$ when $\sigma \to \infty$. Moreover, u_{it} becomes degenerate at 0 when $\varsigma \to 0$; in other words, this reduces to the standard timing whereby firms perfectly observe current productivity. Finally, note that σ_u refers to the amount of subjective uncertainty that firms face over z, and is not the volatility of z.

Night During the night, all uncertainty about current productivity is resolved, and the manager perfectly observes z_{it}. In addition, the manager also observes a signal about *future* innovations to productivity $\epsilon_{i,t+1}$, given by

$$\tilde{s}_{i,t+1} = \tilde{u}_{i,t+1} + \epsilon_{i,t+1},$$

where $\tilde{u}_{i,t+1} \sim N\left(0, \sigma_u^2\right)$. In other words, prior to picking some level of signal clarity, the manager already has some knowledge about future productivity "for free". Given this information, the manager's posterior belief about future productivity $z_{i,t+1}$ is given by

$$\log \tilde{z}_{i,t+1} \sim N\left(\rho \log z_{i,t} + \frac{\sigma_\epsilon^2}{\sigma_u^2 + \sigma_\epsilon^2}\tilde{s}_{i,t+1}, \left(\frac{1}{\sigma_u^2} + \frac{1}{\sigma_\epsilon^2}\right)^{-1}\right).$$

At this point, the manager uses this refined information to decide how much capital to invest into the next period, subject to the aforementioned adjustment cost $\phi(k^p_{i,t+1}, k_{i,t+1})$ in Equation 5. At this point, we see that σ_u arises as a natural upperbound for the amount of uncertainty in the day.[17]

3.4 Bellman Equations

We can now define the problem recursively. Let $J(k, k^p, z)$ denote the value function of the manager after all shocks have been realized (i.e., at *night*), $W(k, s, z_{-1})$ the value function of the manager after a signal has been observed, and finally $V(k, \tilde{z}, z_{-1})$ the value function of the manager at the beginning of the period (i.e., *day*).

The manager's Bellman equation at night can be written as

$$J\left(k, k^p, z\right) = \max_{k'} \pi + \left(1-\delta\right)k - k' - \frac{\chi}{2}\left(\frac{k'}{k^p} - 1\right)^2 k^p + \beta\mathbb{E}\left[V\left(k', \tilde{z}', z\right)|z, \tilde{s}'\right]$$
$$s.t. \quad \log \tilde{z}' \sim N\left(\rho \log z + \frac{\sigma_\epsilon^2}{\sigma_u^2 + \sigma_\epsilon^2}\tilde{s}', \left(\frac{1}{\sigma_u^2} + \frac{1}{\sigma_\epsilon^2}\right)^{-1}\right)$$

[17] Notice that in terms of the evolution of the manager's information set, when $\sigma = 0$ for all (z_{-1}, k), our model collapses to that in David and Venkateswaran (2019). Moreover, σ_u corresponds to their exogenously imposed uncertainty parameter.

when next-period capital k' is chosen by the manager to maximize her expected utility, taking into account that deviations from investment plans are costly. Therefore, the Bellman equation after the signal has been observed is given by

$$\begin{aligned} W(k, s, z_{-1}, \sigma) &= \max_{k^p} \mathbb{E}\left[J(k, k^p, z) \,|s, z_{-1}, \sigma\right] \\ s.t. \quad \log z &\sim N\left(\rho \log z_{-1} + \frac{\sigma_\epsilon^2}{\varsigma^2 + \sigma_\epsilon^2} s, \mathbb{V}\right) \end{aligned}$$

when the manager chooses investment plan k^p, taking into account the potential future need to deviate from the plan. The manager has to set k^p accounting for the full distribution of possible investment choices, since she does not know at this point what current productivity is. Finally, in anticipation of the post-signal and night period, the Bellman equation in the day is given by

$$\begin{aligned} V(k, \tilde{z}, z_{-1}) &= \max_{\sigma} -\xi\left(\frac{1}{\mathbb{V}} - \left(\frac{1}{\sigma_u^2} + \frac{1}{\sigma_\epsilon^2}\right)\right) + \mathbb{E}\left[W(k, s, z_{-1}, \sigma)\right] \\ s.t. \quad & s = u + \epsilon \\ & u \sim N\left(0, \varsigma^2\right), \end{aligned}$$

noting that $\varsigma^2 \equiv \left(\frac{1}{\sigma^2} + \frac{1}{\sigma_u^2}\right)^{-1}$. Here, the manager chooses the signal quality of today's productivity σ, whereby a more precise signal is costlier to acquire but adds value at night, since this reduces the need for ex post deviations from investment plans. The manager knows the productivity from the previous period (z_{-1}), since we assume that all uncertainty regarding productivity is resolved at night.

3.5 Equilibrium Definition

A stationary competitive equilibrium is defined by a measure of firms Λ; a set of policy functions $\{\sigma(z_{-1}, k; w, Y), k^p(z_{-1}, k, s; w, Y), k'(z_{-1}, k, s, z; w, Y)\}$; a set of value functions $\{V, W, J\}$; intermediate good prices $\{p_{i,t}\}_{i \in \Lambda}$; wage w; and a Markov transition function Γ induced by the policy function k' and exogenous productivity z' such that (1) the distribution of firms is invariant; namely, $\Lambda = \Gamma(\Lambda)$, (2) the labor market clears, and (3) the intermediate goods market clears.

4 Forecast Errors and Firm Characteristics

Before moving on to our main quantitative results in Section 5, we focus our discussion here on how our model relates back to our initial stylized facts. We then show how our model derives sharp predictions that connect observable firm characteristics to the size of their forecast errors, following which we directly test these predictions using our panel data. Finally, we connect our model implications to a recent literature on behavioral corporate finance.

4.1 Matching Stylized Facts

To begin, recall the observation that firm investment plans are partially flexible. This effect is taken as given in our model as a primitive in the form of the adjustment cost function for investment plans. However, it is less clear how the adjustment cost parameter χ maps investment plans onto actual investment. To do so, we formally state our first result below.

Claim 1. *The optimal investment plan k^p depends on both expected productivity $\hat{z}$ and the posterior distribution of productivity $\mathbb{V}$.*

We derive an explicit solution for k^p in Appendix B, showing that k^p is indeed a function of both $\hat{z}$ and $\mathbb{V}$. We list this claim here mainly to show that investment plans are not simply a conduit for expected TFP, but rather also reflect the amount of uncertainty under which these plans are made. In fact, as we show later, firms with the same expected TFP can have very different plans depending on their choice of information acquisition. Therefore, both investment plans and expected TFP are jointly significant determinants of actual investment, as we documented in Section 2.

Claim 2. *The optimal next-period capital choice k' is a function of realized productivity z and planned investment k^p.*

This claim results directly from the Euler equation (see Appendix B). An important implication here is that investment plans, in the form of k^p, become a relevant state variable for predicting actual investment. Unfortunately, k' is a nonlinear function of z and k^p and is defined implicitly by the Euler equation. To make headway in deriving a sharper characterization, we proceed to the next claim.

Claim 3. *Under a log-linear approximation around the nonstochastic steady-state, k' is given by the following equation:*

$$\log k' = (1 - \phi_k) \log \bar{k} + \phi_z \log z + \phi_k \log k^p, \tag{6}$$

where $\phi_z \equiv \frac{\Theta_z \rho (r+\delta)}{(1+r)\chi + (r+\delta)(1-\Theta_k)}$ and $\phi_k \equiv \frac{(1+r)\chi}{(1+r)\chi + (r+\delta)(1-\Theta_k)}$, and $\bar{k} = \left[\frac{\Theta_k \mathcal{A}(w,Y)}{r+\delta}\right]^{\frac{1}{1-\Theta_k}}$ is the nonstochastic steady-state capital holdings of the firm.

Appendix B derives the preceding claim.[18,19] Here, we see that the weight on current productivity ϕ_z decreases in the adjustment cost parameter χ, whereas the weight on they ϕ_k increases in χ. Essentially, this result shows that the more irreversible investment plans are, the more powerful investment plans will be as a predictor of actual firm investment, and therefore the more correlated these two variables will be. In contrast, if $\chi = 0$ (reducing the model to the standard timing), investment plans would have no predictive power for actual investment once (expected) productivity is properly controlled for. Moreover, it is clear to see that the dispersion of investment errors is decreasing in χ, holding all else constant, since investment is more flexible.

Therefore, our model set up rationalizes the three stylized facts presented in Section 2. This outcome is unsurprising, as the facts informed our model setup. In the next section, we will derive additional implications of the relation between the dispersion of investment errors, TFP errors, and the costliness of deviating from plans to quantitatively inform how flexible investment is.

[18]While we express the policy function in terms of next-period capital stock, note that since investment is simply $i = k' - (1-\delta)k$, all of our claims here will carry through.

[19]While we log-linearize k' around the nonstochastic steady-state, the investment plan and signal acquisition policy functions remain general nonlinear functions of the natural state variables. We have also numerically solved our model using global methods to verify the accuracy of our approximation. This is not surprising, given that in a frictionless environment, k' is exactly log-linear in z.

4.2 Other Key Relationships

We present two additional testable predictions as a means to validate our model. We first present these model predictions and discuss their intuition, then test the predictions empirically. Appendix B contains all relevant formal proofs.

Proposition 1. *Under a log-linear approximation for next-period capital choice k',*

1. *the ex ante (day) expected value of the firm, gross of signal acquisition cost, is strictly decreasing in the posterior uncertainty $\mathbb{V}$.*

2. *the value of increasing signal precision (i.e., decreasing $\mathbb{V}$) is increasing in initial firm productivity z_{-1}.*

This result leads to the following corollary.

Corollary 1. *The optimal signal precision is increasing in initial firm productivity.*

Heuristically, by building directly from Proposition 1 and our assumption on the cost of signal acquisition being monotone and increasing in $\frac{1}{\mathbb{V}}$, we see that the optimal signal precision must be increasing in z_{-1}. We now make two predictions unique to our model:

Prediction 1. *The dispersion of forecast errors for firm TFP is decreasing in initial firm productivity.*

Prediction 2. *The dispersion of forecast errors for investment rates is increasing in the absolute size of realized TFP shocks.*

It is useful to take stock of the economic mechanism behind our proposition and the resulting predictions. Why does the value of the firm decrease in the posterior uncertainty?[20] To understand this, recall that the value of the firm is maximized when expected internal and external rates of return to capital are equal, and that any "wedge" between them reduces

[20]This prediction contrasts with the usual Abel-Hartmann-Ooi effect, which predicts that increases in real uncertainty are often associated with an increase in the value of the firm.

firm value. With this in mind, we can rewrite the Euler equation in our model as

$$\left(\mathcal{M}^e - (r+\delta)\right)^2 = \Bigg(\underbrace{(1+r)\,\chi\frac{k'}{k^p} - (1+r)\,\chi}_{\equiv\tau}\Bigg)^2, \tag{7}$$

where $\mathcal{M}^e \equiv \exp\left(\frac{1}{2}\Theta_z^2\sigma_\epsilon^2\right)\Theta_k\mathcal{A}\left(w\right)z^{\Theta_z\rho}k'^{\Theta_k-1}$ is the expected next-period marginal revenue product of capital, and τ is interpreted as an investment wedge. By squaring both sides, we have the interpretation of a loss function: Larger absolute values of the wedge are associated with larger reduction in firm value. We can then express the ex ante expected loss in the day period as

$$\mathbb{E}\left[\tau^2|\sigma\right] = \left((1+r)\,\chi\right)^2\left(2 - 2\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\right), \tag{8}$$

where we see that the expected loss increases in the posterior uncertainty. Intuitively, larger uncertainty means the probability of a manager needing to make large adjustments is higher, since she is more likely to make a mistake in her plans. Since adjustment is costly (as captured by χ), this means that she is not able to fully correct her mistake. The combined effect leads to a decrease in the expected value of the firm.

The preceding discussion is simply a heuristic and does not explain why the value of increasing signal precision is increasing in z_{-1}. In Appendix B, we show that the marginal benefit of increasing signal precision (increasing $\mathbb{V}^{-1}$) can be explicitly expressed in the following form:

$$\bar{k}\underbrace{z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}}_{\text{scaling effect}}\Bigg[\underbrace{\chi\left(\mathbb{V}^2\frac{\partial\mathcal{F}^A\left(\mathbb{V}\right)}{\partial\mathbb{V}}\right)}_{\text{expected ex-post adjustment needed}} + \underbrace{\frac{1}{1+r}\exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right)\left(-\mathbb{V}^2\frac{\partial\mathcal{F}^\pi\left(\mathbb{V}\right)}{\partial\mathbb{V}}\right)}_{\text{expected continuation profits}}\Bigg],$$

which decomposes the marginal effect of being "more correct" on the expected magnitude of ex post adjustment that has to be done ($\frac{\partial\mathcal{F}^A(\mathbb{V})}{\partial\mathbb{V}}$) and the expected loss of profit (gross of adjustment cost) due to mistakes being made ($\frac{\partial \mathcal{F}^{\pi}(\mathbb{V})}{\partial \mathbb{V}}$).[21] As we show in the Appendix, $\frac{\partial \mathcal{F}^{A}(\mathbb{V})}{\partial \mathbb{V}} > 0$ and $\frac{\partial \mathcal{F}^{\pi}(\mathbb{V})}{\partial \mathbb{V}} < 0$. This effect relates to our earlier discussion of the Euler equation wedge.

[Figure I Here]

Notice that the term z_{-1} does not enter into either $\mathcal{F}^{A}(\mathbb{V})$ or $\mathcal{F}^{\pi}(\mathbb{V})$. In other words, z_{-1} has a pure scaling effect that amplifies the marginal benefit of being correct. As an example, in Figure I we plot the marginal benefit of increasing signal precision for two levels of productivity against the marginal cost. The figure shows that z_{-1} scales up the marginal benefit of improving signal precision. As a result, higher productivity firms prefer to acquire better signals. In this sense, our model captures the economic intuition that high-productivity firms have more to gain from being "correct."[22]

A final takeaway is that any heterogeneity that arises from information acquisition is due only to the returns on learning being higher for higher productivity firms, since we assumed homogeneous and constant marginal cost of improving signal precision.[23] In contrast, recent studies such as Tanaka et al. (2020) — who also find that higher firm productivity correlates with lower forecast error dispersion empirically — postulate that this arises from better managerial ability. This hypothesis implies that higher productivity firms face a lower marginal cost of learning. Our model clarifies that such a finding does not necessarily require decreasing marginal costs. That said, a decreasing marginal cost hypothesis is complemen-

[21] $\mathcal{F}^{A}(\mathbb{V})$ and $\mathcal{F}^{\pi}(\mathbb{V})$ are complicated functions of model parameters and $\mathbb{V}$, and are defined in Appendix B in the interest of space. The marginal benefit of improving signal precision $\mathbb{V}^{-1}$ due to reduced ex post adjustment is in fact given by $-\chi\frac{\partial \mathcal{F}^{A}(\mathbb{V})}{\partial \mathbb{V}^{-1}}$, while that from improved capital allocation is $\frac{1}{1+r}\exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right)\left(\frac{\partial \mathcal{F}^{\pi}(\mathbb{V})}{\partial \mathbb{V}^{-1}}\right)$, which can be rearranged to the terms above when we use the fact that $d\mathbb{V}^{-1} = -\mathbb{V}^{-2}d\mathbb{V}$.

[22] To put our model mechanism in the context of a real life example, consider the impact of capital budgeting mistakes on the profitability of a large grocery store chain as compared to a "mom-and-pop" corner store. Our model predicts that the same relative mistake (i.e., same percentage investment plan deviation) will naturally have a larger *level* impact on the bottom line of the large chain as compared to corner store. This difference then leads the large chain to preemptively make better investment plans (for instance, by engaging external consultants) relative to the corner store.

[23] Notably, we also do not "bake in" this result on the benefit side of the equation. As briefly discussed earlier, the marginal cost of relative investment plan deviation is homogeneous across firms.

tary to our analysis and suggests that future work could be done to quantify the importance of these two channels.

4.3 Empirical Evidence

We now test the two predictions in our data. In the following analyses, the outcome variables are dispersions of either TFP or investment forecast errors, and regressors are either lagged or concurrent TFP forecast error dispersions, corresponding to Panels A and B in Table VI which test predictions (1) and (2) respectively.

Analyses that relate forecast error dispersions use the log of one plus the absolute level of forecast error. There are three advantages of making this choice. First, this transformation accounts for the fact that the absolute errors follow a distribution with long tails and therefore render our statistical inferences more reliable. Second, it preserves all observations with a forecast error including those with forecast error equal to zero. Third, it permits an interpretation similar to an elasticity of dispersion interpretation of our results.[24] We present analyses using two levels of variation: (1) unconditional pooled regressions and (2) cross-sectional regressions that compare firms with others in the same industry and year. We show both levels of variation, since the intuition of our model applies to all levels of analyses, and indeed we broadly find results consistent across all sources of variation that we consider.

Prediction 1. *Dispersion of forecast errors for firm productivity ("TFP") is decreasing in initial firm productivity.* Panel A in Table VI shows the empirical relation between TFP shocks and the previous period's realized TFP. We find that higher realized TFP in the previous fiscal year is correlated with smaller TFP shocks at both levels of variation. The results in all regressions are statistically significant at the 1% level.

Prediction 2. *Dispersion of forecast errors for investment rates is increasing in the absolute size of realized TFP shocks.* Panel C in Table VI shows that the dispersion of investment

[24] All of the empirical results are robust to a battery of alternative transformations, including the inverse hyperbolic sine transformation, log absolute errors in which we drop firm-year observations with numerically zero forecast error, and simply the absolute errors themselves with no transformation for the whole sample.

deviations increase in TFP error dispersions. The results all regressions are statistically significant at the 1% level.

Altogether, these results are consistent with the predictions in Section 4.2, which are unique to our model with endogenous learning and partial flexibility of investment plans. In untabulated analyses, we find that these results also obtain (i) for Prediction 1 when replacing TFP shocks with sales shocks, as defined in Table III, as well as when considering lagged sales relative to capital rather than TFP as the regressor; and (ii) for Prediction 2 when replacing TFP shocks with sales or value-added shocks, as defined in Table III. Therefore, we argue that our economic mechanism is empirically verified and robust. We build on the empirical findings to deduce additional quantitative implications in the next section.

