A Quantitative Analysis of Distortions in Managerial Forecasts *

Yueran Ma[†] Tiziano Ropele[‡]
David Sraer[§] David Thesmar[¶]

April 18, 2019

Abstract

This paper quantifies the economic implications of systematic forecast errors made by firm managers. Using administrative survey data from Italy, we show that managerial forecast errors on 1-year ahead sales are positively and significantly autocorrelated. This persistence in forecast error is consistent with managerial under-reaction to new information. To investigate the micro- and macro-economic effects of this forecasting bias, we develop a dynamic equilibrium model with heterogeneous firms and distorted expectations. We estimate the model using firm-level production and forecast data. The model matches exactly the significant under-reaction observed in managerial forecast data, as well as other moments related to investment and production. Compared to an equally imperfectly informed, but rational firm, distorted forecasts lead, in our baseline model, to an average profit loss of about 1.785 % at the firm-level and an aggregate TFP loss of 0.325 %. We investigate how additional distortions affect these estimates.

*We thank seminar participants at Princeton University, UC Berkeley, Boston College, Brandeis, Chicago, Dartmouth, MIT and SITE for their insightful comments.

[†]University of Chicago

[‡]Bank of Italy

[§]UC Berkeley, NBER and CEPR

[¶]MIT and CEPR

1 Introduction

The behavioral corporate finance literature has convincingly demonstrated the existence of systematic biases in managerial decision-making. In a large sample of entrepreneurs, optimism correlates with excessive short-term leverage (Landier and Thesmar, 2009). CEO overconfidence tends to correlate with expensive acquisitions (Malmendier and Tate, 2008) and excess investment (Malmendier and Tate, 2005), particularly in R&D (Hirshleifer et al., 2012). A nascent, related, literature emphasizes the importance of systematic errors in managerial forecasting. Ben-David et al. (2013) provide evidence that CEOs have mis-calibrated expectations of returns, and that such overconfidence correlates with investment and leverage in the cross-section of firms. Gennaioli et al. (2016) show that managerial expectations tend to be extrapolative, and that expectations drive investment across firms. While these papers establish a statistically significant relationship between systematic forecast errors and managerial decisions, our paper asks whether such forecasting biases matter quantitatively. Do they create significant distortions in firm-level decisions? Do they contribute significantly to aggregate inefficiency?

We start by documenting systematic forecasting biases by firms' managers. We use an approach similar in spirit to recent contributions in the macroeconomics and finance literature that investigate the dynamics of beliefs using survey data (see e.g. Coibion and Gorodnichenko (2015), Bordalo et al. (2017a), Bordalo et al. (2017b), Malmendier and Nagel (2016), Bouchaud et al. (2018), among others). More precisely, we exploit a large representative sample of Italian firms surveyed by the bank of Italy (INVIND survey). This sample is an unbalanced panel of about 4,000 firms per year over 2002-2017. It contains firm-level forecast on next year sales that we match with administrative data on sales realizations. We show that managerial forecasting errors are highly persistent, with an auto-correlation coefficient that ranges from .17 to .32. This persistence in forecast errors holds even when controlling for firm fixed-effects and is stable across size quintiles. This result is consistent with the hypothesis that firm managers under-react to recent news about their own firm's

output.

We then develop an economic framework to quantify the effect of forecasting biases. We start from a standard neoclassical model of investment with heterogeneous firms. Productivity follows an AR(1) process. Every period, managers observe the realization of productivity and a private signal informative about next-period productivity. Firms face a one period time-to-build for capital investment, so that managerial forecast about future TFP determines current capital expenditures. Managers may have non-rational expectations about next period TFP. We extend the formulation of belief formation in Bordalo et al. (2017b) to accomodate both over- and under-reaction to news about firm's future TFP. We also allow managers to exhibit different biases in the way they process public and private information. This baseline model has no additional friction beyond the one-period time-to-build and distorted forecasts. We obtain simple closed-form solutions that relate firms' input choices and managerial forecast errors.

We estimate this baseline model by combining data on forecasts and production. On the production side, the model's key parameters are the persistence and productivity of TFP shocks. Using both our U.S. and Italian data, we fit an AR(1) process on TFP residuals. We use our forecast data to identify the distortions in beliefs and the amount of private information held by managers – intuitively, private information is necessary to fill the gap between large TFP shocks and relatively smaller forecast errors. We target the persistence in forecast errors, the dispersion in forecast errors, TFP dynamics, and the covariance between TFP innovations and forecast errors. In line with our reduced-form evidence, we estimate significant distortions in forecasting relative to rational expectations. Managers under-react to TFP shocks (both private information and actual surprise). We then use our model to measure the real effects of expectation distortions on firm behavior. Given that the Italian data is more representative of the economy, and less likely to be polluted by financial communication consideration, we use them as the baseline. The direct effect of expectations distortions on capital investment is large. In our estimated model, following a 1 s.d. TFP

shock, fully rational firms increase their capital stock by 18 ppt more than distorted firms. Across all possible shocks, average corporate profit would be only 1.8% larger (with a s.e. of .12%) in a counterfactual where managers hold rational expectations. This sizable effect on profits happens even though our model has no friction – except for the time to build in capital. So the envelope theorem applies, making effects on profits an order of magnitude smaller than effects on investment.

We then experiment with additional (real) frictions that may affect our results: capital adjustments, long-term time-to-build and labor frictions. We find that quadratic adjustment costs tend to make effects smaller, as they make firms react less to forecasting errors. Longer-term time to build also reduces the effect of expectations, as forecasters in our model make relatively fewer systematic errors in the long-run (there is no bias in the very very long-run). Finally, we account for labor frictions by assuming time-to-build in labor, so that firm capital and labor decisions both depend on productivity forecasts. This increases the scope for forecasting mistakes, as distorted forecasters cannot make up for investment errors by adjusting employment ex-post. So in our exploration, capital frictions tend to make our results weaker, while labor frictions make them larger.

We finally consider how distorted beliefs affect macroeconomic outcomes. We show that aggregate efficiency (aggregate TFP) is proportional to the dispersion in log-sales forecast errors, a result reminiscent of Hsieh and Klenow (2009) and David et al. (2016). Intuitively, dispersion in forecast errors implies capital misallocation: firms with positive (resp. negative) TFP shocks relative to expectations end up with too little (resp. too much) capital ex post. Quantitatively, aggregate TFP losses from distorted beliefs are proportional to the difference between the dispersion in log-sales forecast errors in the data and the dispersion in log-sales forecast errors under fully rational expectations. We use our estimated structural model of forecasting biases to compute the dispersion in log-sales forecast errors under fully rational expectations. We estimate aggregate TFP losses from distorted beliefs to be about .36%. This is a sizable effect. For instance, Baqaee and Farhi (2017) estimate that removing all

heterogeneous mark-ups would increase U.S. TFP by about .76%. Catherine et al. (2017) estimate that removing all financing constraints would raise U.S. TFP by 2%. Our exercise is less extreme in the sense that our benchmark is more "policy-relevant" since it allows for rational errors. In other words, it would not be possible to remove all financing constraints in the real world, while it would presumably be more possible to teach managers to form undistorted forecasts.

Our paper builds on a recent literature that uses forecast data to test rationality. Coibion and Gorodnichenko (2015) document that past revisions positively predict forecast errors in macroeconomic forecasts. They argue that this predictability arises from informational frictions. Bouchaud et al. (2018) document under-reaction among security analysts. Also looking at analysts, Bordalo et al. (2017b) show that forecast errors on long-term EPS growth forecasts are positively correlated with past growth, suggesting over-reaction. Bordalo et al. (2018) document over-reaction in macro and financial variables among professional forecasters. Bloom et al. (2017) show that 15% of plant managers cannot form and express subjective probability distributions. Using data on household expectations of inflation, Malmendier and Nagel (2016) find evidence consistent with "experience effects": heavy discounting of pre-birth data combined with recency bias. We contribute to this literature by documenting significant under-reaction to new information in guidance data. More importantly, our paper provides a tractable framework to quantify the economic effects of this forecasting bias. We incorporate non-rational forecasts into an otherwise standard neo-classical model of investment with heterogeneous firms.

Through its aggregation approach, our paper is also related to a small number of papers that investigate the impact of managerial information on long-term output in steady state models. On the theory side, Akerlov and Yellen (1985) show that, in most equilibrium settings, near-rational behavior can have first-order aggregate consequences, even when it has second-order individual effects. Hassan and Mertens (2017) builds on Akerlov and Yellen (1985) and show that near-rational errors lead to first-order distortions in household savings

decisions. On the empirical side, David et al. (2016) develop a steady-state production model similar to ours, but use it to quantify aggregate efficiency improvements that results from a well-developed stock market. Our paper is concerned with inefficiencies arising from non-rational expectations. In parallel work, Barrero (2018) has developped an approach very similar to ours, using expectation survey data to calibrate the economic effect of managerial forecasts biases that he identifies in the data.

The rest of the paper is organized as follow. Section 2 presents reduced-form evidence of persistence in managerial forecast errors. Section 3 builds a production framework with heterogeneous firms and distorted expectations, provides a structural estimation of the model and quantify a number of partial equilibrium counterfactuals. Section 4 provides aggregation results. Section 5 concludes.

2 Evidence on Managerial Biases

2.1 Data and Summary Statistics

Throughout the paper we are using two data sources. The first one comes from a large business survey conducted every year by the Bank of Italy on a representative sample of firms ("Italian Data"). The second dataset consists of guidances given by large publicly held U.S. firms to financial markets ("U.S. Data").

Italian Data: Our Italian data comes from the Survey on Industrial and Service Firms (INVIND, henceforth) conducted every year by the Bank of Italy since 1972 on a representative sample of firms. Since 2002, the reference universe in INVIND consists of firms with at least 20 employees operating in industrial sectors (manufacturing, energy and extractive industries) and in non-financial private services and with administrative headquarters in Italy.[1] In recent years each wave has around 4,000 firms (3,000 industrial firms and 1,000 service firms). Data referring to a specific year are collected by the Bank of Italy's local

[1]For further details see the methodological note about the Survey on industrial and services firms here.

branches between February and April of the following year. Among other things, the survey questionnaire asks firms to report their sales, investments and employment in three periods: the fiscal year just ended (preliminary results), the previous fiscal year (final results) and the current fiscal year (forecasts). To compute the firm-specific total sale forecast error we do not use INVIND to measure actual sale (i.e. final result for previous fiscal year) as this information might be subject to self-reporting bias. Rather we measure the firm-specific actual sale using the Company Accounts Data System (CADS), which managed by the Cerved Group and includes balance-sheet data for all Italian limited liability companies.

The Italian survey-account sample runs from 2002 through 2017 and contains about 37,000 firm-year observations.[2]

U.S. Data: As complementary evidence, we use data on forecasts given by large publicly listed U.S. firms to financial markets. These data come from the IBES Guidance dataset, a feed that offers quantitative (numeric) company expectations from press releases and transcripts of corporate events. We focus on forecasts of fiscal year total sales, which is one of the most common forecast. The forecast horizon ranges 1 quarter out (current quarter corresponds to the fiscal year end quarter) to about 30 quarters out; forecast horizon longer than 12 quarters is rare. For consistency, we restrict the sample to forecasts made in the first fiscal quarter about total sales of that fiscal year. Comprehensive data coverage starts from 2003 onward, and the extract we use in this paper stops in 2015. Since we are interested in the persistence of forecast errors, we also restrict the sample to firms that issue at least 5 guidance over the sample period.

We focus on US non-financial firms (SIC outside of 6000 to 6999). The final sample thus runs from 2003 to 2015 and contains 5,222 firm-year observations, corresponding to about 440 firms per year. All variables are winsorized at the median +/- 5 times the interquartile range.[3]

[2] As for the guidance sample, we keep only firms that report at least 5 forecasts for total sales and winzorize all variables at the median +/- 5 times the interquartile range.

[3] If we only winsorize at the 1% level, some observations still have log-sales forecast errors greater than 50%

We view U.S. evidence as complementary to evidence from our Italian data, but we focus our structural analysis on Italian data. Italian data have three advantages: (1) the sample of firms surveyed by the Bank of Italy is more representative, (2) most Italian firms are likely to be owner-managed, so less subject to agency problems, and (3) firm-level survey results are not disclosed by the Bank of Italy so are unlikly to be polluted by signalling concerns than public guidances to financial markets (Degeorge et al., 1999). U.S. offer though a useful complement since firms stake their reputations on these announcement, probably much more than in survey responses. Also, they offer evidence from a different country, and firms so large that they are likely to matter disproportionately for aggregate outcomes.

Table 1 Panel A provides summary statistics for the Italian sample while Panel B focuses on the U.S. sample. Panels A1 and B1 focus on firms for which forecast data are available. Panel A2 and B2 provide descriptive statistics for the corresponding universe in order to make a comparison. In both countries, the average firm in the forecasting sample is larger and more profitable than the average firm in the corresponding universe.

Throughout the paper, we define empirical log-sales forecast errors as the difference between sales and sales forecast in the previous year:

$$\widehat{FE}_{it} = \log sales_{it} - \log F_{t-1} sales_{it}$$

where we define $F_{t-1} sales_{it}$ as the forecast made by firm i as of $t-1$ about sales of year t. In both datasets, the forecast F_{t-1} is actually issued at the beginning of year t. In Italian data, it is the answer to the survey collected between February and April of year t. In the U.S. data, we take managerial guidances (or analyst forecast consensus) issued in the first quarter of t in order to maximize the number of observations.

In the U.S., managerial log-sales forecast errors (guidances) have a slightly smaller s.d.

in absolute value. These observations mostly correspond to firms experiencing a significant corporate event, such as a merger. These large forecast errors should therefore not be attributed to managerial forecasting bias.

(9%) than analyst forecast errors (11%). Italian managerial forecast errors are much bigger than U.S. guidance errors, with an s.d. of 18% (compared to 11%). This may be due to the fact that Italian firms are smaller, and therefore more volatile than U.S. firms. Figure 2 provides a histogram of log-sales forecast error for both countries and confirms that Italian forecast errors are more dispersed.

2.2 Forecast Informativeness

We now examine the informativeness of managerial forecasts in our data. A concern is that these forecasts are only wekly correlated with actual forecast, for instance because managers are trying to anchor market expectations (for U.S. data) or because they put little information in answering the survey run by the Bank of Italy (for Italian data). We show here that (1) in both datasets forecasts are highly predictive of future sales, (2) they explain investment decisions very well and (3) U.S. guidance data are consistent with firm forecasts collected by the Duke University CFO survey, which is anonymous (this is to address the concern that U.S. data may be polluted by financial market communication).

Predicting Actuals

Are sales forecasts a reliable predictor of realized sales? Table 2 predicts actual sales in fiscal year t using sales forecasts for year t made in the first quarter and additional control variables. We estimate the following regression:

$$\frac{sales_{it}}{assets_{i,t-1}} = \alpha_i + \delta_t + \frac{F_{i,t-1}sales_{it}}{assets_{i,t-1}} + X_{it}'\gamma + \epsilon_{it},$$

where $F_{i,t-1}sales_{it}$ corresponds to the guidance for firm i and fiscal year t made in the first quarter of that year. The set of control variables X includes: contemporaneous (consensus) analyst sales forecasts, lagged sales, beginning-of-year Q (for U.S. data only), beginning-of-year log assets, year fixed effects, industry fixed effects (Fama-French 12 industries for U.S.

data, 2-digit industry classification for Italian data), industry-year fixed effects, and firm fixed effects. Standard errors are clustered by firm and time.