[Table VI Here]

4.4 Relationship to the literature on behavioral corporate finance

We conclude this section by connecting our model to a growing literature in corporate finance that has emphasized the importance of accounting for managerial forecasting biases in determining investment or hiring decisions (e.g., Bordalo et al. (2018, 2020); Ma et al. (2020); Barrero (2021)). This emphasis builds on the multitude of empirical research documenting the predictibility of forecast errors at the idiosyncratic level (e.g., managers or professional forecasters), a finding that more traditional models of full information rational expectations (FIRE) with information rigidity (e.g., Woodford (2003); Reis (2006)) have difficulty matching (Coibion and Gorodnichenko (2015)).

In this context, our paper is more similar to standard FIRE models with one exception: By explicitly incorporating capital budgeting frictions, our model predicts that measured forecast errors are predictable at the idiosyncratic level. For instance, using sales forecast errors (in logs) as an example, we can derive the formula

$$FE_{i,t} \equiv \log sales_{i,t} - \log \mathbb{E}_{i,t-1} sales_{i,t} = \Theta_z \Delta \log z_t - \frac{1}{2}\Theta_z^2 \mathbb{V}\left(z_{i,t-1}\right),$$

where $\mathbb{E}_{i,t-1} sales_{i,t}$ is forecasted sales using last period's information, and $\Delta \log z_t$ is the time-t TFP "shock" and is unforecastable, but the measured forecast error $FE_{i,t}$ *is* forecastable due to the $\mathbb{V}$ term. As we already discussed, firm uncertainty today (i.e., $\mathbb{V}$) depends on lagged productivity due to capital budgeting frictions, which in turn makes $\mathbb{V}$ autocorrelated and forecastable.[25] Note that this result is different from existing firm dynamics models (e.g., David and Venkateswaran (2019)), for which $\mathbb{V}$ would be homogenous across firms and $FE_{i,t}$ would correspondingly be uncorrelated.

5 Quantitative Analysis

In this section, we discuss the quantitative and theoretical implications of our model. We first calibrate our model to identify moments in the data and compute the dollar costs of these two frictions. Subsequently, we compare and contrast the different strategies different types of firms use to mitigate uncertainty when formulating capital budgets. Finally, we discuss how our results relate to firm uncertainty as a source of dynamic misallocation. In this vein, we relate our model findings to broader results in the literature on capital misallocation.

5.1 Calibration

We calibrate our model in general equilibrium. We set $\beta = 0.98$, consistent with an annual real interest rate of 2%, and $\delta = 0.06$; and assume $N = 1$ and $P = 1$ as a normalization. Following standard assumptions, we set $\eta = 4$ and $\sigma_u = \infty$ (that is, firms do not have any additional prior information about future productivity). Finally, we set $(\sigma_\epsilon, \rho, \alpha)$ equal to (0.358, 0.903, 0.280), respectively, based on estimates directly from the data.[26]

[25] As an example, running an AR(1) regression using simulated sales forecast errors from our calibrated model in Section 5 yields a coefficient of around 0.06.

[26] For σ_ϵ and ρ, we fit an AR(1) regression to our measure of TFP with industry and year fixed effects; for α, having imposed $\eta = 4$, we compute the median labor share and back out α via the firm's first-order conditions (See Section 2 and Appendix A.3 for more details).

This leaves two key parameters: the cost of information acquisition ξ and the cost of deviation of investment plans χ. We jointly calibrate ξ and χ to two identifying moments: (i) dispersion of TFP shocks relative to the volatility of TFP and (ii) the dispersion of investment deviations.[27] Finally, we solve for the endogenous wages and aggregate output, such that these values are consistent with our market-clearing conditions (i.e., aggregate labor demand and output is consistent with $N = 1$ and $P = 1$), as well as the calibrated values for ξ and χ.

We note one important concern in calibrating our model pertaining to the empirical measures of TFP and investment. In Section 2, we explained that our empirical measures of TFP shocks and investment plan deviations are mismeasured due to the exact variable definitions in the survey. To address the issue of constructing expected TFP from expected sales, we recreate this exact same bias in our calibration step. When computing TFP shocks in our model, we compute expected TFP as we would in the data. For the truncation bias in our observed gross investment expenditures, for any firms that do negative investment (or planned on doing negative investment), we replace I and I^p with zeros when calibrating.

From our simulated method of moments calibration, we find that $(\chi, \xi) = (0.212, 8.76 \times 10^{-5})$ best fits our data.[28] Panel A of Table V shows our calibration results, which demonstrates a perfect model fit. We provide further validation of our model by running the same regressions as our earlier specifications for Predictions 1 and 2 in Section 4, and comparing the model-predicted regression results with the data, again using the same mismeasured variables as in the data. We report our results in Panel B of Table V. Our preferred comparison uses a regression specification that includes industry-by-year fixed effects, since this matches our model structure best. Given the simplicity of our model, we believe the model fit is close.

[Table V Here]

[27] We define investment deviations as $\frac{k'}{k^p}$. Since we observe capital stock and realized (planned) investment, we simply compute realized (planned) next-period capital directly using a perpetual inventory method using the appropriate parameter values.

[28] To compute model moments in the steady state, we simulate a panel of 10,000 firms for 50 years and drop the first 30 years as a burn-in.

Figure A.1 shows the full distribution of firm uncertainty in terms of the posterior uncertainty $\mathbb{V}$ as a fraction of prior uncertainty σ_ϵ^2. We find that the median (mean) firm has a posterior variance that is approximately 27% (38%) of its prior. There is also large dispersion, with an interquartile range of 41%. In contrast, if we had simply assumed a homogeneous uncertainty, we would have imposed that all firms face a posterior uncertainty that is 39% that of the prior.[29] Translating our model parameter values into dollars, we find that the aggregate cost of signal acquisition and investment plan adjustment cost is approximately US$ 5.1 billion and US$ 1.4 billion, respectively, based on GDP data from 2019 and exchange rates as of November 10, 2020.

5.2 Trading Off Better Ex Ante Information or More Ex Post Adjustments

In our model, firms can get around their imperfect information by either investing in higher-precision signals ex ante (and therefore make precise plans) or by adjusting their actual investments ex post. In both cases, the resulting ex post outcome would be identical in terms of investment decisions. However, the costs of each strategy differ across firms.

[Figure III Here]

To provide greater insight into which strategy firms prefer, Figure III shows the distribution of costs paid out in signal acquisition and investment plan adjustments (henceforth "σ-cost" and "k^p-cost", respectively). The distribution of σ-cost has a much thicker and longer tail than that of k^p-cost: reflecting the endogenous distributions, the average σ-cost paid is almost five times that of the k^p-cost. This result reflects a strong preference by firms to be correct ex ante instead of fixing their errors ex post.

Given our rich heterogeneity in firm characteristics, we also examine which firms are more likely to incur a σ-cost or k^p-cost. In Figures IIa and IIb, we plot the joint distribution of revenue and σ-cost or k^p-cost in the form of a scatterplot. Figure IIa shows that

[29] This can be computed directly from Equation 11 by taking the variance of Δz and inverting the equation to back out the variance of $\tilde{\epsilon}$. With homogeneous uncertainty, the bias term drops out from this calculation.

revenue is strongly positively correlated with σ-cost, while Figure IIb shows that revenue is negatively correlated with k^p-cost. In other words, high-revenue firms endogenously choose better information ex ante and are thus less likely to depend on ex post corrections to their investment plans. This reflects our earlier intuition that *expected* investment plan adjustment cost is increasing in productivity for a given level of uncertainty. Therefore, in equilibrium, high-productivity firms (who are high-revenue on average) want to avoid paying this cost in equilibrium, and do so by preemptively acquiring better signals. Consequently, the average cost paid out in plan adjustments becomes negatively correlated with revenue.

[Figure II Here]

Overall, based on the distribution of costs incurred, firms prefer to make better plans ex ante rather than adjust ex post. While our simple setup does not exactly identify what these costs are in a practical sense, we believe that our model provides a reference point as to what form policy could take. This result suggests that greater improvement of information provision might be more desirable for firms even if they can adjust corporate policies ex post. For instance, one policy intervention could be in the form of improving management practices a la Bloom et al. (2013).

5.3 Imperfect Information and Capital Misallocation

Firm uncertainty has often been proposed as a source of capital misallocation. In this section, we explore the role of these two levers of corporate policy (i.e., endogenous learning and investment plan flexibility) in mitigating misallocation.

We consider a counterfactual in which uncertainty is a model primitive and homogeneous across firms, and that investment plans are completely immutable. This amounts to imposing a fixed $\mathbb{V}$ for every firm and setting $\chi = \infty$. We consider this as our primary counterfactual, because this framework is conceptually most similar to common assumptions used in the literature.[30] We proxy for capital misallocation using two measures: aggregate TFP and aggregate wages.

Our quantitative exercise begins by computing the two aforementioned measures for our baseline calibration in our model and comparing them relative to a frictionless economy (i.e., $\chi = 0$ and $\xi = 0$). This comparison provides a sense of how much misallocation there is in our baseline calibration. We then reparameterize our model by imposing $\mathbb{V}$ as a fixed fraction of σ_ϵ^2 and setting $\chi = \infty$. We set $\mathbb{V}$ to replicate our targeted moment of $\frac{\sigma_{\hat{z}-z}}{\sigma_\epsilon}$. This alternative parameterization provides a counterfactual scenario to study how much misallocation is present when we simply aim to replicate the overall level of uncertainty, ignoring the heterogenous responses of firms to mitigate uncertainty.

Columns 1 and 2 of Panel A in Table VII show that capital misallocation from information imperfection is overstated by around seven to 10 times in the counterfactual, depending on the measure of misallocation. Notably, in our baseline model, misallocation due to uncertainty appears to be economically insignificant, generating TFP losses of only 0.26% compared with 2.59% in the counterfactual.

Why do the alternative assumptions lead to such a large overstatement of misallocation driven by imperfect information? This result relates to Proposition 1. On the one hand, firms with higher initial productivity face a marginally larger distortion from the same uncertainty compared with firms with lower productivity. When we impose a common level of uncertainty, on the one hand, high-productivity firms that would have chosen to pay to lower $\mathbb{V}$ are now forced to face higher uncertainty, therefore generating increased misallocation for this group of firms. On the other hand, while misallocation is lower for low-productivity firms (since they now face lower uncertainty relative to our model), the effect of improving

[30] In our model, firm managers made investment *plans* under uncertainty about *contemporaneous* and *future* productivity, while they make investment *choices* under uncertainty only about *future* productivity. This differs from the standard literature, in which managers either make investment choices (e.g., David and Venkateswaran 2019) or plans (e.g., Tanaka et al. 2020) under uncertainty over *future* productivity. Our timing assumption arises due to the necessity to match how capital budgets are actually formed in practice. Given these differences, we defer a brief discussion, connecting our mechanisms to the literature, to the end of this subsection.

information imperfection for these firms is marginally smaller. In net, misallocation would be overstated. This effect is magnified by the assumption that investment plans are immutable, since high-productivity firms cannot overcome their high amount of initial uncertainty via ex post adjustments.

We further quantify the separate contributions of allowing for learning and allowing for investment flexibility in reducing capital misallocation. Specifically, we consider two additional counterfactuals: one in which we again set a fixed $\mathbb{V}$ for all firms, but χ is set to our calibrated value (henceforth referred to as "Fixed $\mathbb{V}$"); and the other in which $\mathbb{V}$ is endogenous but $\chi = \infty$ (henceforth referred to as "Fixed Plans"). In other words, we quantify the extent to which misallocation is ameliorated if we only allowed for investment flexibility as we estimated using our data; and alternatively, if we allowed for learning but no investment flexibility.

Columns 3 and 4 of Panel A in Table VII report a sharp result: The bulk of the reduction in misallocation arises from firms trying to make correct plans ex post. For instance, the difference in the relative change in wages for our baseline calibration and this counterfactual is only 0.12 percentage points. In other words, simply being given the option to learn — even if costly — sharply reduces misallocation. In contrast, having the flexibility to adjust investment only reduces about one-third of the excess misallocation. For perspective, the difference in the relative change in wages for our baseline calibration and this counterfactual is about 2.14 percentage points.

[Table VII Here]

We now analyze our results in the context of costs incurred to achieve these reductions in misallocation. Panel B of Table VII reports the total σ-cost and k^p-cost incurred in our respective counterfactual economies relative to our baseline. For the economies with fixed $\mathbb{V}$, we assume that firms exogenously pay for information using our imposed value of $\mathbb{V}$. Three results emerge.

First, comparing Column 2 with 1, we see that the total σ-cost incurred, if we assumed homogeneous learning, is only 30% that of the baseline even though both models generate the same average uncertainty. This result speaks to the importance of accounting for heterogeneity in learning. Specifically, this implies that while some firms do choose to buy lower quality information, their cost-reduction choices are overwhelmed by the group of firms who wish to purchase better information and are willing to pay a high price for it. This is most evident in Figure A.1 of Appendix A.4, in which we plot the endogenous distribution of signal precisions and contrast that with the no-learning counterfactual.

Second, when comparing Column 3 with 1, we see that investment plans are in fact highly inflexible. We infer this from the fact that even though firms are able to reduce losses by adjusting their investment, the total cost incurred is only slightly higher than in the baseline. This implies that the marginal benefit of ex post adjustment is small relative to the cost. That said, despite few firms taking advantage of plan flexibility, the gains are still quite substantial. This is because the firms that actually use this margin of adjustment are the ones that matter; specifically, high-productivity firms. In Figure IVb, we plot the same scatterplot as in Figure IIb for this counterfactual. We see that the correlation of revenue and k^p-cost essentially falls to 0, whereas it was strongly negatively correlated in our baseline (elasticity of about -0.36). Critically, if we zoom in on the upper-right quadrant, we see that the bulk of this reduction is driven by high-revenue firms incurring larger costs of ex post adjustment. This happens because high-productivity firms have the most to lose when they are wrong; and since they cannot compensate by acquiring better information in our counterfactual, they compensate by adjusting their actual investment ex post. Because only a small fraction of firms increase their expenses, the total increase in cost is small. However, since these are the firms that matter most for the aggregate economy, aggregate factor allocation improves substantially.

[Figure IV Here]

Third, and related to our second point, learning is comparatively cheap. We infer this from the fact that the increase in σ-cost is large relative to the baseline (about a 16% increase). Unlike in the case with fixed $\mathbb{V}$, this increase is driven by an across-the-board increase in spending on information acquisition. We can see this in Figure IVa, in which the joint distribution of σ-cost and revenue is shifted upward relative to the baseline (that is, firms with the same revenue now spend a larger amount on information acquisition). Since this spending increase is across the board, the overall correlation between cost and revenue does not change much (an elasticity of 0.60 in the counterfactual, compared with 0.62 in the baseline).

5.4 Relation to the literature on capital misallocation

Since our model does not nest the standard framework, our quantitative results cannot be directly compared with prior estimates . That said, our findings are still qualitatively important. As David and Venkateswaran (2019) note, misallocation that arises due to uncertainty is relatively low because some fraction of uncertainty is usually resolved ahead of time. Our paper adds further nuance to this argument in two ways. First, we show that firms do not passively face uncertainty. Critically, the heterogenous response of firms to uncertainty can further mitigate economic losses, even holding fixed the overall aggregate level of economic uncertainty. Second, the fact that firm managers have the ability to adjust capital spending in the face of real-time shocks further dampens the effect of uncertainty on capital misallocation. In this context, a direct translation of our framework into a more conventional model would further reduce uncertainty as a source of capital misallocation.

That said, we note an important caveat. In our model, all resources spent on information acquisition (and investment plan adjustment) are productive, and does not compete for productive labor inputs. This mirrors the assumption of David and Venkateswaran (2019), in which a fraction of uncertainty is costlessly resolved ahead of time. Future work can examine whether alternative market structures for information might lead to different conclusions regarding the impact of uncertainty on capital misallocation.

6 Aggregate Risk

We now extend our baseline model to include aggregate risk. Our goal is to show that our model of endogenous learning in the face of costly investment plan adjustment naturally predicts two well-documented stylized facts: that (i) (subjective) uncertainty is countercyclical and (ii) misallocation is countercyclical.

6.1 Model

Since our model with aggregate risk largely replicates the structure of our stationary model, we will only discuss the key modifications in the interest of brevity, and focus on key changes to our model solution. We modify the physical production function to allow for aggregate shocks to productivity; namely $y_{it} = A_t z_{it} k_{it}^{\alpha} l_{it}^{1-\alpha}$, where A_t is an aggregate shock to productivity. As in the literature, we assume that A_t follows an AR(1) process in logs, given by $\log A_t = \rho_A \log A_{t-1} + \sigma_{\epsilon,A}\epsilon_A$, and that aggregate productivity is orthogonal to the process for idiosyncratic productivity. We extend our assumptions on uncertainty over idiosyncratic productivity to uncertainty over aggregate productivity. In other words, we assume A_t is also not observable to the firm manager in the day, but she can observe a noisy signal s_A of A_t prior to making her investment plans. Moreover, we assume that she can improve on her signal quality by paying a cost given by $\mathcal{C}^A(\sigma_A) = \xi_A \left(\frac{1}{\mathbb{V}_A} - \frac{1}{\sigma_{\epsilon,A}^2}\right)$, where $\mathbb{V}_A^{-1} = \left(\frac{1}{\sigma_A^2} + \frac{1}{\sigma_{\epsilon.A}^2}\right)^{-1}$ is the posterior variance of the aggregate shock given the choice of signal precision $\frac{1}{\sigma_A}$ (over aggregate conditions). Finally, we derive our results here under *partial equilibrium.* This allows us to focus the discussion on directly mapping our micro-level mechanisms, as discussed in the earlier sections, to the aggregate results here. Except for these modifications, the rest of the model remains the same.

Forms of uncertainty in our model At this point, it is useful to take stock of the four forms of uncertainty in our model with aggregate shocks. First, there is *idiosyncratic uncertainty over idiosyncratic fundamentals (i.e., over z_{it}).* Second, there is *idiosyncratic uncertainty over aggregate fundamentals (i.e., over A_t).* These are given by σ and σ_A, respectively, and are heterogeneous across firms. Third, there is *aggregate uncertainty over idiosyncratic fundamentals*. Fourth, there is *aggregate uncertainty over aggregate fundamentals*. We define the latter two forms of aggregate uncertainty as $\bar{\sigma} \equiv \int \sigma d\Lambda = \bar{\sigma}$ and $\bar{\sigma}_A \equiv \int \sigma_A d\Lambda = \bar{\sigma}_A$, respectively; that is, aggregate uncertainty is an unweighted average of all of the individual choices across the distribution Λ.

6.2 Predictions

We now discuss two key predictions. To begin, we first present our second proposition.

Proposition 2. *Under a log-linear approximation for next-period capital choice k',*

1. *the ex ante (day) expected value of the firm, gross of all signal acquisition costs, is strictly decreasing in the posterior uncertainty $\mathbb{V}_A$.*

2. *the value of increasing signal precision over idiosyncratic productivity (i.e., decreasing $\mathbb{V}$) is increasing in initial aggregate productivity A_{-1}.*

3. *the value of increasing signal precision over aggregate productivity (i.e., decreasing $\mathbb{V}_A$) is increasing in initial aggregate productivity A_{-1}.*

Appendix B.2 presents the full proof. Proposition 2 is in fact simply the "aggregate risk" counterpart to Proposition 1, in which we studied a stationary model. With Propositions 1 and 2 in hand, we now state our next two corollaries.

Corollary 2. *The optimal signal precision over idiosyncratic and aggregate productivity is increasing in initial aggregate productivity.*

From Propositions 1 and 2 and our assumptions on the cost of signal acquisition, we see that the optimal signal precision must be increasing in A_{-1}. Appendix B.2.1 shows the formal proof.

Corollary 3. *Aggregate subjective uncertainty, both over idiosyncratic and aggregate TFP, is persistent and countercyclical.*

This corollary follows naturally from the one before. Corollary 2 implies $\frac{\partial \sigma}{\partial A_{-1}} < 0$ and $\frac{\partial \sigma_A}{\partial A_{-1}} < 0$ for all firms, meaning that all firms respond to positive (negative) aggregate shocks by investing in better (poorer) signals. Consequently, aggregate uncertainty is countercyclical. Moreover, since A_{-1} is itself an autocorrelated process, $\bar{\sigma}$ and $\bar{\sigma}_A$ are also trivially autocorrelated; that is, subjective uncertainty is persistent. Appendix B.2.3 shows the formal proof. With our extended proposition and corollaries in hand, we now provide two more sets of predictions:

Prediction 3. The dispersion of firm forecast errors over TFP, at both the individual and aggregate level, is persistent and countercyclical.

Prediction 4. Misallocation, as measured by the dispersion of log ARPK, is countercyclical.

Prediction 3 arises trivially from Corollary 3. Prediction 4 also arises from Corollary 3: Since subjective uncertainty is countercyclical, misallocation will also be countercyclical. Although this is intuitive, this claim requires a formal proof, which we show in Appendix B.2.3. Our main point here is that misallocation arises because firms endogenously choose to acquire bad information in bad times, rather than because uncertainty itself is fundamentally countercyclical.

We verify our predictions, which we report in Appendix C. Our empirical results mirror the multitude of recent studies on firm uncertainty. That said, we emphasize that our findings do not simply verify prior results. Instead, we contribute to the literature by arguing for a relatively simple and intuitive economic mechanism that relates optimal capital budget planning to aggregate uncertainty. Since our model is fully informed by direct evidence, we view this as an important first step toward understanding why uncertainty is countercyclical and persistent.

7 Conclusion

In this paper, we show that accounting for a firm's capital budget formation process is important for understanding how profitability uncertainty translates to capital misallocation. To that end, we first document that investment plans (i) contain more information than simply expectations of profitability, (ii) are partially flexible, but (iii) are costly to adjust. Next, we study the implications of our findings using a parsimonious model of firm dynamics that features endogenous information acquisition and partially flexible investment plans. Importantly, our unique data can be used to directly verify the model's predictions. Our calibrated model shows that although investment plans are flexible, firms prefer to make better plans ex ante as opposed to adjusting investment on the fly in response to new information. This effect is especially salient for high-productivity firms. Quantitatively, we show that this channel significantly mutes the contribution of uncertainty to capital misallocation, due to an improved allocation of information across firms of different productivities.

References

Antle, R. and Eppen, G. D. (1985). Capital Rationing and Organizational Slack in Capital Budgeting. *Management Science*, 31(2):163–174.

Arnold, G. C. and Hatzopoulos, P. D. (2000). The theory-practice gap in capital budgeting: Evidence from the United Kingdom. *Journal of Business Finance and Accounting*, 27(5-6):603–626.

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

Bachmann, R. and Bayer, C. (2014). Investment Dispersion and the Business Cycle. *American Economic Review*, 104(4):1392–1416.

Bachmann, R., Caballero, R. J., and Engel, E. M. (2013). Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model. *American Economic Journal: Macroeconomics*, 5(4):29–67.

Bachmann, R. and Elstner, S. (2015). Firm Optimism and Pessimism. *European Economic Review*, 79:297–325.

Bachmann, R., Elstner, S., and Hristov, A. (2017). Surprise, surprise - Measuring firm-level investment innovations. *Journal of Economic Dynamics and Control*, 83:107–148.

Baik, B., Farber, D. B., and Lee, S. S. (2011). CEO Ability and Management Earnings Forecasts. *Contemporary Accounting Research*, 28(5):1645–1668.

Bar-Ilan, A. and Strange, W. C. (1996). Investment Lags. *American Economic Review*, 86(3):610–622.

Barrero, J. M. (2021). The Micro and Macro of Managerial Beliefs.

Benhabib, J., Liu, X., and Wang, P. (2016). Endogenous Information Acquisition and Countercyclical Uncertainty. *Journal of Economic Theory*, 165:601–642.

Bernardo, A. E., Cai, H., and Luo, J. (2004). Capital Budgeting in Multidivision Firms: Information, Agency, and Incentives. *Review of Financial Studies*, 17(3):739–767.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does Management Matter? Evidence From India. *Quarterly Journal of Economics*, 118(1):1351–1408.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-eksten, I., and Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, 86(3):1031–1065.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Over-reaction in Macroeconomic Expectations. *American Economic Review*.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic Expectations and Credit Cycles. *Journal of Finance*, 73(1):199–227.

Christiano, L. J. and Todd, R. M. (1996). Time to plan and aggregate fluctuations. *Federal Reserve Bank of Minneapolis Quarterly Review*.

Coibion, O. and Gorodnichenko, Y. (2015). Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review*, 105(8):2644–2678.

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

David, J. M. and Venkateswaran, V. (2019). The Sources of Capital Misallocation. *American Economic Review*, 109(7):2531–67.

Gennaioli, N., Ma, Y., and Shleifer, A. (2015). Expectations and Investment. In *NBER Macroeconomics Annual*, number 30, pages 379–442.

Goodman, T. H., Neamtiu, M., Shroff, N., and White, H. D. (2014). Management Forecast Quality and Capital Investment Decisions. *The Accounting Review*, 89(1):331–365.

Harris, M., Kriebel, C. H., and Raviv, A. (1982). Asymmetric Information, Incentives and Intrafirm Resource Allocation. *Management Science*, 28(6):604–620.

Harris, M. and Raviv, A. (1996). The capital budgeting process: Incentives and information. *The Journal of Finance*, 51(4):1139–1174.

Hopenhayn, H. A. (1992). Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica*, 60(5):1127–1150.

Ilut, C. and Saijo, H. (2020). Learning, Confidence, and Business Cycles. *Journal of Monetary Economics*.

Ilut, C. and Valchev, R. (2020). Economic agents as imperfect problem solvers. *Working Paper*.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

Kuehn, L.-A. (2011). Disentangling Investment Returns and Stock Returns: The Importance of Time-to-Build. *SSRN Electronic Journal*.

Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6).

Lamont, O. A. (2000). Investment Plans and Stock Returns. *The Journal of Finance*, 55(6):2719–2745.

Lee, S., Matsunaga, S. R., and Park, C. W. (2012). Management forecast accuracy and CEO turnover. *Accounting Review*, 87(6):2095–2122.

Li, J. and Wang, H. (2020). The Expected Investment Growth Premium. *Working Paper*.

Li, J., Wang, H., and Yu, J. (2020). Aggregate expected investment growth and stock market returns. *Journal of Monetary Economics*, (Forthcoming).

Ma, Y., Sraer, D., Ropele, T., and Thesmar, D. (2020). A Quantitative Analysis of Distortions in Managerial Forecasts. *NBER Working Paper Series*, pages 1–61.

Malenko, A. (2019). Optimal dynamic capital budgeting. *Review of Economic Studies*, 86(4):1747–1778.

Mao, J. C. T. (1970). Survey of Capital Budgeting: Theory and Practice. *Journal of Finance: Papers and Proceedings*, 25(2):349–360.

Myers, S. C. (1974). Interactions of Corporate Financing and Investment Decisions - Implications for Capital Budgeting. *The Journal of Finance*, 29(1):1–25.

Reis, R. (2006). Inattentive producers. *Review of Economic Studies*, 73(3):793–821.

Ryan, P. a. and Ryan, G. P. (2002). Capital Budgeting Practices of the Fortune 1000: How Have Things Changed? *Journal of Business & Management*, 8(4):355.

Schall, L. D., Sundem, G. L., and Geijsbeek, W. R. (1978). Survey and Analysis of Capital Budgeting Methods. *Journal of Finance*, 33(1):281–287.

Schmitt-Grohé, S. and Uribe, M. (2012). What's News in Business Cycles. *Econometrica*, 80(6):2733–2764.

Senga, T. (2018). A New Look at Uncertainty Shocks: Imperfect Information and Misallocation. *Working paper*.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Tanaka, M., Bloom, N., David, J. M., and Koga, M. (2020). Firm Performance and Macro Forecast Accuracy. *Journal of Monetary Economics*, 18(9):1–38.

Van Nieuwerburgh, S. and Veldkamp, L. (2009). Information immobility and the home bias puzzle. *Journal of Finance*, 64(3):1187–1215.

Woodford, M. (2003). Imperfect Common Knowledge and the Effects of Monetary Policy. In Aghion, P., Frydman, R., Stiglitz, J., and Woodford, M., editors, *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pages 25–58. Princeton, NJ: Princeton University Press.

8 Tables and Figures

Table I: Summary Statistics

The table below shows the summary statistics of our firm-year panel. The total number of firms in our sample is 5,989, of which 2,273 are publicly listed and the rest are private companies. To reduce the influence of outliers on these summary statistics, we winsorize variables at the 1% level. Capital stock is the Net Plants, Property & Equipment. Profits are reported ordinary profits according to Japanese Generally-Accepted Accounting Principles (GAAP). The Wage Bill is the sum of total salary cost for employees and company officers as well as the bonus for employees and company officers. Investment plans are represented as a percentage of end-of-previous-fiscal-year capital stock. Employment is the number of employees is represented as the number of full-time equivalent workers and may include fractions. When calculating residualized AR(1) and TFP shocks, we use the MoF industry-level Cobb-Douglas estimated labor cost shares with additional details Appendix Section A.3. All numbers are rounded to three significant digits or three decimal points, whichever results in fewer decimal points.

Panel A: Firm Fundamentals						
					Percentile	
Variable	Mean	SD	Skew	25^{th}	50^{th}	75^{th}
Total Assets (mn)	88,100	205,000	4.81	11,060	26,700	67,300
Capital Stock (mn)	50,000	131,000	5.14	4,590	12,300	33,100
Employment (count, FTE)	1,180	2,100	3.63	162	466	1,180
COGS (mn)	57,100	125,000	4.62	4,950	17,400	49,800
Wage Bill (mn)	1,670	3,120	4.11	229	649	1,650
Sales (mn)	72,700	152,000	4.47	7,810	23,800	66,300
Expected Sales (mn)	75,300	157,000	4.43	8,100	24,700	68,800
Profits (mn)	3,760	9,050	4.56	207	920	3,070
Expected Profits (mn)	3,700	8,650	4.57	247	930	3,000
Total Capital Investment (mn)	3,060	8,500	5.23	121	545	1,990
Total Capital Investment Plan (mn)	3,250	9,160	5.20	100	554	2,030
Panel B: Constructed Variables						
Investment (% of Capital)	7.63	9.55	2.67	1.63	4.58	9.76
Capital Investment Plan (% of Capital)	7.73	9.78	2.67	1.40	4.69	10.1
Investment Plan Deviations (% of Capital)	-0.140	5.55	0.63	-1.64	-0.0610	0.947
log TFP Shock	-0.127	0.480	-1.06	-0.275	-0.043	0.082

Table II: Investment Plans and Expected Firm Performance.

The table below shows the relation between realized investment rates and expected performance measures as well as investment plans. The subscript t denotes a fiscal year and i denotes a firm. VA stands for value added and is defined as total sales minus costs of goods sold which is assumed to be known ex ante (so expected VA is expected sales minus realized costs of goods sold). $\mathbb{E}(x)$ is the expectation based on the beginning of the fiscal year. Where there is no confusion, we drop unnecessary subscripts. Additional controls include the cash to total assets, book leverage ratio, and log total assets from the previous fiscal year. Expectations, actual values, and shocks are winsored at the 1% level. All regressions include industry-year fixed effects. Standard errors are clustered by firm and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Dependent Variable:	i/k					
Performance Measure:	$\log(\mathbb{E}\left[\text{Sales}\right]/k)$		$\log(\mathbb{E}\left[VA\right]/k)$		$\log\left(\mathbb{E}\left[z\right]\right)$	
	(1)	(2)	(3)	(4)	(5)	(6)
Expected Performance	2.610***	0.874***	2.231***	0.676***	1.719***	0.493***
	(0.223)	(0.140)	(0.199)	(0.135)	(0.198)	(0.127)
i^p		0.634***		0.638***		0.644***
		(0.051)		(0.051)		(0.050)
Observations	26,718	26,718	26,718	26,718	26,718	26,718
R^2	0.107	0.476	0.098	0.474	0.080	0.473

Table III: Investment Plans, Investment Errors, and TFP Shocks.

The table below shows the relation between firm-level annual investment errors and sales, value-added, and TFP shocks. The subscript time t denotes a fiscal year. The variable $\Delta i = \frac{i-\mathbb{E}[i]}{k}$ is the annual investment plan deviation relative to the initial full-year investment plan made in the first quarter survey, scaled by previous fiscal year's total net plants, property, and equipment. $\Delta \log Sales$ is defined as $\log(Sales/K) - \log(\mathbb{E}[Sales]/K)$, $\Delta \log VA$ is defined as $\log(VA/K) - \log(\mathbb{E}[VA]/K)$ where VA stands for value added and is defined as total sales minus costs of goods sold which is assumed to be known ex ante (so expected VA is expected sales minus realized costs of goods sold), and $\Delta \log z$ is defined in Section 2.2.1. Where there is no confusion, we drop unnecessary subscripts. Additional controls include the cash to total assets, book leverage ratio, and log total assets from the previous fiscal year. Expectations, actual values, and shocks are winsored at the 1% level. All regressions include industry-year fixed effects. Standard errors are clustered by firm and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Performance Measure:	$\Delta \log Sales$		$\Delta \log VA$		$\Delta \log z$	
Dependent Variable:	$\Delta \frac{i}{k}$	$\frac{i}{k}$	$\Delta \frac{i}{k}$	$\frac{i}{k}$	$\Delta \frac{i}{k}$	$\frac{i}{k}$
	(1)	(2)	(3)	(4)	(5)	(6)
Performance Shock:	2.337***	2.301***	5.023***	3.381***	0.782***	0.532***
	(0.633)	(0.568)	(1.267)	(1.144)	(0.168)	(0.155)
i^p		0.648***		0.649***		0.649***
		(0.050)		(0.050)		(0.050)
Observations	26,718	26,718	26,718	26,718	26,718	26,718
R^2	0.033	0.471	0.034	0.470	0.034	0.471

Table IV: Investment Plan Deviations and Future Profitability

The table below shows the relation between firm-level annual investment plan deviations and future gross profit margins. The subscript time t denotes a fiscal year. Future Gross Profit Margin is the ordinary income divided by sales in the next year. $\Delta \log(z_{-1})$ is the TFP shock from the previous year. Where there is no confusion, we drop unnecessary subscripts. Additional controls include the cash to total assets, book leverage ratio, and log total assets from the previous fiscal year. Expectations, actual values, and shocks are winsored at the 1% level. All regressions include industry-year fixed effects. Standard errors are clustered by firm and shown in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Dependent Variable:	Future Gross Profit Margin (% of Sales)		
	(1)	(2)	(3)
$\log\left(1 + \lvert\Delta \frac{i}{k}\rvert\right)$	-0.334***	-0.355***	-0.351***
	(0.118)	(0.121)	(0.122)
i	-0.028***	-0.033***	-0.030***
	(0.010)	(0.010)	(0.010)
$\log(z)$	2.649***	2.602***	2.189***
	(0.159)	(0.166)	(0.206)
$\Delta \log(z_{-1})$		1.094***	1.116***
		(0.133)	(0.134)
Observations	20,082	20,082	20,082
Fixed Effects	Industry×Year	Industry×Year	Industry×Year
Additional Controls			Yes
R^2	0.322	0.328	0.333

Table V: Model Fit

In panel A, we reported model fit per our targeted moments. Note that we also clear the market for labor and output down to a tolerance of 10^{-8}. In panel B, we report model fit per untargeted moments, namely, our two predictions from the earlier section.

Panel A: Targeted Moments			Panel B: Untargeted Moments		
Moments	Data	Model	Moments	Data	Model
$\frac{\sigma_{\hat{z}-z}}{\sigma_\epsilon}$	1.347	1.347	Prediction 1	-0.051	-0.12
$\sigma\left(\frac{k'}{k^p}\right)$	0.073	0.073	Prediction 2	0.17	0.27
$(N, P) = (1, 1)$	—	$(1, 1)$			

Figure I: Marginal Benefit and Marginal Cost of Improving Signal Precision

Graphs are plotted against $\mathbb{V}^{-1}$ (i.e., increasing signal precision). Marginal cost is a constant ξ by assumption. Note that figure is purely illustrative as $\mathbb{V}^{-1}$ is bounded below by σ_ϵ^{-2}.