Table 2 has two panels: Panel A reports regression results for Italian data, while Panel B reports regression results for U.S. data. In Italian data, the coefficient on forecast is .97 with an R^2 of 90%. In the U.S. data, where firms are larger and less volatile, the coefficient on forecast is 1.005 and the regression's R^2 is 97%. The various specifications estimated in Table 2 confirm the robustness of this finding to including additional predictors. In particular, U.S. data show that managerial guidances are much more informative about future sales than analyst forecasts.

Expectations and Investment

We show here that sales expectations explain investment decisions over and above traditional determinants. Table 3 regresses capital expenditures in fiscal year t on sales forecasts for year t made in the first quarter and additional control variables. We estimate the following regression:

$$\frac{capx_{it}}{assets_{i,t-1}} = \alpha_i + \delta_t + \frac{F_{i,t-1}sales_{it}}{assets_{i,t-1}} + X_{it}'\gamma + \epsilon_{it},$$

where $F_{i,t-1}sales_{it}$ corresponds to the sales forecast for firm i and fiscal year t made in the first quarter of that year. The set of control variables X includes: contemporaneous (consensus) analyst sales forecasts, lagged capex over asset, beginning-of-year Q (for U.S. data only), beginning-of-year log assets, cash flows (income before extraordinary items plus depreciation and amortization, normalized by lagged assets), lagged cash flows year, fixed effects, industry fixed effects (Fama-French 12 industries in U.S. data versus 2-digit industry dummies in Italian data), industry-year fixed effects, and firm fixed effects.

Table 3 shows that sales forecast explain investment decisions above and beyond its traditional determinants. This result is consistent with (Gennaioli et al., 2016), who document that managers earnings expectations have strong explanatory power for investment activities.

This relationship is robust to controlling for standard determinants of capital expenditures used in the literature, such as Q or cash flows. Obviously, a threat to interpretation is reverse causality: when firms are investing a lot, managers expect sales to increase. Note that capital expenditures capture long-term capital investments, which are unlikely to pay off immediately. Quantitatively, if we switch in U.S. data the left-hand side and right-hand side variables (i.e. use capital expenditures to explain sales forecast), the coefficient becomes greater than 2.5, implying that if the reverse causality explanation was the entire story, $1 of capital expenditure immediately turns into $2.5 of sales. This would be unreasonably large. So while we cannot rule out reverse causality, it is unlikely to explain 100% of our findings.

Cross-Check with Duke CFO Survey Data

Last, we cross-check sales forecasts measured using guidance data and forecasts collected in the Duke CFO Survey led by John Graham and Campbell Harvey. The survey is anonymous and takes place on a quarterly basis; respondents come from around 400 major firms in the US.[4] Each quarter, respondents report expectations of the future 12 month growth of sales and other key corporate variables. The aggregate statistics are published on the CFO Survey's website.

Figure 1 presents the time series of aggregate (sales-weighted) expected sales growth using IBES guidance data and Duke CFO Survey data. Note that the forecast horizon for the CFO Survey is rolling 12 months, while the guidance data is every fiscal year. Thus every quarter we use guidance data with forecast horizon ranging from 2 to 6 quarters as an approximate. Figure 1 shows managerial forecasts are consistent across these two sources. The guidance data is not systematically biased upward or downward. The guidance aggregate also appears to be slightly more volatile than the CFO Survey aggregate.

[4]Other segments of the CFO Survey also have firms from other countries. Here we focus on US firms for both CFO Survey data and guidance data.

2.3 Persistence of Forecast Errors

Under rational expectations, managerial forecast errors should be i.i.d. We document instead that managerial forecasts errors are persistent in both our samples. This finding is consistent with under-reaction to new information. We estimate the following model:

$$\widehat{FE}_{it} = \alpha + \delta_t + \beta \widehat{FE}_{i,t-1} + \epsilon_{it}, \tag{1}$$

where $\widehat{FE}_{it}$ is log-sales forecast error defined above and δ_t correspond to year fixed-effects. Figure 3 provides a binned scatter plot of this relationship between past log-sales forecast errors and current log-sales forecast errors. Panel A shows Italian results and panel B shows U.S. results. In both cases, it restricts the sample to firms with at least 5 forecasts. To construct this figure, we split the sample in vingtiles of lagged log-sales forecast error (x-axis) and represent, on the y-axis, the average current log-sales forecast error: the relationship between lagged and current log-sales forecast error is increasing and close to linear. It appears to be about two times steeper in Italian data. This is consistent with the idea that underreaction is more pronounced in the data. This is also consistent with the possibility that U.S. managerial guidances are "noisier" in the sense that they face a bigger measurement error (possibly coming from strategic considerations of publicly listed firms in releasing their forecasts). If that were the case, U.S. data would underestimate underreaction as past errors are noisily estimated. To partly account for this, we will assume in our estimation procedure that such a measurement error exist and will identify it using forecast error variance.

Table 4 reports the equivalent regression results. Again, Panel A presents Italian results and Panel B. U.S. ones. Column (1) estimates Equation 1 on the baseline sample using OLS. Standard errors are clustered at the firm and year level. The estimated β is .32 in Italy and .18 in the U.S. Both are statistically significant at the 1% confidence level. Column (2) repeats this analysis, using analysts consensus forecasts to compute log-sales forecast errors. We can only perform this analysis for U.S. firms as Italian firms are non-listed and

therefore have no analyst coverage. We include all firms with at least 5 analysts consensus forecasts, and since analyst coverage is extensive in the U.S., the number of observation goes from 4k to 18k. The estimated β remains the same as with managerial guidances at .17 (significant at 1%). Columns (3) and (4) augment the models of columns (1) and (2) by including firm fixed-effects. Such fixed-effects can arise, for example, for strategic reasons or goal-setting purposes. They are a priori more likely to arise in U.S. guidance data, which are public and part of the firm's financial communication, than in Italian data. This augmented model cannot be estimated consistently using OLS given the short time period in our sample (Nickell (1981)). As a result, we further restrict the sample to firms with at least 9 forecasts and estimate the model using dynamic panel GMM (Arellano and Bover (1995)). Using both guidance and analysts consensus forecasts, these augmented models lead to an estimate of .17 for Italy and .12 for U.S. , both significant at the 1% confidence level.

These estimates are consistent with the idea that managers underreact to news: Good news, when they arise, are slowly incorporated in forecasts, leading to forecast error persistence. How important are these errors? The quantitative interpretation of these persistence estimates does require a structural model, which we describe in Section 3. But before doing this, we provide one more piece of evidence consistent with underreaction in managerial expectations.

2.4 Autocorrelation by Size Group

We ask here how forecast error persistence varies as a function of size. One possible concern is that stickiness is mostly a small firm phenomenon – since U.S. firms are larger than Italian firms in our sample. To check this, it is most sensible to use the Italian sample, which is more representative of the population of firms. We split the Italian sample into 5 size groups, and re-run regression (1) for each of these 5 subgroups separately. We report results in Appendix Table B.1. Across size groups, estimates are quantitatively consistent and strongly significant. We do not find evidence that forecast persistent decreases with size.

If anything, the opposite happens. But given the confidence intervals, we cannot reject the null hypothesis that the autoregression coefficient is the same across size groups.

2.5 Link between Error and Revisions

We also document the link between forecast errors and forecast revisions, drawing on the specification in (Coibion and Gorodnichenko, 2015). When a forecaster receive positive news, she revise her forecast upward. When she under-reacts to the information, however, the revision fail to fully account for the news content. The new forecast therefore leads to a positive forecast error, creating a positive correlation between forecast revisions and errors.

We focus here on U.S. data as the Italian survey does not collect the term structure of expectation. Even in the U.S., one limitation of the guidance data is the forecast horizon is rarely longer than six quarters and the forecasts are also not necessarily regular. Thus it is difficult to run regressions that utilize forecast error of year t sales made at the beginning of t, together with forecast revision of year t sales from the beginning of $t-1$ to the beginning of year t. To utilize more data, we mix forecast horizons and revision duration: we use all available forecasts, together with the revision since the previous forecast of the same outcome (we require the forecast horizon to be at least 2 quarters). Again we perform this test using both firm guidance data and contemporaneous analyst forecast data.