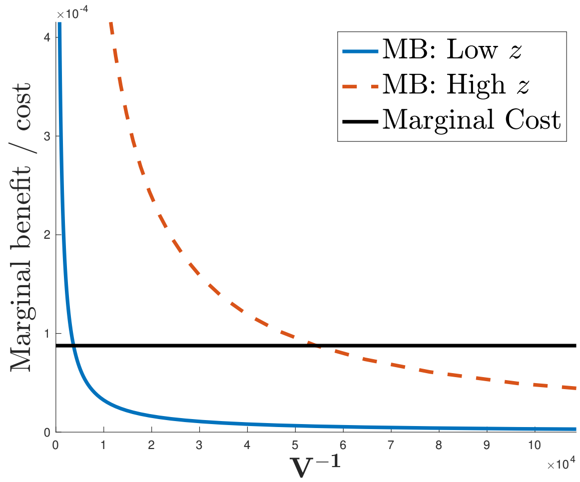

Figure II: Revenue and Distributions of σ-cost & k^p-cost

The joint probability distribution of cost and revenue. The grey solid line represents a best fit line from ordinary least squares.

(a) σ-cost and Revenue

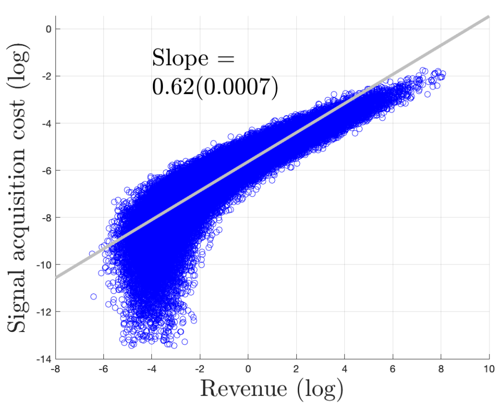

(b) k^p-cost and Revenue

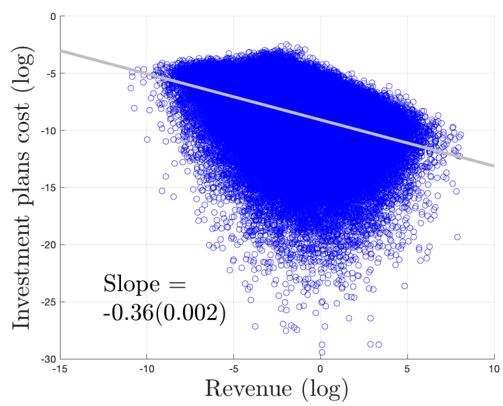

Figure III: Cost Distributions

Distribution of signal acquisition (blue bars connected with circles) and investment adjustment costs (red bars connected with lines) incurred by firms.

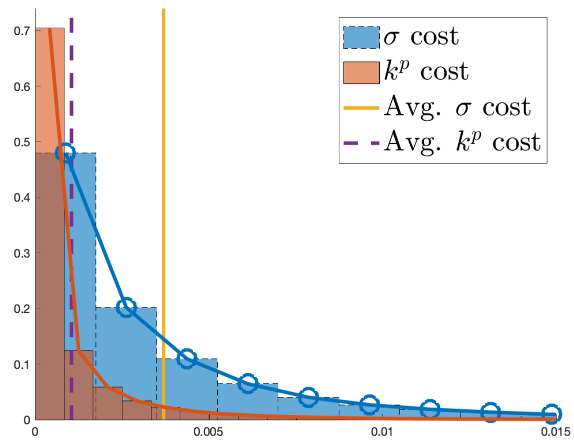

Table VI: Model Predictions

The table below tests predictions 1 and 2 from Section 4.2. Where there is no confusion, we drop unnecessary subscripts. The measure of dispersion of a variable is the log of one plus the absolute value of the variable. The estimated constants in the regressions with no fixed effects are suppressed for space. Expectations, actual values, and shocks are winsored at the 1% level. Standard errors are clustered by firm and shown in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Panel A: Prediction 1 – TFP Shock Dispersion and Productivity		
Dependent Variable:	Dispersion of $\Delta \log z$	
	(1)	(2)
$\log(z_{-1})$	-0.011***	-0.051***
	(0.003)	(0.003)
Observations	26,718	26,718
Fixed Effects		Industry×Year
R^2	0.002	0.234
Panel B: Prediction 2 – Investment Plan Deviation and Productivity		
Dependent Variable:	Dispersion of Δi	
	(1)	(2)
Dispersion of $\Delta \log z$	0.194***	0.170***
	(0.029)	(0.030)
Observations	26,718	26,718
Fixed Effects		Industry×Year
R^2	0.003	0.079

Table VII: Counterfactual Moments

Panel A: Aggregate TFP, and wages, relative to a frictionless model. The measures are reported as *decreases* relative to frictionless reference. Panel B: Change in costs incurred relative to baseline model. For columns 2 and 3, the change in σ cost is the same by construction.

	Baseline	Fixed $\mathbb{V}$, $\chi = \infty$	Fixed $\mathbb{V}$, $\chi < 0$	Endogenous $\mathbb{V}$, $\chi = \infty$
	(1)	(2)	(3)	(4)
		Panel A: Effect on misallocation		
ΔTFP	0.26%	2.59%	1.75%	0.34%
Δw	0.39%	3.74%	2.53%	0.51%
		Panel B: Costs (relative to baseline)		
k^p-cost	1	0.000	1.025	0.000
σ-cost	1	0.297	0.297	1.155

Figure IV: Revenue and Distributions of σ-cost & k^p-cost

Scatterplot of (left) σ-cost and (right) k^p-cost against revenue. Blue circles and solid grey reference line are from the baseline model as in Figure IIa; Red crosses and dashed grey reference line are from the relevant counterfactual comparison models. Reference lines correspond to predicted values of a univariate regression of log cost on log revenue. For σ-cost, the counterfactual is assuming fixed plans. For k^p-cost, the counterfactual is assuming fixed $\mathbb{V}$.

(a) σ-cost and Revenue

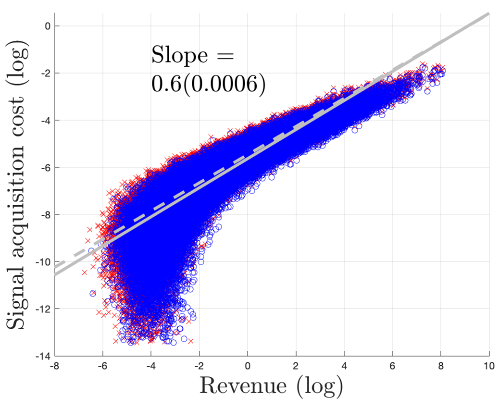

(b) k^p-cost and Revenue

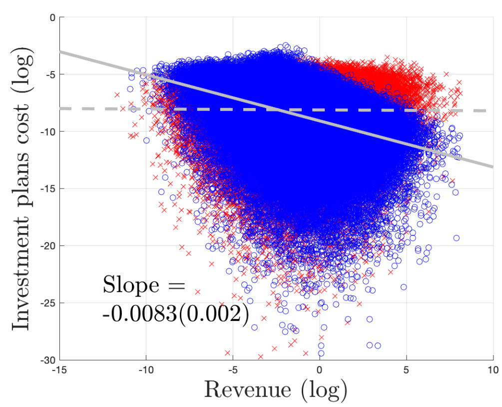

Online Appendix

A Empirical Methodology and Data Details

A.1 Detailed Summary Statistics

Table A.1: Additional Summary Statistics

The table below shows the summary statistics of our firm-year panel. The total number of firms in our sample is 5,989, of which 2,273 are publicly listed and the rest are private companies. To reduce the influence of outliers on these summary statistics, we winsorize variables at the 1% level. Capital stock is the Net Plants, Property & Equipment. Profits are reported ordinary profits according to Japanese Generally-Accepted Accounting Principles (GAAP). The Wage Bill is the sum of total salary cost for employees and company officers as well as the bonus for employees and company officers. Investment plans are represented as a percentage of end-of-previous-fiscal-year capital stock. Employment is the number of employees represented as the number of full-time equivalent workers and may include fractions. When calculating residualized AR(1) and TFP shocks, we use the MoF industry-level Cobb-Douglas estimated labor cost shares with additional details Appendix Section A.3. All numbers are rounded to three significant digits or three decimal points, whichever results in fewer decimal points.

Panel A: Firm Fundamentals						
				Percentile		
Variable	Mean	SD	Skew	25^{th}	50^{th}	75^{th}
Land Investment (mn)	135	583	5.88	0.000	0.000	0.000
Land Investment Plan (mn)	59.7	308	6.34	0.000	0.000	0.000
Software Spending (mn)	190	562	4.89	0.000	15	100
Software Spending Plan (mn)	191	580	4.99	0.000	9	100
Panel B: Constructed Variables						
Investment Plan Deviations (% of Capital)	-0.140	5.55	0.63	-1.64	-0.0610	0.947
Land Purchasing Plan (% of Assets)	0.127	0.631	6.05	0.000	0.000	0.000
Land Purchasing Plan Deviations (% of Assets)	0.187	0.964	5.38	0.000	0.000	0.000
Software Spending Plan (% of Assets)	0.867	2.650	5.64	0.000	0.0690	0.557
Software Spending Plan Deviations (% of Assets)	0.003	1.160	0.725	-0.080	0.000	0.083

Table A.2: Additional Summary Statistics on Gross Investment Rates From Financial Statements

The table below shows the summary statistics of our firm-year panel. The total number of firms in our sample is 5,989, of which 2,273 are publicly listed and the rest are private companies. To calculate gross investment rates from financial statements (which are reported annually), we assume a capital depreciation rate of 6%. All investment rates in this table come from financial statements and not the BOS. All numbers are rounded to three decimal points when in percentages and four decimal point sin probabilities.

Panel A: Across Years

	(%)	(%)		Capital Sales		Inaction Region		From BOS	
Fiscal Year	$\mathbb{E}[\frac{i}{k}]$	$\mathbb{E}[\frac{i}{k} \mid \frac{i}{k} < 0]$	$P(\frac{i}{k} < 0)$	$P(\frac{i}{k} < -1\%)$	$P(\frac{i}{k} < -5\%)$	$P\left(\frac{\lvert i\rvert}{k} < 1\%\right)$	$P\left(\frac{\lvert i\rvert}{k} < 5\%\right)$	$P(\frac{i}{k} < 1\%)$	$P(\frac{i}{k} < 5\%)$
2005	4.655	-4.261	0.1809	0.1056	0.0396	0.1684	0.4807	0.2070	0.5246
2006	4.491	-3.428	0.1962	0.1216	0.0322	0.1656	0.4956	0.1932	0.5041
2007	5.697	-2.771	0.1440	0.0757	0.0206	0.1477	0.4493	0.1917	0.5034
2008	5.901	-5.400	0.1297	0.0716	0.0244	0.1311	0.4300	0.1776	0.5036
2009	3.970	-3.493	0.1603	0.0826	0.0303	0.1905	0.5274	0.2500	0.6075
2010	2.947	-2.794	0.2351	0.1305	0.0285	0.2252	0.5996	0.2238	0.5868
2011	3.631	-2.277	0.1763	0.0844	0.0144	0.2135	0.5654	0.2203	0.5750
2012	3.885	-3.612	0.1866	0.1006	0.0273	0.1930	0.5363	0.1974	0.5366
2013	5.359	-4.079	0.1355	0.0771	0.0237	0.1524	0.4618	0.1820	0.5177
2014	5.196	-3.693	0.1288	0.0690	0.0226	0.1539	0.4659	0.1931	0.5176
2015	5.414	-4.333	0.1230	0.0736	0.0256	0.1326	0.4501	0.1817	0.5033
2016	4.406	-3.508	0.1617	0.0953	0.0285	0.1693	0.5022	0.1827	0.5087
Total	4.591	-3.579	0.1631	0.0908	0.0267	0.1695	0.4957	0.2000	0.5314

Panel B: Public versus Private Companies

	(%)	(%)		Capital Sales		Inaction Region		From BOS	
	$\mathbb{E}[\frac{i}{k}]$	$\mathbb{E}[\frac{i}{k} \mid \frac{i}{k} < 0]$	$P(\frac{i}{k} < 0)$	$P(\frac{i}{k} < -1\%)$	$P(\frac{i}{k} < -5\%)$	$P\left(\frac{\lvert i\rvert}{k} < 1\%\right)$	$P\left(\frac{\lvert i\rvert}{k} < 5\%\right)$	$P(\frac{i}{k} < 1\%)$	$P(\frac{i}{k} < 5\%)$
Publicly-Listed	4.916	-3.265	0.1326	0.0785	0.0216	0.1321	0.4836	0.1525	0.5410
Private Companies	4.232	-3.767	0.1891	0.1013	0.0310	0.2015	0.5061	0.2401	0.5232

A.2 How Good are Forecasts and Plans?

We show that the forecasts and plans that a firm reports do in fact predict its realized counterparts; in other words, these forecasts and plans are relevant. To do so, we use regressions of the form:

$$y_{i,t} = \alpha_{j(i),t} + \alpha_i + \beta \mathbb{E}\left[y_{i,t}\right] + \varepsilon_{i,t} \tag{9}$$

where the i subscript indexes a firm, $j(i)$ indexes the MoF industry of the firm, and t indexes a fiscal year. The variables $\mathbb{E}\left[y_{i,t}\right]$ is the forecast or plan of the outcome variable $y_{i,t}$ made from the first quarter of the same fiscal year, α_i denotes firm fixed effects, and $\alpha_{j(i),t}$ denotes industry-by-year fixed effects. Standard errors are clustered by firm. The outcome variables considered are realized capital investment, land investments, software expenses, profits, and sales. As discussed, for investment, land investment, and software expenses, the expectation variable is interpreted as a spending plan; while for profits and sales, the expectation variable is interpreted as a forecast.

The empirical specification in Equation 9 controls for industry and macroeconomic shocks and compares firms with higher expected spending plans with those in the same industry and year with lower investment plans. The coefficient of interest is β – capturing the importance of plans on actual realizations across different fiscal years. $\beta = 1$ corresponds to perfectly sticky plans which do not permit ex post adjustments while $\beta = 0$ corresponds to completely uninformative plans for actual spending. Finally, in addition to reporting the point estimate, we also consider the statistical importance represented by the relative improvement of the R^2 in the regressions compared against a regression specification without the expectation variable.

Our empirical results suggest the BOS corporate plans and forecasts are good predictors of actual realizations, and hence economically relevant. Columns 1 through 3 of Panel A in Table A.3 study spending plans for physical capital investment, land investment, and software spending, while Columns 4 and 5 evaluate profits and sales forecasts. For realized physical capital investment, plans have an estimated coefficient of 0.666 and increases the R^2 by 28% relative to a model with only firm and industry-by-quarter fixed effects. For land investment, plans have an estimated coefficient of 0.904 and increases the R^2 by 32%. For software expenses, plans have an estimated coefficient of 0.494 and increases the R^2 by only 7%. These results for both the point estimates and relative statistical fit improvements are consistent with an intuitive ranking of how costly it is to adjust your investment plans: changing a firm's land purchasing plan incurs the highest cost, followed by physical capital investment, and then by software spending. Notably, spending for which deviation from plans is more costly will carry a coefficient closer to one and account for a larger statistical variation of actual realized investment.

Table A.3: Spending Plans and Forecasts

The table below shows the relation between expected firm measures as the explained variable and actual realized firm measures as the explanatory variable. Investment and the change in software spending is scaled by previous end of period capital while changes in land investment, profits, and sales are scaled by previous end of period total assets. The subscript t denotes a fiscal year and i denotes a firm. VA stands for value added and is defined as total sales minus costs of goods sold which is assumed to be known ex ante (so expected VA is expected sales minus realized costs of goods sold). $\mathbb{E}(x)$ is the expectation based on the beginning of the fiscal year. Expectations, actual values, and shocks are winsored at the 1% level. Where there is no confusion, we drop unnecessary subscripts. All regressions include firm and industry-by-year fixed effects. Standard errors are clustered by firm and shown in parentheses. * p $<$ 0.10, ** p $<$ 0.05, *** p $<$ 0.01.

Dependent Variable:	y_t				
y_t =	PP&E	Land Purchases	Software Spending	Profits	Sales
	(1)	(2)	(3)	(4)	(5)
$\mathbb{E}[y_t]$	0.666***	0.904***	0.494***	0.759***	0.915***
	(0.015)	(0.0401)	(0.028)	(0.030)	(0.012)
Observations	26,718	26,588	26,412	26,718	26,718
R^2	0.739	0.442	0.839	0.885	0.995
R^2 without Expected Value	0.578	0.335	0.787	0.802	0.940
Relative % Increase	28%	32%	7%	10%	36%

A.3 Estimating α, and TFP

In this section, we detail our estimation strategy. For the most part, we follow Asker et al. (2014) in our measurement strategy. The basic assumption in our model is that we have a physical production function given by $y = zk^{\alpha}l^{1-\alpha}$. Under assumptions of constant elasticity of substitution (CES) across goods and monopolistic competition, we get the "sales" production function $py = z^{\theta} k^{\alpha\theta} l^{(1-\alpha)\theta}$, where $\theta \equiv \frac{\eta-1}{\eta}$. However, this is a "value-added" production function. Therefore, to map the data to the model, we need to compute value added in the data. We consider three specifications of economic value added: (1) Total sales minus the costs of goods sold (which includes materials), (2) total sales scaled by one minus the fraction of material costs to total sales based on aggregate statistics, following David and Venkateswaran (2019), and (3) total sales. Our main specification in the paper uses method (1), but our reported empirical results are robust to using (2) or (3).