We estimate the following model:

$$\log \text{sales}_{it} - \log F_{t-h}\text{sales}_{it} = cst + \gamma \left(\log F_{t-h} \log \text{sales}_{it} - \log F_{t-k}\text{sales}_{it}\right) + \epsilon_{iht}$$

where we lump together different horizons of revisions h (as long as h is no less than 2 quarters), and k is the time when the previous forecast of sales_{it} is made. Table B.2 in Appendix reports the results. We estimate a significant γ of around 0.33 in both guidance and analyst forecast data. Without further assumption on the sales process, this coefficient cannot directly be compared to the persistence estimate we discussed above. But both

estimates go in the same direction: Manager forecasts seem to behave in a way consistant with underreaction to news. Such bias may arise because of inattention, or because of informational or institutional rigidities.

In the rest of the paper, we investigate the quantitative consequences of such biases for production and efficiency.

3 Partial Equilibrium Model of Investment with Distorted Beliefs

3.1 Baseline Model

We start from a standard neoclassical model of investment with two frictions: (1) 1 period time to build (2) distorted beliefs.

Economy

Time is discrete. At date t, firm i combines capital k_{it} and labor l_{it} to generate sales with a Cobb-Douglas technology:

$$p_{it}y_{it} = Ae^{\nu_{it}}\left(k_{it}^{\alpha}l_{it}^{1-\alpha}\right)^{\theta},$$

where ν_{it} is revenue-based log-productivity, α is the capital share and θ captures decreasing returns to scale in revenues, which may arise from technology or market power. Input markets are competitive. w is the wage on the labor market and R is the rental rate of capital.[5] At date t, firms hire l_{it} employees after observing ν_{it}. However, we assume a one-period time-to-build in capital: firms invest in the capital stock k_{it} before ν_{it} is realized. As a result, managers need to form expectations about next-period productivity before investing. We assume an AR(1) process for ν_{it}:

[5]We omit the time subscript for w and R, as we only consider steady-state economies.

$$\nu_{it} = \rho\nu_{it-1} + \psi_{it} + \omega_{it} \quad \text{with:} \quad (\omega_{it}, \psi_{it}) \sim \mathcal{N}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\omega^2 & 0 \\ 0 & \sigma_\psi^2 \end{pmatrix}\right], \tag{2}$$

where ω_{it} is realized at date t, after k_{it} has been purchased; in contrast, ψ_{it} is observed by the manager at date $t-1$, but not by the econometrician. Our model thus allows for managerial private information about future productivity.

Distorted Expectations

Managers may exhibit distorted expectations about future productivity. In the spirit of Bordalo et al. (2017a), we assume managers use the following conditional productivity:

$$h^s_{t-1}(\nu_{it+T}) = h(\nu_{it+T}|\nu_{it-1}, \psi_{it}) \left[\frac{h(\nu_{it+T}|\nu_{it-1}, \psi_{it})}{h(\nu_{it+T}|\widehat{\nu_{it-1}} = \rho\nu_{it-2} + \psi_{it-1}, \psi_{it})}\right]^\gamma \left[\frac{h(\nu_{it+T}|\nu_{it-1}, \psi_{it})}{h(\nu_{it+T}|\nu_{it-1}, \widehat{\psi_{it}} = 0)}\right]^\lambda \frac{1}{Z} \tag{3}$$

where Z is a normalization coefficient and $T \geq 0$ is the forecast horizon. This subjective density allows for under- or overreaction to private information and true innovation. The first term is the true (rational) conditional distribution. The second term is the over/under-reaction to the surprise. The "diagnostic" is large when the past realization of productivity is high compared to expectation. Those states are overweighted when $\gamma > 0$ (representativeness bias in Bordalo et al. (2017a)). They are underweighted when $\gamma < 0$. This story can be rationalized as lack of attention: Large surprises are less representative of potential realization and therefore more likely to be overlooked. The third term follows the same logic but focuses on over/underreaction to private information. The diagnostic ratio is high when the private information is large compared to its expected value $\widehat{\psi} = 0$.

Given the above distribution, it is straightforward (see for instance the proof in Bordalo

et al. (2017a)) to show that the above subjective distribution has conditional expectation:

$$\begin{aligned} F_{t-1}\nu_{it+T} = & E\left(\nu_{it+T}|\nu_{it-1},\psi_{it}\right) \\ & + \gamma\left(E\left(\nu_{it+T}|\nu_{it-1},\psi_{it}\right) - E\left(\nu_{it+T}|\rho\nu_{it-2}+\psi_{it-1},\psi_{it}\right)\right) \\ & + \lambda\left(E\left(\nu_{it+T}|\nu_{it-1},\psi_{it}\right) - E\left(\nu_{it+T}|\nu_{it-1},0\right)\right) \end{aligned}$$

and has conditional variance $var\left(\nu_{it+T}|\nu_{it-1},\psi_{it}\right)$. So this formulation only introduces a bias in expected TFP, not in the other moments.

We focus here on short horizon forecasts (we return to longer horizon forecast in Section 3.5.2), as these are the ones we see in the data. So we set $T = 0$. We obtain the following expression for the conditional subjective expected TFP:

$$F_{t-1}\nu_{it} = \left(\rho\nu_{it-1} + \psi_{it} + \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2\sigma_\omega^2\right) + \underbrace{\gamma\rho\omega_{it-1} + \lambda\psi_{it}}_{\text{distorted expectations}}$$

This formulation entails two deviations from rational expectations: (1) when $\gamma > 0$ (resp. < 0), managers are over-reacting (resp. under-reacting) to the date t-1 innovation in productivity ω_{it-1} (2) when $\lambda > 0$ (resp. < 0), managers are over-reacting (resp. under-reacting) to their date t-1 private information about date-t productivity ψ_{it}. When $\gamma = \lambda > 0$, our model has the same properties as Bordalo et al. (2017a): the forecaster overweights "exceptional" past realizations when forming beliefs (i.e. date t-1 realizations that deviate from date t-2 forecasts). But our formulation extends the framework of Bordalo et al. (2017a) in two directions: (1) We allow errors on private and public information to differ (γ and λ may differ) and (2) the agent may underweight as well as overweight such realizations (γ and λ may be negative).

Finally, we allow reported managerial forecasts about future sales, $\widehat{\mathbb{F}_{i,t-1}[p_{it}y_{it}]}$, to differ from true managerial forecasts $\mathbb{F}_{i,t-1}[p_{it}y_{it}]$. Specifically:

$$\ln\left(\widehat{\mathbb{F}_{i,t-1}[p_{it}y_{it}]}\right) = \ln\left(\mathbb{F}_{i,t-1}[p_{it}y_{it}]\right) + \zeta_{it}, \quad \text{where:} \quad \zeta_{it} \sim \mathcal{N}\left(0,\sigma_\zeta^2\right) \tag{4}$$

This is to account for the fact that managers may not reveal their true expectations, for strategic considerations (as in U.S. guidance data) or for internal of external goal-setting purposes.

We summarize firm behavior in the following proposition:

Proposition 1. Let $\Phi = \frac{\alpha\theta}{1-(1-\alpha)\theta}$. Firm i's optimal capital stock at date t is:

$$k_{it} = \tilde{\Omega}_1 \left(\mathbb{F}_{it-1} \left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}} \right] \right)^{\frac{1}{1-\Phi}} = \tilde{\Omega}_1 e^{\frac{1}{1-\Phi}\left(\frac{\Phi}{\alpha\theta}(\rho\nu_{it-1} + \gamma\rho\omega_{it-1} + (1+\lambda)\psi_{it}) + \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2 \sigma_\omega^2 \right)},$$

where $\tilde{\Omega}_1$ is constant across firms. Firm i's sales at date t are simply given by:

$$p_{it}y_{it} = \tilde{\Omega}_2 e^{\frac{\Phi}{\alpha\theta}\nu_{it}} k_{it}^{\Phi}$$

where $\tilde{\Omega}_2$ is constant across firms. Log-sales forecast error at date t using reported forecasts at date t-1 are:

$$\widehat{FE}_{it} = \ln\left(\frac{p_{it}y_{it}}{\widehat{\mathbb{F}_{i,t-1}[p_{it}y_{it}]}} \right) = -\frac{\Phi}{\alpha\theta}\left(\gamma\rho\omega_{it-1} + \lambda\psi_{it}\right) + \frac{\Phi}{\alpha\theta}\omega_{it} - \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2 \sigma_\omega^2 - \zeta_{it} \qquad (5)$$

Proof. See Appendix A.1. □

Firm i's investment decision at date t-1 depends on firm i's forecast of date-t productivity. This forecast is distorted: with $\gamma \neq 0$ and $\lambda \neq 0$, managers put non-rational weights on date t-1 innovations in productivity (ω_{it-1} and ψ_{it}). These distorted forecasts lead to predictable forecast errors that depend systematically on these past innovations.

3.2 Model Estimation Procedure

We calibrate two parameters related to production: $\alpha = .33$ and $\theta = .8$ ((Broda and Weinstein, 2006)). We estimate 6 parameters $(\rho, \sigma_\omega, \sigma_\psi, \gamma, \lambda, \sigma_\zeta)$. We follow David et al. (2016) and construct revenue-based productivity for firm i at date t as: $\hat{v}_{it} = \frac{\alpha\theta}{\Phi}\left(\ln(p_{it}y_{it}) - \Phi\ln(k_{it})\right)$.

We use the net value of property, plant and equipment in year $t-1$ as our measure of k_{it}[6] and $p_{it}y_{it}$ is firm i's total sales for fiscal year t. We use dynamic panel GMM to estimate the following process for $\hat{v}_{it}$ (Arellano and Bover (1995)):

$$\hat{v_{it}} = \delta_i + \delta_t + \chi \hat{v_{it}} + \tau_{it}$$

The estimated persistence, $\hat{\chi}$ is a consistent of ρ provided the production function and demand system are true. This is a strong assumption but one we are making throughout this paper. We estimate the remaining structural parameters through a moment estimator that targets five sample moments from both production and forecast data. Since the model does not allow for aggregate shocks, all the moments used in the estimation control for year fixed-effects. The five moments we target are:

1. the variance of the estimated productivity innovations $\widehat{\sigma}^2_\tau$. In the model,

$$\sigma^2_\tau = \sigma^2_\omega + \sigma^2_\psi$$

   the variance of TFP innovation is the sum of the part that is anticipated by managers (via private information) and the true surprise.

2. The variance of residuals from a regression of log-sales forecast errors on year fixed-effects ($\widehat{\mathbb{V}\text{ar}}[\widehat{FE}_{it}]$). In the model,

$$\mathbb{V}\text{ar}[\widehat{FE}_{it}] = \sigma^2_\zeta + \left(\frac{\Phi}{\alpha\theta}\right)^2 \left((1+\gamma^2\rho^2)\sigma^2_\omega + \lambda^2\sigma^2_\psi\right)$$

   The variance of forecast errors contains: forecast measurement error, the effect on sales of the true productivity surprise, and the effect on sales of errors coming from expectations distortions.

[6] In the model, k_{it} is determined at date $t-1$ but can only be used for production in period t. PPE observed in year $t-1$ include to the capital expenditures made in year t and thus corresponds to our definition of k_{it} in the model.

3. the estimated coefficient $\hat{\kappa}_1$ from a regression of log-sales forecast error on lagged forecast error. This regression corresponds to results shown in column 1, Table 4. As discussed in Section 2.3, a positive $\hat{\kappa}_1$ implies that managers are under-reacting to TFP innovations ($\gamma < 0$). In the model,

$$\kappa_1 = -\frac{\left(\frac{\Phi}{\alpha\theta}\right)^2 \gamma\rho\sigma_\omega^2}{\mathbb{V}\text{ar}[\widehat{FE}_{it}]}$$

where it appears that the autocorrelation of forecast errors is positive only when $\gamma < 0$, i.e. when managers underreact to past positive news. As expected, forecast measurement error σ_ζ^2 makes the coefficient smaller. Note that λ, overreaction to private information, does not affect the regression coefficient. This is because under/overreaction to private information has no impact on error persistence.[7]

4. the estimated coefficient $\hat{\kappa}_2$ of a regression of date t productivity residual, $\widehat{\tau_{it}}$ on date-t reported log-sales forecast error, controlling for year and firm fixed-effects. Intuitively, under-reaction to private information (i.e. $\lambda < 0$) creates a positive correlation between what the econometrician observes as a positive innovation to TFP and managerial forecast error. In the model:

$$\kappa_2 = \frac{\left(\frac{\Phi}{\alpha\theta}\right)\left(\sigma_\omega^2 - \lambda\sigma_\psi^2\right)}{\mathbb{V}\text{ar}[\widehat{FE}_{it}]}$$

TFP and forecast error comove less when λ is large, i.e. when managers overreact to private information. When private information is good, TFP is high, but since managers overreact, their forecast error is low.

5. the estimated coefficient $\hat{\kappa}_3$ of a regression of the log sales-to-capital ratio, $\ln(p_{it}y_{it}) - \ln(k_{it})$, on date-t reported log-sales forecast error, controlling for year and firm fixed-

[7] This happens in the model because distortions have no memory in this model. Future biases only depend on recent information, not on old one. This is admittedly a strong assumption that we may relax in future research.

effects. In the model:

$$\kappa_3 = 1 - \frac{\sigma_\zeta^2}{\mathbb{V}\text{ar}[\widehat{FE}_{it}]}$$

We know that in the model, realized MRPK is directly equal to forecasting error (large positive error leads to "inefficiently" high MRPk – too little capital). So in the absence of forecast measurement error the regression coefficient should be 1. So measurement error on forecasts σ_ζ^2 makes it lower than 1.

Intuitively identification operates the following way. κ_3 allows to pin down the forecast measurement error σ_ζ. The role of κ_1 – the persistence in log-sales forecast errors – is to pin down γ, the managerial bias on TFP surprise. κ_2 is here to estimate the excess reaction to managerial private information λ. Then, the variances of surprise σ_ω and private information σ_ψ are jointly measured through $\mathbb{V}\text{ar}[\widehat{FE}_{it}]$ and the TFP innovation variance σ_τ.

3.3 Model Estimation Results

Table 5 reports the 6 moments used in the estimation. As before, Panel A stands to Italian results and Panel B for U.S. results. Standard errors are obtained by bootstrapping on the estimation sample using a block bootstrap at the firm-level. The persistence of log-TFP is estimated to be .72 in the U.S. (.79 in Italy), while the variance of log-TFP innovations is .008 (.072 in Italy). The variance of log-sales forecast errors after projecting on year fixed-effects is .008 (versus .031 in Italy). The two new regression coefficients, $\widehat{\kappa}_2$ and $\widehat{\kappa}_3$, are both positive and significant at the 1% level in both countries.

Table 5 also reports the estimated structural parameters of our model. We estimate γ – the parameter governing the distortion in expectations related to public information – at -.32 in the U.S. (-.78 in Italy): In both countries, managers put negative weights on recent innovations to the public component of TFP innovations, ω_{it}. We also estimate λ to be slightly negative: there is some limited over-reaction to private information, but it does not need to be large to account for the empirical comovement between forecast error and

productivity.

Two additional interesting patterns emerge. First, the model estimates sizable forecast measurement error. This is because MRPK variations can only be imperfectly attributed to changes in forecasting error. Given such a big measurement error, forecast errors should be big, but they are small compared to TFP innovations. This can only be reconciled with significant private information, to which managers do not overreact too much. This leads to a large amount of private information, especially in Italy.