A.3.1 Estimating α

As in Asker et al. (2014), we assume that capital is quasi-fixed (like in our model) but labor is free to adjust every period after productivity z has been observed. Let profits net of labor cost be $py - wl$. Then the optimality condition for labor is

$$\begin{aligned} MRPL \equiv (1-\alpha)\,\theta z^{\theta} k^{\alpha\theta} l^{(1-\alpha)\theta} &= w \\ \implies \frac{wl}{py} &= (1-\alpha)\,\theta. \end{aligned}$$

In other words, we can identify $(1-\alpha)\,\theta$ by simply computing the labor share in value added. We follow Asker et al. (2014) by computing the industry median from the individual firm-year level labor cost shares:

$$\widehat{(1-\alpha)}\,\theta = median\left\{\frac{wl}{py}\right\}.$$

Finally, we estimate α by assuming that $\theta = 0.75$ (i.e., the elasticity of demand η is 4), and then directly back out α.[31]

A.3.2 Estimating TFP, Expected TFP and TFP Shocks

With θ and α in hand, we compute realized TFP (in levels) as

$$z = \left(\frac{py}{k^{\alpha\theta} l^{(1-\alpha)\theta}}\right)^{\frac{1}{\theta}},$$

[31]Our choice of $\eta = 4$ follows from the literature, e.g., Asker et al. (2014); Bloom et al. (2018).

where py is value added for the fiscal year, used in computing α. To compute expected TFP z^e (in levels), we assume that

$$z^e = \left(\frac{py^e}{k^{\alpha\theta} l^{(1-\alpha)\theta}} \right)^{\frac{1}{\theta}},$$

where py^e is forecasted value added for the fiscal year. As discussed in the main text, forecasted value added is computed by simply subtracting forecasted sales from realized costs of goods sold. Finally, we define TFP "shocks" as

$$\Delta \log z \equiv \log z - \log z^e$$

It is clear that our measured z^e is only an approximation of the "true" expected TFP. To be precise, our measurement strategy gives a value of expected TFP that can be expressed as

$$z^e = \left(\frac{\mathbb{E}\left[z^{\theta + \frac{\theta^2(1-\alpha)}{1-(1-\alpha)\theta}} \right]}{z^{\frac{\theta^2(1-\alpha)}{1-(1-\alpha)\theta}}} \right)^{\frac{1}{\theta}},$$

which is not $\mathbb{E}[z]$. In the next section, we proceed to explain why this mismeasurement is not an area of concern for us.

A.3.3 Mismeasured Expected TFP: Why It Happens and Why It Does Not Matter For Our Results

We structure our discussion here in three stages. First, we show that under certain circumstances, measured expected TFP will always exhibit a negative bias, but the bias is small and irrelevant. Second, we show why, due to our data limitations, the bias in z^e can be large, and cannot be signed in general. In particular, we formally derive an expression to equantify this (see equation 10 in this appendix). Finally, we show that the bias in $\Delta \log z$ (i.e., TFP shocks) can always be signed. Specifically, the bias is only in one direction: our mismeasured TFP shocks is always smaller (or more negative) than the true TFP shock. Importantly, because our focus is on studying $\Delta \log z$, not z^e, this means that our qualitative results in Section 2, and quantitative results in Section 5, are not affected by the mismeasurement in z^e.

Best case scenario: Bias in z^e is small and can be signed We begin by emphasizing here that it is impossible to directly estimate an unbiased measure of expected TFP using just balance sheet data alone, even if we observe all possible expectations of these balance sheet variables. For instance, in our setup, suppose we observed expected labor or $\alpha = 1$. In both cases, the expression above is reduced to

$$z^{e^*} = \left(\mathbb{E}\left[z^\theta\right]\right)^{\frac{1}{\theta}},$$

where we denote z^{e^*} as the expected TFP one would back out under the assumptions above. Due to Jensen's inequality, $z^{e^*} < \mathbb{E}[z]$ since $\theta < 1$. That is to say, the mismeasured expected TFP is always smaller than the true expected TFP. For example, suppose we assume that the firm's expectations follow our model, then

$$\begin{aligned} z^{e^*} &= \left[z_{-1}^{\rho\theta} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2}\theta s\right) \exp\left(\frac{1}{2}\theta^2 \mathbb{V}\right)\right]^{\frac{1}{\theta}} \\ &= z_{-1}^{\rho} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2} s\right) \exp\left(\frac{1}{2}\theta \mathbb{V}\right), \end{aligned}$$

where s is the private signal as observed by the firms in our model, and σ (and corresponding $\mathbb{V}$) are the endogenous choice of uncertainty. The unbiased expected value of z is

$$\mathbb{E}[z] = z_{-1}^{\rho} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2} s\right) \exp\left(\frac{1}{2}\mathbb{V}\right).$$

This gives us a biased estimated of expected TFP, specifically,

$$z^{e^*} - \mathbb{E}\left[z\right] = z_{-1}^{\rho} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2} s\right) \left[\exp\left(\frac{1}{2}\theta \mathbb{V}\right) - \exp\left(\frac{1}{2}\mathbb{V}\right)\right] < 0,$$

where the bias shows up because of the terms in the square brackets. However, because the volatility of TFP (i.e, σ_ϵ) is typically small, and $\mathbb{V} < \sigma_\epsilon^2$, the bias will be relatively small. In other words, $z^{e^*} \approx \mathbb{E}\left[z\right]$.

The generic case: Bias in z^e is large and cannot be signed In our case, we do not observe expected labor, and α is clearly less than unity. However, if we follow our model's assumptions (as in Section 3), we can make further headway into understanding the source of the bias, by expressing expected TFP as

$$z^e = \left(\frac{z_{-1}^{\rho\theta + \rho\hat{\theta}} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2}\left(\theta + \hat{\theta}\right) s\right) \exp\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)^2 \mathbb{V}\right)}{z^{\hat{\theta}}}\right)^{\frac{1}{\theta}},$$

where $\hat{\theta} \equiv \frac{\theta^2(1-\alpha)}{1-(1-\alpha)\theta}$. The above expression can be further reduced to

$$\begin{aligned}
z^e &= z_{-1}^{\rho} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2} s\right) \left(\frac{\exp\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)^2 \mathbb{V}\right)}{\exp\left(\hat{\theta}\tilde{\epsilon}\right)}\right)^{\frac{1}{\theta}} \\
&= z^{e^*} \left(\frac{\exp\left(\left(\theta\hat{\theta} + \frac{1}{2}\hat{\theta}^2\right)\mathbb{V}\right)}{\exp\left(\hat{\theta}\tilde{\epsilon}\right)}\right) \\
\Leftrightarrow \log z^e &= \log z^{e^*} + \left(\theta\hat{\theta} + \frac{1}{2}\hat{\theta}^2\right)\mathbb{V} - \hat{\theta}\tilde{\epsilon} \qquad (10)
\end{aligned}$$

where we utilize the fact that $z = z_{-1}^{\rho} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2} s\right) \exp\left(\tilde{\epsilon}\right)$ with $\tilde{\epsilon} \sim N\left(0, \mathbb{V}\right)$ being the mean 0 innovations arising under the posterior distribution, per our model assumptions; and we further substituted in the definition for z^{e^*}. Notice that even after taking z^{e^*} as our reference, there is no clear direction for the bias in expected TFP, which depends on the exact innovation the firm receives. Notably, the implication here is that $\tilde{\epsilon}$ introduces attenuation bias into our regression framework (specifically, in Section A.2). That said, as we reported in Table II, we find that investment is positively and significantly correlated with $\log z^e$ despite the attenuation bias, suggesting that the bias might not be that severe.

Large unsigned bias in z^e does not matter for our results We now turn to relating mismeasurement in z^e to mismeasurement in the forecast errors. Specifically, we can derive a bias for forecast errors as

$$
\begin{aligned}
\frac{z^e}{z} &= \left(\frac{z_{-1}^{\rho\theta+\rho\hat{\theta}} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} \left(\theta+\hat{\theta}\right) s \right) \exp\left(\frac{1}{2} \left(\theta+\hat{\theta}\right)^2 \mathbb{V} \right)}{z^{\hat{\theta}+\theta}} \right)^{\frac{1}{\theta}} \\
&= \left(\frac{z_{-1}^{\rho\theta+\rho\hat{\theta}} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} \left(\theta+\hat{\theta}\right) s \right) \exp\left(\frac{1}{2} \left(\theta+\hat{\theta}\right)^2 \mathbb{V} \right)}{z_{-1}^{\rho\left(\hat{\theta}+\theta\right)} \exp\left(\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} \left(\theta+\hat{\theta}\right) s \right) \exp\left(\left(\theta+\hat{\theta}\right) \tilde{\epsilon} \right)} \right)^{\frac{1}{\theta}} \\
&= \left(\frac{\exp\left(\frac{1}{2} \left(\theta+\hat{\theta}\right)^2 \mathbb{V} \right)}{\exp\left(\left(\theta+\hat{\theta}\right) \tilde{\epsilon} \right)} \right)^{\frac{1}{\theta}} \\
\Leftrightarrow \Delta \log z \equiv \log z - \log z^e &= \frac{1}{\theta} \left(\left(\theta+\hat{\theta}\right) \tilde{\epsilon} - \frac{1}{2} \left(\theta+\hat{\theta}\right)^2 \mathbb{V} \right) \\
&= \frac{\theta+\hat{\theta}}{\theta} \left(\tilde{\epsilon} - \frac{1}{2} \left(\theta+\hat{\theta}\right) \mathbb{V} \right). \qquad (11)
\end{aligned}
$$

Note that the corresponding unbiased measure of TFP shocks is simply $\tilde{\epsilon}$.

It is clear now why neither our main empirical results in Section 2 nor that in our calibration in Section 5 are affected by the bias in expected values — the bias in $\Delta \log z$ is constant (it is shifted by $\frac{1}{2}\left(\theta+\hat{\theta}\right)\mathbb{V}$). For our empirical results, we are interested in mapping deviations from expected value, that is to say how higher than expected TFP (or lower than expected TFP) affects investment. Since the bias is constant, it is absorbed empirically by the fixed effects.

Relationship to Section 2 As an example, consider two firms with $\mathbb{V}$, one that receives an innovation of $\tilde{\epsilon} = 0$ and another with an innovation of $\tilde{\epsilon} = \frac{1}{2}\left(\theta + \hat{\theta}\right)^2 \mathbb{V}$. Our shocks measure is biased: For the first firm, we record it as having a negative shock, whereas for the second firm, it has a shock of zero. However, since our regressions use cross-sectional variations in identifying the effect of shocks to investment deviations (or investment itself), relative to the first firm, the second firm still receives a "larger" shock, so our regressions are still consistent. To be precise, suppose we fit a regression of the form

$$y = \alpha + \beta \Delta \log z + u,$$

where u is the usual error term, and y is either $\frac{\Delta i}{k}$ or $\frac{i}{k}$. Our interest is in using $\tilde{\epsilon}$ as a regressor, which we do not observe, and is thus replaced with $\Delta \log z$ as in our empirical strategy. Then this gives us,

$$\begin{aligned}
\hat{\beta} &= \frac{cov\left(\Delta \log z, y\right)}{var\left(\Delta \log z\right)} \\
&= \frac{cov\left(\frac{1}{\theta}\left(\left(\theta + \hat{\theta}\right)\tilde{\epsilon} - \frac{1}{2}\left(\theta + \hat{\theta}\right)^2 \mathbb{V}\right), y\right)}{var\left(\frac{1}{\theta}\left(\left(\theta + \hat{\theta}\right)\tilde{\epsilon} - \frac{1}{2}\left(\theta + \hat{\theta}\right)^2 \mathbb{V}\right)\right)} \\
&= \frac{\theta}{\theta + \hat{\theta}} \frac{var\left(\tilde{\epsilon}\right)}{var\left(\tilde{\epsilon}\right) + var\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)\mathbb{V}\right) - 2cov\left(\tilde{\epsilon}, \frac{1}{2}\left(\theta + \hat{\theta}\right)^2 \mathbb{V}\right)} \left[\beta - \frac{cov\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)\mathbb{V}, y\right)}{var\left(\tilde{\epsilon}\right)}\right].
\end{aligned}$$

Note that $\tilde{\epsilon} \sim N(0, \mathbb{V})$; therefore, $\mathbb{E}[\tilde{\epsilon} \mid \mathbb{V}] = \mathbb{E}[\tilde{\epsilon}]$, implying $cov(\tilde{\epsilon}, \mathbb{V}) = 0$. Therefore, the expression above reduces to

$$\hat{\beta} = \frac{\theta}{\theta + \hat{\theta}} \frac{var(\tilde{\epsilon})}{var(\tilde{\epsilon}) + var\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)\mathbb{V}\right)} \left[\beta - \frac{cov\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)\mathbb{V}, y\right)}{var(\tilde{\epsilon})}\right]. \tag{12}$$

If we assume that $\mathbb{V}$ is homogeneous across firms, as is typically done in the literature (for example, Bloom et al. (2018); Tanaka et al. (2020)), equation 12 reduces to

$$\hat{\beta} = \frac{\theta}{\theta + \hat{\theta}} \beta.$$

In other words, our estimated $\hat{\beta}$ will be smaller than the true unbiased elasticity of investment with respect to TFP shocks β; however, the qualitative correlation will always remain the same (i.e., investment is positively correlated with realized shocks).

However, as we show in our model, $\mathbb{V}$ is heterogeneous across firms. That said, if y is investment deviations ($\frac{\Delta i}{k}$), $cov\left(\mathbb{V}, \frac{\Delta i}{k}\right) \approx 0$. This is because, through the lens of our model, firms with higher uncertainty make larger ex post mistakes in both directions. In other words, while the absolute size of investment deviations are increasing in $\mathbb{V}$, it is not correlated with $\mathbb{V}$. Therefore, the mismeasurement is not an issue for us when studying investment deviations.

For the case when y is just investment rates ($\frac{i}{k}$), our model does predict that firms that face higher uncertainty will have lower investment (i.e., $cov\left(\frac{1}{2}\left(\theta + \hat{\theta}\right)\mathbb{V}, \frac{i}{k}\right) < 0$). This negative bias will then bias us toward finding a positive correlation between investment and our measure of TFP shocks. Unfortunately, it is not possible for us to directly address this concern, given our data limitations. That said, as we note in the main text, our findings are robust to both sales and value added shocks, as well as when we include investment plans as a regressor. The latter in practice controls partly for the extra covariance term, since investment plans are themselves a proxy for $\mathbb{V}$, as we show in our theory. More importantly, our overall message is that firms are able to adjust to ex post news in a consistent way (i.e., positive surprises lead to higher-than-planned investment, and negative surprises lead to lower-than-planned investment). This is indeed what we find across all our measures for "shocks".

Relationship to Section 5 In the case of our calibration, a similar logic follows. For simplicity, again first assume that all firms have the same $\mathbb{V}$. Our calibration strategy then depends solely on the dispersion of TFP shocks being proportional to the posterior variance. Specifically, if we could observe "correct" TFP shocks, then our identification strategy would simply be to map $var(\tilde{\epsilon})$ to our model parameters (i.e., we directly observe the posterior variance). For our shocks measure, we observe, assuming that firms have the same $\mathbb{V}$, $var(\Delta \log z) = \left(\frac{\theta+\hat{\theta}}{\theta}\right)^2 var(\tilde{\epsilon})$ — but this is simply a scaled measure of the true posterior variance. As such, our indirect inference strategy will remain consistent in estimating the true amount of posterior variance.

However, as we note, $\mathbb{V}$ is heterogenous across firms, and so we cannot directly invert out $var(\tilde{\epsilon})$. Specifically, as we already derived,

$$var(\Delta \log z) = \left(\frac{\theta+\hat{\theta}}{\theta}\right)^2 \left[var(\tilde{\epsilon}) + var\left(\frac{1}{2}\left(\theta+\hat{\theta}\right)\mathbb{V}\right)\right].$$

As such, in our model calibration, we choose $var(\Delta \log z)$ as a target for calibration, rather than simply invert $var(\tilde{\epsilon})$ directly from the data.

A.3.4 Issues with the investment variable

We briefly discuss the issue with our investment variable here.

To be precise, there are three cases when this happens, as briefly summarized in Table A.4 below.

	Case I	Case II	Case III
True i and i^p	$i > 0$ and $i^p < 0$	$i < 0$ and $i^p < 0$	$i < 0$ and $i^p > 0$
Observed i and i^p	$i > 0$ and $i^p = 0$	$i = 0$ and $i^p = 0$	$i = 0$ and $i^p > 0$

Table A.4: Three cases where ΔI is mismeasured.

We account for this issue using three strategies. First, the crudest of the three, we simply drop any observation for which i or i^p is reported as zero.[32] We find that our main results are robust to this data treatment. Second, we use the fact that we observe actual capital expenses in the Annual Financial Statistics survey to impute actual (dis)investment done by the BOS firms. This would in theory address the bias generated in Case III. However, the capital expenses in the Financial Statements do not line up perfectly with the BOS due to discrepancies in accounting treatment in the two surveys. Moreover, this approach does not address Case I or II. As such, we consider this imputation method only as a robustness check. We do find that our results are robust to this alternative source of capital expenses data.[33] Third, when calibrating our model, we use an indirect inference approach similar to how we address the bias in estimating expected TFP. This method addresses all three cases.

A.4 Additional Model-Implied Figures and Tables

We report here the additional figures referenced in the main text, namely the distribution of firm uncertainty and associated descriptive statistics. Figure A.1 below plots the distribution of firm uncertainty, and the table below reports broad descriptive statistics associated with this distribution.

[32]This strategy borrows from Bachmann et al. (2017) with details from footnote 12 of their paper.

[33]This does points to an advantage of our data relative to prior data sources like the German IFO data, which does not appear to provide any information about the firm's balance sheet. Consequently, such a correction would not be possible.

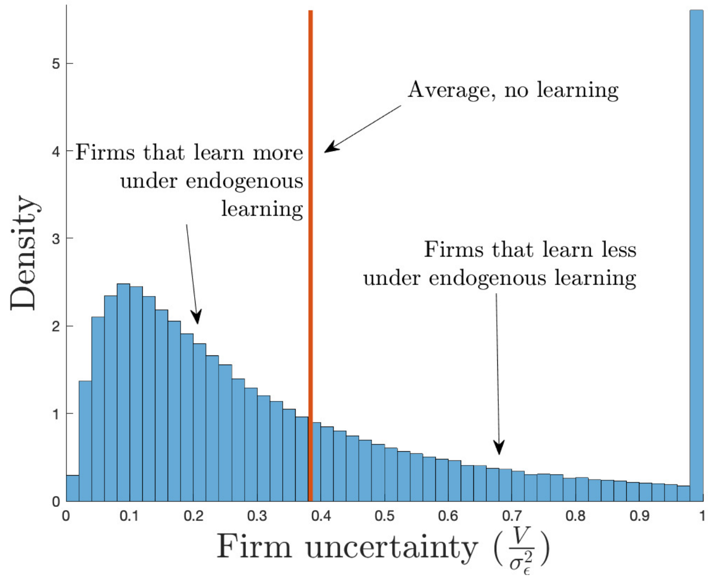

			Percentile		
Mean	SD	Skew	25^{th}	50^{th}	75^{th}
0.3756	0.3069	7.1806	0.1326	0.2665	0.5439

Figure A.1: The distribution of endogenously acquired signal precision. Solid red line is a reference for the signal precision under the "no learning" case. Note that the spike at $\frac{\mathbb{V}}{\sigma^2} = 1$ is a result of the fact that marginal benefit of learning is bounded above, and our assumption that marginal costs of increasing signal precision is constant (i.e, linear); as a result, for firms with low enough productivity, they will choose not to pay for information and simply rely on their priors.

B Model Proofs

B.1 Proof of proposition 1

To prove our proposition, we will first prove the following two lemmas, which will establish that the value of the firm, gross of the signal acquisition cost but net of adjustment costs, is strictly increasing in the posterior uncertainty. For the entirety of the proof, we will assume that $\sigma_u = \infty$ for algebraic clarity, but our proposition does not hinge on this assumption.

Lemma 1. *The expected ex ante adjustment cost is increasing in* $\mathbb{V}$.

Lemma 2. *The expected ex ante value of the firm, gross of adjustment costs, is decreasing in* $\mathbb{V}$.

We first derive some common terms that will be useful in proving the two lemmas. We begin by deriving the solution to k'. To do so, first recall that the Euler equation for investment reduces to

$$\exp\left(\frac{1}{2}\Theta_z^2\sigma_\epsilon^2\right)\Theta_k\mathcal{A}(w)\,z^{\Theta_z\rho}k'^{\Theta_k-1}-(1+r)\,\chi\frac{k'}{k^p}=r+\delta-(1+r)\,\chi. \tag{13}$$

The log-linear approximate solution for k', around the non-stochastic steady state, is therefore given by

$$\begin{aligned}\Delta k' &= \phi_z\Delta z+\phi_k\Delta k^p\\ \implies k' &= \bar{k}^{1-\phi_k}z^{\phi_z}\left(k^p\right)^{\phi_k},\end{aligned} \tag{14}$$

where for some generic variable x and steady state value $\bar{x}$, $\Delta x\equiv x-\bar{x}$; $\phi_z=\frac{\Theta_z\rho(r+\delta)}{(1+r)\chi+(r+\delta)(1-\Theta_k)}$ and $\phi_k=\frac{(1+r)\chi}{(1+r)\chi+(r+\delta)(1-\Theta_k)}$; and $\bar{k}=\left[\frac{\Theta_k\mathcal{A}(w)}{r+\delta}\right]^{\frac{1}{1-\Theta_k}}$ (with $\exp\left(\frac{1}{2}\Theta_z^2\sigma_\epsilon^2\right)\Theta_k\mathcal{A}(w)\,z^{\Theta_z\rho}k'^{\Theta_k-1}=r+\delta$ in the non-stochastic steady-state). We see that $\phi_z>0$ and $\phi_k\in(0,1)$, where in particular, $\lim_{\chi\to0}\phi_z=\frac{\rho}{1-\Theta_k}$ and $\lim_{\chi\to0}\phi_k=0$ returns us to the usual frictionless model, and $\lim_{\chi\to\infty}\phi_z=0$ and $\lim_{\chi\to\infty}\phi_k=1$ moves us to a model where only plans matter.

Next, we can substitute this solution for k' into the Bellman equation in the main text, and obtain

$$\begin{aligned}W(k,s,z_{-1},\sigma) &= \max_{k^p}\mathbb{E}\left[\pi+(1-\delta)\,k-k'-\frac{\chi}{2}\left(\frac{k'}{k^p}-1\right)^2k^p+\beta\mathbb{E}\left[V\left(k',\tilde{z}',z\right)|z\right]|s,z_{-1},\sigma\right]\\ &\quad s.t.\\ &\quad \log z\sim N\left(\rho\log z_{-1}+\frac{\sigma_\epsilon^2}{\varsigma^2+\sigma_\epsilon^2}s,\mathbb{V}\right),\end{aligned}$$

where we write k' as a function of k^p and z. This gives us the expected value of the firm after a signal has been observed, but before a plan has been made. Taking first order conditions, and following some algebra, we can obtain the optimal k^p as

$$
\begin{aligned}
k^p &= \sqrt{\mathbb{E}\left[k'^2|s\right]} \\
\implies k^p &= \sqrt{\mathbb{E}\left[\left(\bar{k}^{1-\phi_k} z^{\phi_z} \left(k^p\right)^{\phi_k}\right)^2\right]} \\
\implies k^p &= \bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \left(\sqrt{\mathbb{E}\left[\exp\left(2\phi_z \epsilon\right)|s\right]}\right)^{\frac{1}{1-\phi_k}},
\end{aligned}
$$

where ϵ is the underlying innovations of z. Noting that the posterior distribution of ϵ (i.e., after the signal s has been observed) is given by $\epsilon \sim N\left(\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2}s, \mathbb{V}\right)$, we can express k^p as

$$
\begin{aligned}
k^p &= \bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \left(\sqrt{\exp\left(2\phi_z \frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} s + 2\phi_z^2 \mathbb{V}\right)}\right)^{\frac{1}{1-\phi_k}} \\
&= \bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \exp\left(\frac{\phi_z}{1-\phi_k} \frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} s\right) \exp\left(\frac{\phi_z^2}{1-\phi_k} \mathbb{V}\right).
\end{aligned}
$$

Before moving on, it is useful to note that the planned investment is increasing in $\mathbb{V}$ for small enough s (in particular, it is always increasing in $\mathbb{V}$ when $s \to 0$). This reflects a precautionary term coming from insurance against any upside risk, which becomes increasingly dominant as s becomes smaller (in the limit, there is only upside risk and no downside risk). This term exists, in part, because our specific formulation of the adjustment cost is bounded below while unbounded above. We can now formally derive a proof for Lemma 1.

Proof. We begin by substituting the solution for k' and k^p back into the original cost function, which gives

$$
\begin{aligned}
\phi\left(k', k^p\right) &= \frac{\chi}{2}\left(\frac{k'}{k^p} - 1\right)^2 k^p \\
&= \frac{\chi}{2}\left(\exp\left(\phi_z \epsilon\right) \exp\left(-\frac{\phi_z \sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} s\right) \exp\left(-\phi_z^2 \mathbb{V}\right) - 1\right)^2 \ldots \\
&\quad \ldots \left(\bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \exp\left(\frac{1}{1-\phi_k} \frac{\phi_z \sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2} s\right) \exp\left(\frac{\phi_z^2}{1-\phi_k} \mathbb{V}\right)\right).
\end{aligned}
$$

The ex ante cost function prior to the realization of signals is therefore

$$
\begin{aligned}
\mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right] &= \mathbb{E}\left[\frac{\chi}{2}\left(\frac{k'}{k^p}-1\right)^2 k^p|\sigma\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\frac{\chi}{2}\left(\frac{k'}{k^p}-1\right)^2 k^p|s\right]|\sigma\right] \\
&= \mathbb{E}\left[\frac{\chi}{2}\mathbb{E}\left[\left(\exp\left(-\phi_z\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2}s\right)\exp\left(-\phi_z^2\mathbb{V}\right)\exp\left(\phi_z\epsilon\right)\right)^2 ...\right.\right. \\
&\qquad \left. ... - 2\exp\left(-\phi_z\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2}s\right)\exp\left(-\phi_z^2\mathbb{V}\right)\exp\left(\phi_z\epsilon\right)+1|s...\right] \\
&\qquad \left. ...\left(\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\exp\left(\frac{1}{1-\phi_k}\frac{\phi_z\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2}s\right)\exp\left(\frac{\phi_z^2}{1-\phi_k}\mathbb{V}\right)\right)|\sigma\right] \\
&= \chi\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\mathcal{F}^A\left(\mathbb{V}\right), \qquad (15)
\end{aligned}
$$

where we define $\mathcal{F}^A\left(\mathbb{V}\right) \equiv \exp\left[\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right]\left(1-\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\right)\exp\left(\frac{\phi_z^2}{2}\left(\frac{1-2\phi_k}{(1-\phi_k)^2}\right)\mathbb{V}\right)$ for notational convenience, and noting that $\chi\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} > 0$.

We see that the $\mathcal{F}^A$ term is the only term in the expression that depends on $\mathbb{V}$. To study the impact of $\mathbb{V}$ on the expected adjustment cost, it therefore suffices to study the marginal effect of changing $\mathbb{V}$ on $\mathcal{F}^A$. To do so, we can take the derivative of $\mathcal{F}^A$ with respect to $\mathbb{V}$, obtaining,

$$
\begin{aligned}
\frac{\partial\mathcal{F}^A\left(\mathbb{V}\right)}{\partial\mathbb{V}} &= \exp\left[\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right]\left(\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\frac{1}{2}\phi_z^2\exp\left(\frac{\phi_z^2}{2}\left(\frac{1-2\phi_k}{(1-\phi_k)^2}\right)\mathbb{V}\right)+...\right. \\
&\qquad \left. ...\left(1-\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\right)\exp\left(\frac{\phi_z^2}{2}\left(\frac{1-2\phi_k}{(1-\phi_k)^2}\right)\mathbb{V}\right)\frac{\phi_z^2}{2}\left(\frac{1-2\phi_k}{(1-\phi_k)^2}\right)\right) \\
&= \exp\left[\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right]\exp\left(\frac{\phi_z^2}{2}\left(\frac{1-2\phi_k}{(1-\phi_k)^2}\right)\mathbb{V}\right)\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 ... \\
&\qquad ...\left(2\phi_k-1\right)\left(\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\left(\frac{\phi_k^2}{2\phi_k-1}\right)-1\right).
\end{aligned}
$$

We now proceed to show our proof for three cases.

Case 1: $\phi_k \leq \frac{1}{2}$**.**

In this case, $2\phi_k - 1 < 0$ and likewise $\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\left(\frac{\phi_k^2}{2\phi_k-1}\right) - 1 < 0$. Therefore, it is clear from the above expression that $\frac{\partial \mathcal{F}^A(\mathbb{V})}{\partial \mathbb{V}} > 0$ for all $\mathbb{V}$. In other words, the expected adjustment cost is always strictly increasing in $\mathbb{V}$ if $\phi_k \leq \frac{1}{2}$.

Case 2: $\frac{1}{2} < \phi_k < 1$

In this case, $2\phi_k - 1 > 0$ and likewise $\frac{\phi_k^2}{2\phi_k-1} > 1$. To show that $\frac{\partial \mathcal{F}^A(\mathbb{V})}{\partial \mathbb{V}} > 0$, we need to show that $\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\left(\frac{\phi_k^2}{2\phi_k-1}\right) - 1 > 0$. This implies that $\mathbb{V} < \frac{2}{\phi_z^2}\log\left(\frac{\phi_k^2}{2\phi_k-1}\right)$, or equivalently, $\mathbb{V} < 2\left(\frac{1}{\rho(\eta-1)}\right)^2\left(\frac{1}{1-\phi_k}\right)^2\log\left(\frac{\phi_k^2}{2\phi_k-1}\right)$. Now recall that because $\mathbb{V} < \sigma_\epsilon^2$, as long as $\sigma_\epsilon^2 < 2\left(\frac{1}{\rho(\eta-1)}\right)^2\left(\frac{1}{1-\phi_k}\right)^2\log\left(\frac{\phi_k^2}{2\phi_k-1}\right)$ is satisfied, we can establish case 2. Keeping in mind that $\left(\frac{1}{1-\phi_k}\right)^2\log\left(\frac{\phi_k^2}{2\phi_k-1}\right) > 1$ when $\frac{1}{2} < \phi_k < 1$, this implies that for most reasonable calibrations, $\sigma_\epsilon^2 \ll 2\left(\frac{1}{\rho(\eta-1)}\right)^2\left(\frac{1}{1-\phi_k}\right)^2\log\left(\frac{\phi_k^2}{2\phi_k-1}\right)$. Therefore, $\frac{\partial \mathcal{F}^A(\mathbb{V})}{\partial \mathbb{V}} > 0$ for all feasible choices of $\mathbb{V}$. In other words, the expected adjustment cost is again always strictly increasing in $\mathbb{V}$ if $\frac{1}{2} < \phi_k < 1$.

Case 3: $\phi_k = 1$

Note that this is the limiting case for which $\chi \to \infty$, where we also have $\phi_z = 0$. In this limiting case, we see from the expression that $\lim_{\chi\to\infty}\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\left(\frac{\phi_k^2}{2\phi_k-1}\right) - 1 = 0$, so $\lim_{\chi\to\infty}\frac{\partial \mathcal{F}^A(\mathbb{V})}{\partial \mathbb{V}} = 0$. Intuitively, since adjustment costs are infinitely large, the marginal effect of improving information is trivially zero. In other words, we have shown that the expected cost of violating the adjustment friction is increasing in the posterior uncertainty $\mathbb{V}$. This concludes the proof for Lemma 1. □

We can also now formally derive a proof for Lemma 2.

Proof. To begin, recall that since the choice of $\mathbb{V}$ only affects next-period profits, this implies that $\mathbb{V}$ only affects the value of the firm through this channel. Recalling that expected profits net of investment cost (but gross of the adjustment cost), π^e, is given by the expression

$$\pi^e \equiv \mathbb{E}\left[-k' + \frac{1}{1+r}\left(\mathbb{E}\left[\mathcal{A}(w,Y)\, z'^{\Theta_z} k'^{\Theta_k} | z\right] + (1-\delta)\, k'\right) | \sigma\right],$$

it suffices to show that $\frac{\partial \pi^e}{\partial \mathbb{V}} < 0$ to show that higher posterior uncertainty has a negative impact on the firm's value gross of adjustment cost. To do so, it is convenient to rewrite the expected next-period profit as

$$\pi^e = \frac{1}{1+r}\mathbb{E}\left[\mathbb{E}\left[\mathcal{A}(w,Y)\, z'^{\Theta_z} k'^{\Theta_k} | z\right] - (r+\delta)\, k' | \sigma\right].$$

We will now proceed to derive this as a function of $\mathbb{V}$ (and other initial conditions and parameters). We can show that

$$\mathbb{E}\left[k'|\sigma\right] = \bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2\right) \exp\left(-\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \phi_k^2 \mathbb{V}\right)$$

and

$$\mathbb{E}\left[\mathbb{E}\left[\mathcal{A}(w)\, z'^{\Theta_z} k'^{\Theta_k} | z\right] | \sigma\right] = \mathcal{A}(w,Y) \exp\left(\frac{1}{2}\Theta_z^2 \sigma_\epsilon^2\right) \bar{k}^{\Theta_k} z_{-1}^{\rho\frac{\phi_z}{1-\phi_k}} \exp\left(\frac{1}{2}\vartheta\sigma_\epsilon^2\right) ...$$
$$... \exp\left(-\frac{1}{2}\Theta_k^2 \left(\frac{\phi_z}{1-\phi_k}\right)^2 \phi_k^2 \frac{2-\Theta_k}{\Theta_k}\mathbb{V}\right),$$

with $\vartheta \equiv (\rho\Theta_z)^2 + 2\rho\Theta_z\Theta_k\frac{\phi_z}{1-\phi_k} + \left(\Theta_k\frac{\phi_z}{1-\phi_k}\right)^2$.

Before proceeding, it is worth discussing briefly the economic intuition here. Notice that the $\mathbb{E}[k'|\sigma]$ term is decreasing in $\mathbb{V}$, that is, expected investment is decreasing in the posterior uncertainty. This contrasts with the $\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2$ term, which says that expected investment is increasing in the dispersion of productivity. This term is the usual volatility effect that predicts investment increasing with uncertainty, whereas subjective uncertainty $\mathbb{V}$ drives down expected investment. Notice however that the term here is pre-multiplied by ϕ_k, which is the weight on the investment plan in the manager's investment policy function. A more standard model of Bayesian learning would imply that $\phi_k = 1$. In our model, we nest this standard framework by allowing partial flexibility of plans. As such, the ability to weakly deviate from planned investment (i.e., $\phi_k < 1$) dampens the effect of subjective uncertainty.