3.4 Partial equilibrium counterfactual

In this section, we use our structural estimates to quantitatively compare average corporate behavior for firms with rational managers relative to firms with managers that use distorted forecasts. We do this in partial equilibrium, i.e. without clearing product and labor markets.

Investment conditional on TFP shocks

In Figure 4, we first consider the case of investment. From Proposition 1, we find that log investment is given by:

$$\Delta \log k_{it} = \frac{1}{1-\theta}\left(\rho \Delta \nu_{it-1} + \gamma\rho\Delta\omega_{it-1} + (1+\lambda)\Delta\psi_{it}\right)$$

so that the regression coefficient of investment on log TFP innovation is given by:

$$\beta_{\Delta \log k_{it}/\omega_{it-1}} = \frac{\rho}{1-\theta}(1+\gamma)$$

which shows that investment is less sensitive to productivity shocks when the manager underreacts (i.e. when $\gamma < 0$).

We use the estimated parameters in Table 5 to quantify this effect. In Figure 4, we plot log investment against 20 deciles of TFP shocks for managers with rational forecasts (dark square – $\gamma = \lambda = 0$) and managers with distorted forecasts (white diamonds – γ and λ at

their estimated values). We see that the conditional investment policy of a rational manager significantly differs from that of a manager with distorted forecasts. For high (resp. low) realization of ω_{t-1}, the distorted manager significantly under-invest (resp. over-invest). In Italian data, the effect is large: for a +1 s.d. TFP shock, capital growth is about 23% for a rational manager, compared to only 5% for an underreacting manager. In the U.S. the difference is smaller, yet still sizable. Facing a shock of the the same size (thus more than 1 s.d. given that U.S. are less volatile), rational U.S. firms grow their capital stock by 21%, vs 14% if they are distorted. The effect of expectation distortions is about 2.5 times larger in Italy (18ppt vs 7ppt) because the estimated γ coefficient is also about 2.5 times larger there, which is itself directly coming from more error persistence in reduced-form regressions.

Profits conditional on TFP shocks

This significant difference in investment policy leads to a difference in conditional profits which is an order of magnitude smaller. This is the direct result of the envelope theorem. We do not have simple closed forms for profits, but can easily simulate their expected values, and sho them in Figure 5. Within each bucket of past TFP innovations (ω_{it-1}), we compute the difference between log expected profits made by distorted managers and managers with distorted forecasts, conditional on TFP level ν_{it-1}. On average, distorted managers are close to being rational for zero TFP firms – this comes from the fact that the log forecast error is zero on average in our model. As TFP moves away from zero, distorted managers make bigger and bigger mistakes. In Italian data, the difference in profit appears non-trivial: for a negative shock of -.5 in log TFP (slightly more than one s.d.), distorted managers do not scale investment down enough, and there profits are about 3% lower. On U.S. data, the effect is here too, but much smaller. This is because, as we have seen above, the investment distortion is smaller for U.S. firms. The envelope theorem applies and makes profit losses an order of magnitude smaller than investment differences.

Unconditional Profits

We compute here the unconditional gain from having rational expectations. For each realization of productivity, we calculate the difference between (1) the log-profit of a rational manager $\ln(\pi^*_{it})$ and (2) the log-profit of a manager with distorted forecasts $\ln(\pi^F_{it})$. We then compute Δ, the average of this log-difference:[8]

$$\begin{aligned}\Delta = \mathbb{E}\left[\ln(\pi^*_{it}) - \ln(\pi^F_{it})\right] &= \underset{(0.123\%)}{1.785\%} \text{ on Italian data} \\ &= \underset{(0.020\%)}{0.058\%} \text{ on U.S. data}\end{aligned}$$

This calculation confirms that relative profit differences are an order of magnitude smaller than investment differences. In U.S. data, this leads to negligible profit difference. In Italian data, the envelope theorem does not work as forecfully, as investment differences are larger. Then, non-bayesian expectations lead to a 1.75% loss in profits. We now explore if these profit losses translate into misallocation losses. Note that our estimation allows us to give confidence intervals to these estimates. The Italian estimate of profitability loss is both sizable and strongly significant.

3.5 Extensions

3.5.1 Adjustment costs

We now assume that firms face a fixed cost of adjusting their capital stock, as well as quadratic adjustment costs. With our formulation of distorted expectations, we can write the Bellman representation of the firm dynamic optimization problem:

[8] Note that Δ is independent of w, R and A since, in our Cobb-Douglas environments, realized profits are log-linear in the wage w, the user cost of capital R and the average productivity A

$$V(k_{it}, \nu_{it}, \nu_{it-1}, \psi_{it}) = \left| \begin{array}{l} \int_{\tilde{\psi}} \max_{(k_{it+1})} \left\{ \left(\pi\left(\nu_{it}, k_{it}, k_{it+1}\right) + \frac{1}{1+r} \int_{\tilde{\omega}} V(k_{it+1}, \rho(\nu_{it} + \gamma(\nu_{it} - \psi_{it} - \rho\nu_{it-1})) + (1+\lambda)\tilde{\psi} + \tilde{\omega}, \nu_{it}, \tilde{\psi}) G(\tilde{\omega}) d\tilde{\omega} \right) G(\tilde{\psi}) d\tilde{\psi} \right\} \\ \pi(k, k', \nu) = \Omega e^{\frac{\Phi}{\alpha\theta}\nu} k^{\Phi} - (k' - (1-\delta)k) - c_k \frac{(k' - (1-\delta)k)^2}{k} - f_k \mathbb{1}_{k' \neq \delta k} \end{array} \right. , \tag{6}$$

where c_k parameterizes quadratic adjustment costs to capital and f_k fixed costs of investment. The implicit assumption here is that biased managerial expectations differ from rational ones through their mean only. All other moments are rational. This can be rationalized in a framework similar to ?.

To gauge whether adjustment costs have the potential to increase the effect of distorted forecasts on firm-level efficiency, we proceed in the following steps. First, we numerically solve the Bellman problem (6) using U.S. parameters from the estimation of our baseline model (Section 5), and varying values of the adjustment cost paramater c_k (going from 0 (the baseline model) to $c_k = .2$).[9] We then re-solve this Bellman problem keeping the same parameters except for λ and γ, which we both set to 0: this corresponds to our rational expectation benchmark. Third, for each set of parameters, we simulate firm-level data from the solution to the Bellman problem. Fourth, we compute, for each realization in the state space, the log-difference between the profits realized by a firm with rational expectations and a firm with distorted forecasts.

We report the results on Figure C.1: as quadratic adjustment costs increase, we see that the benefits of rational forecasts in terms of firm-level profits decrease. Intuitively, adjustment costs play two roles in our model: (1) they make forecast errors more distortive, since the firm carries the inefficent capital longer (2) they make the firm less sensitive to productivity forecasts, making forecast errors less distortive. Quantitatively, the second effect dominates. Quadratic adjustment costs are therefore unlikely to make biases in forecasts more important from a quantitative standpoint.

[9] Using US plant-level data, Cooper and Haltiwanger (2006) estimate $c_k = .05$ for a neoclassical model of investment with quadratic and fixed adjustment costs.