Moreover, notice that $\mathbb{E}\left[\mathbb{E}\left[\mathcal{A}(w)\, z'^{\Theta_z} k'^{\Theta_k} | z\right] | \sigma\right]$ (expected revenue) is also decreasing in

$\mathbb{V}$. Since expected investment is decreasing in $\mathbb{V}$, this result is not surprising. Now with these two terms in hand, we can rewrite the expected profits as

$$\begin{aligned}\pi^{e} &= \frac{1}{1+r}\left[\mathcal{A}\left(w,Y\right)\exp\left(\frac{1}{2}\Theta_{z}^{2}\sigma_{\epsilon}^{2}\right)\bar{k}^{\Theta_{k}}z_{-1}^{\rho\frac{\phi_{z}}{1-\phi_{k}}}\exp\left(\frac{1}{2}\vartheta\sigma_{\epsilon}^{2}\right)...\right.\\ &\quad ...\exp\left(-\frac{1}{2}\Theta_{k}^{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\phi_{k}^{2}\frac{2-\Theta_{k}}{\Theta_{k}}\mathbb{V}\right)-\\ &\quad ...\left(r+\delta\right)\bar{k}z_{-1}^{\frac{\rho\phi_{z}}{1-\phi_{k}}}\exp\left(\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\sigma_{\epsilon}^{2}\right)\\ &\quad \left....\exp\left(-\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\phi_{k}^{2}\mathbb{V}\right)\right]\\ &= \frac{1}{1+r}\bar{k}z_{-1}^{\frac{\rho\phi_{z}}{1-\phi_{k}}}\exp\left(\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\sigma_{\epsilon}^{2}\right)\mathcal{F}^{\pi}\left(\mathbb{V}\right),\end{aligned} \tag{16}$$

where $\mathcal{F}^{\pi}\left(\mathbb{V}\right)$ is a function of the posterior uncertainty and other model parameters. Critically, this function does not include z_{-1}, which pre-multiplies this function. Therefore, the effect of $\mathbb{V}$ is always scaled by z_{-1}. To show that increasing posterior uncertainty decreases expected profits, we simply need to show that $\frac{\partial\mathcal{F}^{\pi}(\mathbb{V})}{\partial\mathbb{V}}<0$. We derive

$$\begin{aligned}\frac{\partial\mathcal{F}^{\pi}\left(\mathbb{V}\right)}{\partial\mathbb{V}} &= \left(\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\phi_{k}^{2}\right)\exp\left(-\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\phi_{k}^{2}\mathbb{V}\right)...\\ &\quad \left[\mathcal{A}\left(w,Y\right)\exp\left(\frac{1}{2}\Theta_{z}^{2}\sigma_{\epsilon}^{2}\right)\bar{k}^{\Theta_{k}-1}\exp\left(\frac{1}{2}\left(\vartheta-\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\right)\sigma_{\epsilon}^{2}\right)...\right.\\ &\quad \left....\left(-\Theta_{k}\left(2-\Theta_{k}\right)\right)\exp\left(\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\phi_{k}^{2}\left(1-\Theta_{k}\left(2-\Theta_{k}\right)\right)\mathbb{V}\right)+\left(r+\delta\right)\right],\end{aligned}$$

which means that for the previous condition to hold, we need the following condition

$$\begin{aligned}&\mathcal{A}\left(w,Y\right)\exp\left(\frac{1}{2}\Theta_{z}^{2}\sigma_{\epsilon}^{2}\right)\bar{k}^{\Theta_{k}-1}\exp\left(\frac{1}{2}\left(\vartheta-\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\right)\sigma_{\epsilon}^{2}\right)...>r+\delta\\ &...\left(\Theta_{k}\left(2-\Theta_{k}\right)\right)\exp\left(\frac{1}{2}\left(\frac{\phi_{z}}{1-\phi_{k}}\right)^{2}\phi_{k}^{2}\left(1-\Theta_{k}\left(2-\Theta_{k}\right)\right)\mathbb{V}\right).\end{aligned}$$

But recall that $\bar{k} = \left[\frac{\Theta_k \mathcal{A}(w)}{r+\delta}\right]^{\frac{1}{1-\Theta_k}}$, which means that $\bar{k}^{\Theta_k - 1} = \frac{r+\delta}{\Theta_k \mathcal{A}(w,Y)}$. Substituting this back into the equation above, we get

$$\exp\left(\frac{1}{2}\Theta_z^2\sigma_\epsilon^2\right)\exp\left(\frac{1}{2}\left(\rho\Theta_z\right)^2\frac{\Theta_k}{1-\Theta_k}\sigma_\epsilon^2\right)\left(2-\Theta_k\right)\exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\phi_k^2\left(1-\Theta_k\left(2-\Theta_k\right)\right)\mathbb{V}\right) > 1.$$

Further substituting in $\vartheta - \left(\frac{\phi_z}{1-\phi_k}\right)^2 = \left(\rho\Theta_z\right)^2\frac{\Theta_k}{1-\Theta_k}$, we reduce the expression to,

$$\exp\left(\frac{1}{2}\Theta_z^2\sigma_\epsilon^2 + \frac{1}{2}\left(\rho\Theta_z\right)^2\frac{\Theta_k}{1-\Theta_k}\sigma_\epsilon^2 + \frac{1}{2}\left(\frac{\rho\Theta_z}{1-\Theta_k}\right)^2\phi_k^2\left(1-\Theta_k\left(1-\frac{1}{2}\Theta_k\right)\right)\mathbb{V}\right) > \frac{1}{2-\Theta_k}.$$

Notice that since $\Theta_k \in (0,1)$, $\frac{1}{2-\Theta_k} \in \left(\frac{1}{2},1\right)$, the above relation is trivially true for all parameter values, so $\frac{\partial\mathcal{F}^\pi(\mathbb{V})}{\partial\mathbb{V}} < 0$. Therefore, expected profits, gross of adjustment cost, is decreasing in the posterior variance. This concludes the proof for Lemma 2. □

With Lemma 1 and 2 in hand, we can now derive a proof for our proposition.

Proof. First, Lemma 2 tells us that the ex ante value of the firm, gross of signal acquisition costs but net of adjustment cost, is strictly decreasing in $\mathbb{V}$. This gives us point 1 in Proposition 1.

To show part 2 of Proposition 1, we can simply combine both the expected adjustment cost and expected profits and take the first derivative with respect to $\mathbb{V}$. To recall, the first derivative of the expected adjustment cost is

$$\frac{\partial\mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial\mathbb{V}} = \chi\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\frac{\partial\mathcal{F}^A\left(\mathbb{V}\right)}{\partial\mathbb{V}},$$

and recall $\frac{\partial\mathcal{F}^A(\mathbb{V})}{\partial\mathbb{V}} > 0$. The marginal effect of $\mathbb{V}$ on firm value, net of adjustment cost, can then be expressed as

$$-\frac{\partial \mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial \mathbb{V}}+\frac{\partial \pi^e}{\partial \mathbb{V}} = -\chi \bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \frac{\partial \mathcal{F}^A\left(\mathbb{V}\right)}{\partial \mathbb{V}} + \frac{1}{1+r}\bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2\right)\frac{\partial \mathcal{F}^\pi\left(\mathbb{V}\right)}{\partial \mathbb{V}}$$

$$= \bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} \left(-\chi \frac{\partial \mathcal{F}^A\left(\mathbb{V}\right)}{\partial \mathbb{V}} + \frac{1}{1+r} \exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2\right)\frac{\partial \mathcal{F}^\pi\left(\mathbb{V}\right)}{\partial \mathbb{V}}\right)$$

where we see that the term $-\chi \frac{\partial \mathcal{F}^A(\mathbb{V})}{\partial \mathbb{V}} + \frac{1}{1+r} \exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2\right)\frac{\partial \mathcal{F}^\pi(\mathbb{V})}{\partial \mathbb{V}} < 0$, and does not depend on z_{-1}. In other words, initial productivity z_{-1} has a pure scaling effect—that is—firms with higher initial productivity face a steeper cost of having a more dispersed signal. Conversely, the benefits to get a better signal is increasing in initial productivity. Formally, this statement is seen in the cross-derivative, which is given by

$$\frac{\partial}{\partial z_{-1}}\left(-\frac{\partial \mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial \mathbb{V}}+\frac{\partial \pi^e}{\partial \mathbb{V}}\right) = \frac{\rho\phi_z}{1-\phi_k}\bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}-1} \dots$$

$$\dots \left(\chi \frac{\partial \mathcal{F}^A\left(\mathbb{V}\right)}{\partial \mathbb{V}} + \frac{1}{1+r} \exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2\right)\frac{\partial \mathcal{F}^\pi\left(\mathbb{V}\right)}{\partial \mathbb{V}}\right)$$

$$< 0$$

Therefore, the benefits of having a lower posterior uncertainty is increasing in initial firm productivity z_{-1}. This thus concludes the proof of point 2 in Proposition 1. □

B.1.1 Proof of Corollary 1

From Proposition 1, we already showed that the marginal benefit of a lower posterior variance is increasing in z_{-1}. The marginal cost is given by

$$-\frac{\partial}{\partial \mathbb{V}}\xi\left(\frac{1}{\mathbb{V}}-\frac{1}{\sigma_\epsilon^2}\right) = \xi\left(\frac{1}{\mathbb{V}}\right)^2$$

where the cost of lowering posterior uncertainty is increasing, but does not depend on z_{-1}.[34] Equating the marginal cost and benefit, and with some trivial rearrangement of terms,

[34]Here, we directly impose that $\sigma_u = \infty$.

gives us

$$\bar{k} z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} = -\xi \left(\frac{1}{\mathbb{V}}\right)^2 \left(-\chi \frac{\partial \mathcal{F}^A(\mathbb{V})}{\partial \mathbb{V}} + \frac{1}{1+r} \exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2 \sigma_\epsilon^2\right) \frac{\partial \mathcal{F}^\pi(\mathbb{V})}{\partial \mathbb{V}}\right)^{-1},$$

where $\mathbb{V}$ solves the implicit equation above. Note that the left-hand side term is strictly increasing in z_{-1}; and the right-hand side term is strictly decreasing in $\mathbb{V}$ (i.e., decreasing marginal benefit), and importantly, not a function of z_{-1}. Therefore, an application of the inverse function theorem tell us that the choice of $\mathbb{V}$ is decreasing in z_{-1}. Since the posterior variance is decreasing in z_{-1}, trivially then, the dispersion of forecast errors is decreasing in initial productivity.

B.2 Proof Of Proposition 2

Like in our stationary model, we begin our proof by deriving explicit analytical forms for the ex ante expected gross profits at the firm level. First, we derive the log-linear solution to k'. From the Euler Equation,

$$\exp\left(\frac{1}{2}\Theta_z^2\sigma_\epsilon^2 + \frac{1}{2}\Theta_A^2\sigma_{\epsilon,A}^2\right)\Theta_k \mathcal{A}(w,Y)\, z^{\Theta_z\rho} A^{\Theta_A\rho_A} k'^{\Theta_k - 1} - (1+r)\chi\frac{k'}{k^p} = r+\delta-(1+r)\chi. \tag{17}$$

The log-linear approximate solution for k', around the non-stochastic steady state, is therefore given by

$$\begin{aligned} \hat{k}' &= \phi_A \hat{A} + \phi_z \hat{z} + \phi_k \widehat{k^p} \\ \implies k' &= \bar{k}^{1-\phi_k} A^{\phi_A} z^{\phi_z} \left(k^p\right)^{\phi_k}, \end{aligned} \tag{18}$$

where $\phi_A = \frac{\Theta_A\rho_A(r+\delta)}{(1+r)\chi+(r+\delta)(1-\Theta_k)}$, $\phi_z = \frac{\Theta_z\rho(r+\delta)}{(1+r)\chi+(r+\delta)(1-\Theta_k)}$ and $\phi_k = \frac{(1+r)\chi}{(1+r)\chi+(r+\delta)(1-\Theta_k)}$, and $\bar{k} = \left[\frac{\Theta_k\mathcal{A}(w)}{r+\delta}\right]^{\frac{1}{1-\Theta_k}}$. Notice that this solution is very similar to the solution for the model aggregate risk. With this solution for k', we can substitute the solution back into the Bellman equation, obtaining

$$W\left(k,s,A,z_{-1},\sigma\right)=\max_{k^p}\mathbb{E}\left[\pi+\left(1-\delta\right)k-k'-\frac{\chi}{2}\left(\frac{k'}{k^p}-1\right)^2k^p...\right.$$
$$\left....+\beta\mathbb{E}\left[V\left(k',\tilde{z}',\tilde{A}',z,A\right)|z,A\right]|s,s_A,z_{-1},A_{-1},\sigma,\sigma_A\right]$$
$$s.t.$$
$$\log z\sim N\left(\rho\log z_{-1}+\frac{\sigma_\epsilon^2}{\varsigma^2+\sigma_\epsilon^2}s,\mathbb{V}\right)$$
$$\log A\sim N\left(\rho_{-1}\log A_{-1}+\frac{\sigma_{\epsilon,A}^2}{\sigma_A^2+\sigma_{\epsilon,A}^2}s_A,\mathbb{V}_A\right),$$

where s_A is the signal associated with aggregate conditions, and $\mathbb{V}_A$ is the endogenously chosen posterior variance of A. Taking first order conditions and following some algebra, we can obtain the optimal k^p as

$$k^p=\sqrt{\mathbb{E}\left[k'^2|s,s_A\right]}$$
$$\implies k^p=\sqrt{\mathbb{E}\left[\left(\bar{k}^{1-\phi_k}A^{\phi_A}z^{\phi_z}\left(k^p\right)^{\phi_k}\right)^2\right]}$$
$$\implies k^p=\bar{k}^{1-\phi_k}A_{-1}^{\rho_A\phi_A}z_{-1}^{\rho\phi_z}\left(k^p\right)^{\phi_k}\sqrt{\mathbb{E}\left[\exp\left(2\phi_A\epsilon_A\right)|s_A\right]}\sqrt{\mathbb{E}\left[\exp\left(2\phi_z\epsilon\right)|s\right]}$$
$$\implies k^p=\bar{k}A_{-1}^{\frac{\rho_A\phi_A}{1-\phi_k}}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\left(\sqrt{\mathbb{E}\left[\exp\left(2\phi_A\epsilon_A\right)|s_A\right]}\right)^{\frac{1}{1-\phi_k}}\left(\sqrt{\mathbb{E}\left[\exp\left(2\phi_z\epsilon\right)|s\right]}\right)^{\frac{1}{1-\phi_k}},$$

where ϵ_A is the underlying innovations of A. Similar to ϵ, the posterior distribution of ϵ_A is given by $\epsilon_A\sim N\left(\frac{\sigma_{\epsilon,A}^2}{\sigma_A^2+\sigma_{\epsilon,A}^2}s_A,\mathbb{V}_A\right)$, and we can express k^p as

$$k^p=\bar{k}A_{-1}^{\frac{\rho_A\phi_A}{1-\phi_k}}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\left(\sqrt{\exp\left(2\phi_A\frac{\sigma_{\epsilon,A}^2}{\sigma_A^2+\sigma_{\epsilon,A}^2}s_A+2\phi_A^2\mathbb{V}_A\right)}\sqrt{\exp\left(2\phi_z\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2}s+2\phi_z^2\mathbb{V}\right)}\right)^{\frac{1}{1-\phi_k}}$$
$$=\bar{k}A_{-1}^{\frac{\rho_A\phi_A}{1-\phi_k}}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\exp\left(\frac{\phi_A}{1-\phi_k}\frac{\sigma_{\epsilon,A}^2}{\sigma_A^2+\sigma_{\epsilon,A}^2}s_A+\frac{\phi_z}{1-\phi_k}\frac{\sigma_\epsilon^2}{\sigma^2+\sigma_\epsilon^2}s+\frac{\phi_A^2}{1-\phi_k}\mathbb{V}_A+\frac{\phi_z^2}{1-\phi_k}\mathbb{V}\right).$$

Like the solution for k', because of our assumption that aggregate conditions are orthogonal to idiosyncratic conditions, we have an expression for k^p that is very much similar to that in the model without aggregate risk. We obtain (similar to our stationary model) expressions for the ex ante cost function,

$$\begin{aligned}\mathbb{E}\left[\phi\left(k', k^p\right)|\sigma\right] &= \mathbb{E}\left[\frac{\chi}{2}\left(\frac{k'}{k^p}-1\right)^2 k^p|\sigma\right] \\ &= \chi\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}} A_{-1}^{\frac{\rho_A\phi_A}{1-\phi_k}} \mathcal{F}^s\left(\mathbb{V}\right)\mathcal{F}^A\left(\mathbb{V}_A\right),\end{aligned} \tag{19}$$

where

$$\begin{aligned}\mathcal{F}^s\left(\mathbb{V}\right) &= \exp\left[\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right]\left(1-\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)\right)\exp\left(\left(\frac{\phi_z^2-2\phi_z^2\phi_k}{2\left(1-\phi_k\right)^2}\right)\mathbb{V}\right) \\ \mathcal{F}^A\left(\mathbb{V}_A\right) &= \exp\left[\frac{1}{2}\left(\frac{\phi_A}{1-\phi_k}\right)^2\sigma_{\epsilon,A}^2\right]\left(1-\exp\left(-\frac{1}{2}\phi_A^2\mathbb{V}_A\right)\right)\exp\left(\left(\frac{\phi_A^2-2\phi_A^2\phi_k}{2\left(1-\phi_k\right)^2}\right)\mathbb{V}_A\right).\end{aligned}$$

Note that $\chi\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\mathcal{F}^s\left(\mathbb{V}\right)$ is simply our expression for the ex ante cost function when there is no aggregate risk. Likewise, expected profits net of investment cost (but gross of the adjustment cost), π^e, can be expressed as

$$\begin{aligned}\pi^e = &\frac{1}{1+r}\bar{k}z_{-1}^{\frac{\rho\phi_z}{1-\phi_k}}\exp\left(\frac{1}{2}\left(\frac{\phi_z}{1-\phi_k}\right)^2\sigma_\epsilon^2\right)\mathcal{F}^\pi\left(\mathbb{V}\right)... \\ &...A_{-1}^{\frac{\rho_A\phi_A}{1-\phi_k}}\exp\left(\frac{1}{2}\left(\frac{\phi_A}{1-\phi_k}\right)^2\sigma_{\epsilon,A}^2\right)\mathcal{F}_A^\pi\left(\mathbb{V}_A\right),\end{aligned} \tag{20}$$

with

$$\begin{aligned}\mathcal{F}_A^\pi\left(\mathbb{V}_A\right) = &\mathcal{A}\left(w,Y\right)\exp\left(\frac{1}{2}\Theta_A^2\sigma_{\epsilon,A}^2\right)\bar{k}^{\Theta_k-1}\exp\left(\frac{1}{2}\left(\vartheta_A-\left(\frac{\phi_A}{1-\phi_k}\right)^2\right)\sigma_{\epsilon,A}^2\right)... \\ &...\exp\left(-\frac{1}{2}\Theta_k^2\left(\frac{\phi_A}{1-\phi_k}\right)^2\phi_k^2\frac{2-\Theta_k}{\Theta_k}\mathbb{V}_A\right)... \\ &...\left(r+\delta\right)\exp\left(-\frac{1}{2}\left(\frac{\phi_A}{1-\phi_k}\right)^2\phi_k^2\mathbb{V}_A\right),\end{aligned}$$

and $\vartheta_A \equiv (\rho\Theta_A)^2 + 2\rho\Theta_A\Theta_k\frac{\phi_A}{1-\phi_k} + \left(\Theta_k\frac{\phi_A}{1-\phi_k}\right)^2$. $\mathcal{F}^\pi(\mathbb{V})$ is the same function as that in the stationary case. The expressions in equations 19 and 20 are simply the formula in the stationary case, extended to account for aggregate shocks. It becomes trivial, given our earlier derivations, to see that $\frac{\partial\mathbb{E}[\phi(k',k^p)|\sigma]}{\partial\mathbb{V}_A} > 0$, and $\frac{\partial\pi^e}{\partial\mathbb{V}_A} < 0$, and therefore,

$$-\frac{\partial\mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial\mathbb{V}_A} + \frac{\partial\pi^e}{\partial\mathbb{V}_A} < 0.$$

That is, the ex ante value of the firm (gross of signal acquisition cost) is increasing in signal precision over aggregate conditions. This gives us point 1 of Proposition 2. For the next two points, it is also straightforward, given our earlier derivations, to see that

$$\frac{\partial}{\partial A_{-1}}\left(-\frac{\partial\mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial\mathbb{V}} + \frac{\partial\pi^e}{\partial\mathbb{V}}\right) < 0$$

and

$$\frac{\partial}{\partial A_{-1}}\left(-\frac{\partial\mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial\mathbb{V}_A} + \frac{\partial\pi^e}{\partial\mathbb{V}_A}\right) < 0.$$

That is, the value of increasing signal precision over idiosyncratic and aggregate productivity is increasing in initial aggregate productivity. This gives us points 2 and 3 of Proposition 2, and therefore concludes the proof.