3.5.2 Longer time-to-build

Can a longer time-to-build for capital investment lead to greater inefficiencies from distorted forecasts? Consider a simple extension to our baseline model where the capital purchased at date t can only be used for production at date $t+T$. The baseline model corresponds to T=1. The data generating process for TFP (ν_{it}) remains similar to our baseline model. The firm's investment decision at date t then becomes:

$$\max_{k_{it+T}} \left\{ \Omega \mathbb{F}_{t-1} \left[e^{\frac{\Phi}{\alpha\theta}\nu_{it+T}} \right] k_{it+T}^{\Phi} - R k_{it+T} \right\}$$

The model of distorted forecasts specified in Equation 4 implies that the T-horizon forecasts is given by:

$$\ln\left(\mathbb{F}_{t-1} \left[e^{\frac{\Phi}{\alpha\theta}\nu_{it+T}} \right] \right) = \frac{\Phi}{\alpha\theta} \left(\rho^{T} (\nu_{it-1} + \gamma\omega_{it-1}) + (1+\lambda)\rho^{T-1}\psi_{it} \right) + \frac{1}{2} \left(\frac{\Phi}{\alpha\theta} \right)^{2} \sigma_{\omega}^{2}$$

Intuitively, as the forecast horizon increases, the forecasting bias is reduced: TFP is mean-reverting and managers understand that irrespective of their current forecast, the long-term mean of the log-productivity process is 0. Therefore, inattention to current levels of productivity does not really matter for long-term forecasts. Long-horizon investments are, if anything, less affected by managerial error than short horizon ones. This comes from the fact that this model embeds no "long-term bias". The diagnostic ratios in the density (3) tend to 1 as horizon T increase. In this model, managers have no reason to under- or overreact to recent news as they know they are not relevant for long horizon forecasts anyways. A model with pure backward-looking extrapolation would not have this property. Evidence from analyst forecasts (Bordalo et al. (2017b)) is consistent with that.

3.5.3 Time-to-build in labor

In our baseline model, capital is the only quasi-fixed factor. We consider the case where labor also needs to be hired one period ahead of production. Both capital and labor depends on productivity forecasts, potentially giving a larger role for distorted forecasts to reduce efficiency.

Appendix A.2 provides the derivation for the firm's optimal investment and hiring decision under this assumption. It also details how the model is estimated with this alternative assumption.

Table 6 shows the moments used in the estimation and the corresponding parameter estimates. The only difference with Table 5 comes from the moments related to TFP residuals: in the model with time-to-build in both labor and capital, TFP residuals can simply be computed as $\ln(p_{it}y_{it}) - \theta \ln(k_{it})$.[10]. While the persistence of TFP residuals remain the same, their overall variance is much higher (.051 vs. 008 in U.S. data). The estimated forecasting biases – $\hat{\gamma}$ and $\hat{\lambda}$ – are almost similar to what they are in the baseline model (-.31 and -.037 in this model vs. -.32 and -.032 in the baseline U.S. estimation, respectively). However, the volatility of both real innovations to TFP (σ_ω) and private information (σ_η) is much larger in this model to account for the larger volatility of overall TFP residuals.

We use the estimation in Table 6 to conduct partial equilibrium counterfactuals, simmilar to those performed in Section 3.4. We compute the unconditional average realized profits for managers with rational expectations $\mathbb{E}[\pi^*_{it}]$ and managers with distorted forecasts $\mathbb{E}[\pi^F_{it}]$, when the production exhibits time-to-build in both labor and capital. Δ is the percentage

[10] In the baseline model, there is only time-to-build in capital and labor is optimized at date t, TFP residuals are computed as: $\frac{\alpha\theta}{\Phi}\left(\ln(p_{it}y_{it}) - \Phi \ln(k_{it})\right)$.

increase in average firm-level profits obtained due to rational expectations:

$$\begin{aligned}\Delta &= \frac{\mathbb{E}[\pi^*_{it}] - \mathbb{E}[\pi^F_{it}]}{\mathbb{E}[\pi^F_{it}]} \\ &= \underset{(1.218\%)}{4.610\%} \text{ in Italian data} \\ &= \underset{(1.218\%)}{4.610\%} \text{ in U.S. data}\end{aligned}$$

Clearly, effects are much larger (about 10 times larger) under this extension, as the firm cannot as easily undo expectation errors once the shock materializes.

4 Aggregation

We nest the firm-level investment model of Section 3 into a general equilibrium framework. This allows sto explore the cost of distortions induced by non-bayesian forecasts.

4.1 Aggregation in the baseline model

We consider a simple market structure following Dixit and Stiglitz (1977). There is a continuum of intermediate input producers: at date t, firm i is a monopoly and produces a quantity y_{it} of an intermediary input at a price p_{it}. These inputs are used in the production of a final good. The final good market is perfectly competitive, and aggregates intermediate inputs with a CES technology:

$$Y_t = \left(\int_i y_{it}^{\theta} di \right)^{\frac{1}{\theta}}, \tag{7}$$

The price of the final good is normalized to 1. Profit maximization in the final good market implies that the demand for product i is given by: $p_{it} = \left(\frac{Y}{y_{it}}\right)^{1-\theta}$. There is a single labor market from which all firms hire. w_t is the wage, which firms take as given. Households have GHH preferences over leisure and consumption: $u(c_t, l_t) = \left(c_t - \frac{w_0}{L_0^{\frac{1}{\epsilon}}} \frac{l_t^{1+\frac{1}{\epsilon}}}{1+\frac{1}{\epsilon}} \right)$. As a result, labor supply is $L_0 \left(\frac{w}{w_0}\right)^{\epsilon}$ and ϵ is the constant labor supply elasticity.

We start by showing how the firm-level model of investment of Section 3 can be nested into this framework. Assume firm i production combines labor and capital with a Cobb-Douglas technology: $y_{it} = e^{z_{it}} k_{it}^{\alpha} l_{it}^{1-\alpha}$. The capital good is the final good. Log-productivity z_{it} is stochastic and follows an AR(1) process:

$$z_{it} = \rho z_{it-1} + \epsilon_{it} + \eta_{it} \quad \text{with:} \quad (\epsilon_{it}, \eta_{it}) \sim \mathcal{N}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\epsilon}^2 & 0 \\ 0 & \sigma_{\eta}^2 \end{pmatrix}\right],$$

where η_{it} is privately observed by the manager in period $t-1$. In particular, we assume no aggregate uncertainty so that aggregate output is constant $Y_t = Y$ and the equilibrium wage on the labor market is also constant $w_t = w$. Given the input producers are monopolists, profit maximization implies that firms revenue exhibit decreasing returns to scale:

$$p_{it} y_{it} = \underbrace{Y^{1-\theta}}_{=A} \underbrace{e^{\theta z_{it}}}_{=e^{\nu_{it}}} \left(k_{it}^{\alpha} l_{it}^{1-\alpha}\right)^{\theta}$$

Therefore, this model is equivalent to the baseline firm-level model discussed in Section 3: $A = Y^{1-\theta} B^{\theta}$, $\nu_{it} = \theta z_{it}$, $\omega_{it} = \theta \epsilon_{it}$ and $\psi_{it} = \theta \eta_{it}$. As in our baseline firm-level model, we assume a one period time-to-build in capital and a user cost of capital R.

Because of the time-to-build in capital, firms need to form expectations about next period productivity. Let $\mathbb{F}_{t-1}[p_{it} y_{it}]$ be the managerial forecast of date t total sales made at date $t-1$. The log-sales forecast error is $FE_{it} = \ln(p_{it} y_{it}) - \ln\left(\mathbb{F}_{t-1}[p_{it} y_{it}]\right)$.

4.2 Measuring distortions

The following proposition provides two results on aggregate efficiency loss resulting from non-rational expectations:

Proposition 2. Assume either:

- managerial log-sales forecasts are log-normally distributed

- variations in log-sales forecast error FE_{it} and log-sales forecast $\ln\left(\mathbb{F}_{t-1}\left[p_{it}y_{it}\right]\right)$ are small around their respective mean

Then aggregate TFP is simply given by:

$$\ln(TFP) = -\frac{\alpha}{2}\left(1+\frac{\alpha\theta}{1-\theta}\right)\mathbb{V}\text{ar}\left[FE_{it}\right] \tag{8}$$

Proof. See Appendix A.4. □

Proposition 2 is based on the observation that forecast errors made when investing are formally equivalent to a wedge between the real cost of capital R and the marginal productivity of capital $\alpha\theta\frac{p_{it}y_{it}}{k_{it}}$. When the realization of productivity is lower than the forecast, the firm has too much capital compared to the frictionless benchmark where there is no time to build and the firm would observe its productivity when investing. Hence, the marginal productivity of capital is lower than R in this case. As in Hsieh and Klenow (2009), when this wedge is assumed to follow a log-normal distribution, log-aggregate TFP is proportional to the dispersion in the log-wedge. In our case this assumption is equivalent to managerial forecasts following a log-normal distribution. Alternatively, one can assume small variations of wedges and TFP realizations around their means (or equivalently here, small variations in log-sales forecast errors and log-sales forecasts) and obtain a similar formula through a second-order Taylor expansion around the population average. Importantly, note that the TFP formula in Proposition 8 holds for all forecasting rules $\mathbb{F}_{t-1}$ as long as it either follows a log-normal distribution or generates only small variations around the average forecast in the population.