B.2.1 Proof Of Corollary 2

This proof is exactly the same as that for Corollary 1. To be precise, because we assume that the signal acquisition cost over idiosyncratic and aggregate productivity are linearly additive, we can derive similar first-order conditions to the earlier proof, namely,

$$\frac{\partial}{\partial\mathbb{V}}\xi\left(\frac{1}{\mathbb{V}} - \frac{1}{\sigma_\epsilon^2}\right) = \frac{\partial\mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial\mathbb{V}} + \frac{\partial\pi^e}{\partial\mathbb{V}},$$

where we equate the marginal cost and marginal benefit of acquiring better signals over idiosyncratic productivity. Likewise, for aggregate productivity, we have

$$\frac{\partial}{\partial \mathbb{V}_A}\xi_A\left(\frac{1}{\mathbb{V}_A}-\frac{1}{\sigma_{\epsilon,A}^2}\right)=\frac{\partial \mathbb{E}\left[\phi\left(k',k^p\right)|\sigma\right]}{\partial \mathbb{V}_A}+\frac{\partial \pi^e}{\partial \mathbb{V}_A}.$$

Similar to the proof earlier, these two relations give us an implicit solution for $\mathbb{V}$ and $\mathbb{V}_A$ where the optimal choice is decreasing in A_{-1}, that is to say, $\frac{\partial \mathbb{V}}{\partial A_{-1}}<0$,and $\frac{\partial \mathbb{V}_A}{\partial A_{-1}}<0$. This concludes the proof of Corollary 2.

B.2.2 Proof Of Corollary 3

To prove this corollary, we follow from Corollary 2 and note that the inequalities $\frac{\partial \mathbb{V}}{\partial A_{-1}}<0$ and $\frac{\partial \mathbb{V}_A}{\partial A_{-1}}<0$ are true for all z_{-1} (i.e., all firms). Therefore, this also implies that

$$\mathbb{E}\left[\frac{\partial \mathbb{V}}{\partial A_{-1}}|A_{-1}\right]<0\implies\frac{\partial \mathbb{E}\left[\sigma|A_{-1}\right]}{\partial A_{-1}}<0\implies\frac{\partial \mathbb{E}\left[\mathbb{V}|A_{-1}\right]}{\partial A_{-1}}<0,$$

where the term $\mathbb{E}\left[\sigma|A_{-1}\right]$ is nothing but the unweighted average idiosyncratic uncertainty. In other words,

$$cov\left(\mathbb{E}\left[\sigma|A_{-1}\right],A_{-1}\right)<0.$$

Recall that A follows an AR(1) process. Therefore, we have

$$\begin{aligned}cov\left(\mathbb{E}\left[\sigma|A_{-1}\right],\log A\right)&=cov\left(\mathbb{E}\left[\sigma|A_{-1}\right],\rho_A\log A_{-1}+\sigma_{\epsilon,A}\epsilon_A\right)\\&=cov\left(\mathbb{E}\left[\sigma|A_{-1}\right],\rho_A\log A_{-1}\right)\\&=\rho_A cov\left(\mathbb{E}\left[\sigma|A_{-1}\right],\log A_{-1}\right)\\&<0,\end{aligned}$$

that is to say, average idiosyncratic uncertainty is countercyclical. To show that $\bar{\sigma}$ is persistent, recall that $\sigma=\sigma\left(z_{-1},A_{-1}\right)$, implying that $\mathbb{E}\left[\sigma|A\right]=\bar{\sigma}\left(A_{-1}\right)$. Trivially, $cov\left(\bar{\sigma}_{-1},\bar{\sigma}\right)=$

$cov\left(\bar{\sigma}\left(A_{-2}\right),\bar{\sigma}\left(A_{-1}\right)\right)>0$,which shows that average idiosyncratic uncertainty is persistent. The proof of countercyclical uncertainty for aggregate conditions, as well as its persistence, are exactly the same as that for the idiosyncratic case.

B.2.3 Proof of Prediction 5

Here, we prove that misallocation, as measured by the dispersion in marginal product of capital, is countercyclical. To begin, note $\mathbb{E}\left[\tau^2\right]$ (the implicit "wedge" as discussed in the main text), can be reinterpreted as the dispersion of ARPK given (A_{-1}, z_{-1}). Next, note that $\mathbb{E}\left[\tau^2|A\right]=\mathbb{E}\left[\tau^2|A_{-1}\right]$ because $\frac{k'}{k^p}$ does not depend on A. Therefore, the dispersion of ARPK (in levels), which is $\mathbb{E}\left[\mathbb{E}\left[\tau^2|A,z_{-1}\right]|A\right]$ is given by,

$$\mathbb{E}\left[\mathbb{E}\left[\tau^2|A,z_{-1}\right]|A\right]=\left(\left(1+r\right)\chi\right)^2\left(2-2\mathbb{E}\left[\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)|A_{-1}\right]\right).$$

We now need to establish that $\frac{\partial\mathbb{E}\left[\mathbb{E}\left[\tau^2|A,z_{-1}\right]|A\right]}{\partial A}<0$. Note that

$$\frac{\partial\mathbb{E}\left[\mathbb{E}\left[\tau^2|A,z_{-1}\right]|A\right]}{\partial A}=-2\frac{\left(\left(1+r\right)\chi\right)^2}{\rho A_{-1}^{\rho-1}\exp\left(\sigma_{\epsilon,A}\epsilon_A\right)}\frac{\partial\mathbb{E}\left[\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)|A_{-1}\right]}{\partial A_{-1}},$$

so it is sufficient to just show that $\frac{\partial\mathbb{E}\left[\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)|A_{-1}\right]}{\partial A}>0$. As we already showed in deriving the proof of Corollary 3, $\frac{\partial\mathbb{E}\left[\mathbb{V}|A_{-1}\right]}{\partial A_{-1}}<0$. Therefore, this implies that

$$\begin{aligned}\frac{\partial\mathbb{E}\left[\exp\left(-\frac{1}{2}\phi_z^2\mathbb{V}\right)|A_{-1}\right]}{\partial A_{-1}}&<\frac{\partial\exp\left(-\frac{1}{2}\phi_z^2\mathbb{E}\left[\mathbb{V}|A_{-1}\right]\right)}{\partial A_{-1}}\\&=-\frac{1}{2}\phi_z^2\exp\left(-\frac{1}{2}\phi_z^2\mathbb{E}\left[\mathbb{V}|A_{-1}\right]\right)\frac{\partial\mathbb{E}\left[\mathbb{V}|A_{-1}\right]}{\partial A_{-1}}\\&>0,\end{aligned}$$

where the first line follows from Jensen's inequality, and the second line uses the fact that $\exp\left(-\frac{1}{2}\phi_z^2\mathbb{E}\left[\mathbb{V}|A_{-1}\right]\right)>0$ while $\frac{\partial\mathbb{E}\left[\mathbb{V}|A_{-1}\right]}{\partial A_{-1}}<0$. Therefore, the dispersion of ARPK is decreasing in A. That is, misallocation is countercyclical.

C Evidence On countercyclicality of Uncertainty

We test our predictions in our data. Dispersion of TFP forecast errors and ARPK, at the aggregate level, is computed as the log of the inter-quartile range of TFP forecast errors (or log ARPK). For dispersion of forecast errors at the firm level, we proxy for dispersion using the log of one plus the absolute value of firm TFP forecast errors. Finally, for "aggregate" variables, we compute them at two levels of aggregation: (i) at the annual level, and (ii) at the industry-year level. Our second specification allows us to address the fact that we only have twelve observations at the annual level.

To test our predictions at the aggregate industry-year level, we run regressions of the form:

$$y_{j,t} = \alpha_{j(i)} + x_t + \varepsilon_{j,t}, \tag{21}$$

where i indexes a firm, t indexes a fiscal year, $j(i)$ indexes the MoF industry group that a firm is in, and $\alpha_{j(i)}$ capture industry fixed effects. The dependent variable x_t is either a recession indicator from the St. Louis Federal Reserve when studying countercyclicality, or lagged $y_{i,t}$ when studying persistence. For regressions at the yearly aggregation, we simply modify equation 21 as appropriate.

To test our predictions at the firm-level, we use a specification of the form:

$$y_{i,t} = \alpha_i + \alpha_{j(i)} + x_t + \varepsilon_{i,t}, \tag{22}$$

where the subscripts follow from Equation 21, and the dependent variable x_t is again either a recession indicator or lagged $y_{i,t}$, where appropriate. We consider specifications with both industry or firm fixed effects (α_i), and standard errors are clustered by firm and year to account for within-firm correlated error dispersions and cross-sectionally correlated error dispersions due to common aggregate shocks.

Table C.1: Cyclicality and Persistence of TFP Shock Dispersions

The table below shows the dispersion of TFP shocks depending on whether the economy is in a recession. Observations are either at the year, industry-by-year level, or firm-by-year level. Recessions are defined according to data from the St. Louis Federal Reserve. The estimated constants in the regressions with no fixed effects are suppressed for space. Where there is no confusion, we drop unnecessary subscripts. Expectations, actual values, and shocks are winsored at the 1% level. Standard errors are clustered by industry and year in Panel A and firm and year in Panel B, and are shown in parentheses below the estimated coefficient. * p < 0.10, ** p < 0.05, *** p < 0.01.

Panel A: countercyclicality of TFP Shock Dispersion			
Dependent Var:	Dispersion of $\Delta \log z$		
	(1)	(2)	(3)
Recession	0.102*	0.066***	0.036***
	(0.047)	(0.020)	(0.003)
Observation Level	Year	Industry-Year	Firm-Year
Fixed Effects		Industry	Firm
Observations	12	103	26,718
R^2	0.323	0.475	0.540
Panel B: Autocorrelation of TFP Shock Dispersion			
Dependent Var:	y = Dispersion of $\Delta \log z$		
	(1)	(2)	(3)
y_{-1}	0.140	0.058	0.004***
	(0.259)	(0.083)	(0.001)
Observation Level	Year	Industry-Year	Firm-Year
Fixed Effects		Industry	Firm
Observations	11	94	24,541
R^2	0.031	0.595	0.264

Prediction 3. The dispersion of forecast errors is persistent and countercyclical, both at the firm- and aggregate- levels. Columns 1 and 2 of Panel A of Table C.1 show that dispersion in forecast errors at aggregate levels is countercyclical, while column 3 of the same panel shows that dispersion in forecast errors at the firm level is countercyclical. Specifically, we see that in all three cases, the dispersion in TFP forecast errors is higher in recessionary times relative to times of expansion.

Columns 1 and 2 of Panel B of Table C.1 shows that the dispersion in forecast errors at the aggregate level is persistent, while column 3 of the same panel shows that dispersion in forecast errors at the firm level is persistent. Unfortunately, our short time series makes it impossible for us to verify this (in a statistically significant sense) at the aggregate level, but the coefficient remains positive. Moreover, since this behavior at the firm level is what ultimately drives our aggregate results, we see this as corroborating model Prediction 4.

Prediction 4. The dispersion of ARPK is countercyclical. Table C.2 shows that the dispersion of the ARPK is higher in recessionary times relative to times of expansion. Unfortunately, given the length of our time series is only 12 years, column 1 using observations at the year level is not statistically significant with a noisy point estimate. However, column 2 shows that the dispersions at the industry-by-year level is statistically significant and positive at the 10% level.

Table C.2: Cyclicality of ARPK Dispersions

The table below shows the dispersion of log ARPK depending on whether the economy is in a recession. Observations are either at the year or industry-by-year level. When using industry-by-year level observations, we include industry fixed effects. Recessions are defined according to data from the St. Louis Federal Reserve. Where there is no confusion, we drop unnecessary subscripts. Expectations, actual values, and shocks are winsored at the 1% level. Standard errors are are shown in parentheses below the estimated coefficient. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Fact 5 – countercyclicality of ARPK Dispersions		
Dependent Var:	Dispersion of ARPK	
	(1)	(2)
Recession	-0.047	0.012*
	(0.036)	(0.007)
Observation Level	Year	Industry-Year
Fixed Effects		Industry
Observations	12	103
R^2	0.147	0.841

Not Intended For Publication

D Details of the Business Outlook Survey

The Ministry of Finance (MoF) and Cabinet Office of Japan conduct the Business Outlook Survey (BOS) in order to help forecast the economy and inform fiscal planning. The survey is implemented primarily to evaluate the investment condition of firms and is conducted separately from other Japanese survey data that have been used by existing researchers such as the Tankan survey conducted by the Bank of Japan that "aims to provide an accurate picture of business trends of enterprises in Japan, thereby contributing to the appropriate implementation of monetary policy". The MoF uses some information from the BOS in producing its Monthly Economic Report, which is made publicly available in the form of aggregated statistics. None of the disaggregated individual firm-level data ever become public.

The BOS has been conducted from the first quarter of Fiscal Year (FY) 2004 as a General Statistical Survey under the Statistics Act. The first Statistics Act (Act No. 18 of 1947) was later revised in Act No. 53, which was passed on May 23, 2007 and implemented in April 2009. The survey targets are non-financial corporations with paid-in-capital of at least 10 million yen (approx. 100,000 USD) and utilities and financial institutions with paid-in-capital of at least 100 million yen (approx. 1 million USD). Responses are collected both via mail and online, and the sampling is based on the corporations covered by quarterly surveys of Financial Statements Statistics of Corporations by Industry. Industries in the aggregated survey result are based on Japanese SIC (J-SIC) codes (the 2-digit figure for manufacturing firms and larger "alphabet" group classifications for non-manufacturing industries) and grouped into 45 MoF industries. This means that some 2-digit J-SIC industries are grouped into a larger classification for reporting. In all our analyses, we use these MoF industry definitions as our definition of industries.

The survey is administered as repeated cross sections with stratified sampling across seven categories by registered capital and industry to be representative of the Japanese economy. The seven categories are (1) 0.01 - 0.02 billion, (2) 0.02 - 0.05 billion, (3) 0.05 - 0.1 billion, (4) 0.1 - 0.5 billion, (5) 0.5 - 1 billion, (6) 1 - 2 billion, (7) over 2 billion. Before FY2010, the border between (4) and (5) was 0.6 billion yen. In the aggregated results, firms are grouped by size into three categories: "small-medium" corporations (1, 2, 3), "medium-sized" corporations (4, 5) and "large" corporations (6, 7). In addition, the Twelfth revision of the J-SIC (November 2007) was enforced on April 1, 2008, and the stratification by industry changed from the first quarter (April - June) of FY 2009. Originally, the sampling was stratified into 43 categories by industry for non-financial corporations.

Table D.1 shows the sampling probabilities for each broad size strata. Due to the high sampling probability among middle to large firms, we are able to construct a panel. The overall response rate is nearly 80% on average and nears 90% for large firms. Table D.2 shows the average annual response rates by year from FY 2005 to FY 2016. Firms are sampled on an annual basis, aligned with standard fiscal period ends for Japanese firms. Firms that are sampled are assigned a unique company identifier that is only for use in the MOF. The disaggregated information and responses at the firm-level do not leave the MoF building and are stored on air-gapped computers.

The BOS asks both qualitative and quantitative questions quarterly. Surveys start in the late of first month in each quarter. Firms are required to answer their forecasts and plans in the middle of the second month in each quarter as a reference date. Qualitative questions ask about the business condition, domestic economic conditions, employment, and others. These questions provide options like "up", "same", "down", and "unknown". For example, "what is your business condition of the current quarter (the next quarter, the one after next) compared with the previous one?" Respondents can answer one of the four choices "up, same, down and unknown."

Table D.1: Survey Sampling

The table below shows the sampling procedure for non-financial and financial corporations. Financial corporations include both banks, insurance companies, and other financial institutions. * means that 60% of overall small corporations selected by quarterly surveys of Financial Statements Statistics for Corporations by Industry must be sampled and the target number in the sample is around 6,000 firms. In addition, when the number of total sampled corporations with capital less than 500 million yen is less than 30 for each stratum, more firms will be additionally sampled to increase the total number in that stratum to 30.

Panel A: Non-Financial Corporations			
Corporation Type	Size (Yen)	Approx. Size (USD)	Sample Probability
Large	$\geq$ 2 billion	$\geq$ 18 million	100%
	1 - 2 billion	9-18 million	50%
Medium-Sized	0.5 - 1 billion	4.5 - 9 million	50%
	0.1 - 0.5 billion	1 - 4.5 million	Remaining
Small-Medium	0.01 - 0.1 billion	0.1 - 1 million	to hit 6,000 firms*
Panel B: Financial Corporations			
Corporation Type	Size (Yen)	Approx. Size (USD)	Sample Probability
Large	$\geq$1 billion	9 million	100%
Medium-Sized	0.5 - 1 billion	4.5 - 9 million	50%
	0.1 - 0.5 billion	1 - 4.5 million	Remaining
Small-Medium	0.01 - 0.1 billion	0.1 - 1 million	0%

Quantitative questions are about realized and expected values of sales, profits, and investment spending. A key feature of the survey is that full fiscal year values are always reported at all horizons. Questions about intra-year forecasts of quantitative items are different according to quarter. Regarding sales and current profit, the surveys in the first quarter (April - June) and second quarter (July - September) asks the first and second half-year forecasts for the current fiscal year. The survey in the third quarter (October - December) has realizations in the first half-year and forecasts for the second half-year for the current fiscal year. The survey in the fourth quarter (January - March) has realizations in the first half-year and forecasts for the second half-year for the current fiscal year as well as the first and second half-year forecasts in the next fiscal year. All surveys ask the semi-annual realizations in the previous fiscal year.

Table D.2: Response Rates

The table below shows the average annual response rate of the BOS, which was administered from FY 2005 (April 2005 to March 2006) to FY 2016 (April 2016 to March 2017).

Fiscal Year	Response Rate (%) (All Firms)	Response Rate (%) (Large Firms)	Fiscal Year	Response Rate (%) (All Firms)	Response Rate (%) (Large Firms)
2005	78.9	88.8	2011	78.6	87
2006	78.7	88.1	2012	79.2	88.1
2007	78.6	86.8	2013	80.2	88.4
2008	79.2	87.3	2014	81.0	88.1
2009	79.4	87.2	2015	80.9	88.6
2010	79.0	87.1	2016	80.9	88.1