We derive two corollaries from the above proposition. The first one offers an upper bound for TFP loss of irrationality. This upper bound does not require estimating the structural model of Section 3. The second corollary explicitly computes the TFP loss from non-bayesian expectations defined in equation (3).

Corollary 1. For any forecasting rule satistfying the assumptions of Equation (3), the TFP losses due to imperfect foresight are:

$$\Delta \ln(TFP)^0 = \ln(TFP^{\text{perfect foresight}}) - \ln(TFP) = \frac{\alpha}{2}\left(1 + \frac{\alpha\theta}{1-\theta}\right) \mathbb{V}\text{ar}\left[FE_{it}\right]$$

In particular, the TFP losses due to observed forecasts are bounded by:

$$\Delta \ln(TFP) = \ln(TFP^{\text{rational forecasts}}) - \ln(TFP) < \Delta \ln(TFP)^0 = \frac{\alpha}{2}\left(1 + \frac{\alpha\theta}{1-\theta}\right) \mathbb{V}\text{ar}\left[\widehat{FE}_{it}\right]$$

Proof. See Appendix A.5 □

Corollary 1 is important because it provides us with a simple bound for the TFP losses due to non-rational forecasts, which requires minimal assumptions on the forecasting rules used by managers. This bound is also easily implementable, since the variance of log-sales forecast errors is directly observed in the data.

The second corollary provides us with

Corollary 2. Assume that managers form expectation using the distorted belief model of equation (3). Forecasts follow a log-normal distribution, and the TFP losses due to non-rational forecasts are given by:

$$\Delta \ln(TFP) = \frac{1}{2}\left(\frac{\theta}{1-\theta}\right)\left(\frac{\theta\alpha}{1-(1-\alpha)\theta}\right)\left(\gamma^2\rho^2\sigma_\epsilon^2 + \lambda^2\sigma_\eta^2\right) \tag{9}$$

Proof. See Appendix A.6 □

The above formula is easy to interpret: it increases with α the share of distorted input. It increases with the amount of forecast variance generated by non-biased expectations.

4.3 Estimation

We can combine the results in this section with the structural estimates recovered in Section 3.2 to quantify the TFP losses from non-rational forecasts in our context.

We can start by bounding the TFP losses using Corollary 1, which, again, does not rely on a particular forecasting rules:

$$\begin{aligned}\Delta \ln(TFP) &\leq \underbrace{\frac{\alpha}{2}\left(1+\frac{\alpha\theta}{1-\theta}\right)}_{=0.383}\mathbb{V}\text{ar}\left[\widehat{FE}_{it}\right] \\ &\leq 1.177\% \text{ in Italian data} \\ &\leq 0.298\% \text{ in U.S. data}\end{aligned}$$

As for expected profits, distortions induced by forecast errors are bigger in Italian than in US data. This is directly reflecting the fact that the variance of forecast errors in Italy is 4 times as large as in the U.S. This is due to the fact that (1) Italian firms are more volatile to start with and (2) Italian firms make bigger forecast error.

To get to the counterfactuals where managers have rational expectations, we use instead 2, combined with the structural estimates obtained in Panel B of Table 5:

$$\begin{aligned}\Delta \ln(TFP) &= \underbrace{\frac{1}{2}\left(\frac{\theta}{1-\theta}\right)\left(\frac{\theta\alpha}{1-(1-\alpha)\theta}\right)}_{=1.138}\left(\gamma^2\rho^2\sigma_\epsilon^2+\lambda^2\sigma_\eta^2\right) \\ &= 0.325\% \text{ in Italian data} \\ &= 0.014\% \text{ in U.S. data}\end{aligned}$$

The effect on Italian data is about 30 times larger than on U.S. data. This is to be expected as γ is about 2.5 times larger and σ_ω about 2 times larger. This directly comes from the fact that the persistence of forecasts errors and TFP shock variances are much

larger in Italy. Given that the TFP formula is a function of the square of $\gamma\sigma_\omega$, such a large difference is fully in line with the data.

Our conclusions remains broadly unaffected when we consider alternative calibrations for α and θ. We provide comparative statics w.r.t. these two parameters (using U.S. data) in Figure C.2. On the left panel of Figure C.2, we fix $\alpha = 1/3$ and let θ vary from .69 to .99. On the right panel, we fix $\theta = .8$ and let α vary from .1 to .7. For each calibrated value, we estimate the moments in Panel A, Table 5 and then reestimate the model following the methodology developed in Section 3.2. We then compute the TFP losses using Equation 9. Variations in α lead to modest variations in the estimated TFP: from .02% when α is equal to .1 to .5% when α is .7. For θ below .95, we get the same conclusion: going from $\theta = .7$ to $\theta = .95$ leads to TFP losses ranging from .1% to .5%. When θ becomes closer to 1, however, TFP losses can become quite significant: a high θ implies a very high degree of substitutions across firms output such that optimist firms (i.e. a firm with negative TFP news ω_{t-1}) may end up representing a large fraction of total output, distorting allocating efficicency.

4.4 Aggregation with Time-to-build in labor

With time-to-build in labor, it is direct to see that the model is equivalent to a model with a capital wedge and a labor wedge that are both equal to the log-sales forecast error:

$$\ln(1+\tau^k_{it}) = \ln(1+\tau^l_{it}) = \nu_{it} - \ln\left(\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right) = \ln(p_{it}y_{it}) - \ln\left(\mathbb{F}_{t-1}\left[p_{it}y_{it}\right]\right) = FE_{it}$$

We can apply Proposition B4 in Sraer and Thesmar (2018) and obtain:

$$\ln(TFP) = -\left(\alpha + \frac{1}{2}\frac{\theta}{1-\theta}\right)\mathbb{V}\text{ar}\left[FE_{it}\right]$$

Using the formula for the variance in the log-sales forecast error in Section A.2, it is direct to show that the TFP losses due to distorted forecast when there is time-to-build in labor are simply given by:

$$\Delta \ln(TFP) = \left(\alpha + \frac{1}{2}\frac{\theta}{1-\theta}\right)\left(\gamma^2\rho^2\sigma_\omega^2 + \lambda^2\sigma_\psi^2\right)$$
$$= 2.917\% \text{ in Italian data}$$
$$= 0.140\% \text{ in U.S. data}$$

As expected, the TFP losses are much larger with time-to-build in labor. This happens because labor is not optimally reallocated across firms in response to shocks, thereby increasing ex-post distortions.

4.5 Incorporating aggregate shocks

[TO BE DONE]

5 Conclusion

This paper incorporates forecasts into an otherwise standard neoclassical model of investment. Our model allows for managerial private information, distortion in the reports of forecasts as well as distorted forecasts due to non-rational expectations. We first document significant bias in managerial forecasts in guidance data: managerial forecast errors are persistent in a statistically significant way. This predictability is consistent with underreaction to new information by managers. When nested into a model of production, the estimated forecasting bias lead to sizable efficiency losses, both in partial and general equilibrium.

We think two deviations from the standard model could potentially allow for a greater role of distorted forecasts. First, a longer time-to-build is a clear candidate: however, our managerial forecasts data does not cover long-term forecasts, so that we cannot estimate a model with a longer time-to-build. Financial frictions may also lead to amplifications of

distorted forecasts. We leave this analysis for future research.

References

Akerlov, George and Janet Yellen, "Can Small Deviations from Rationality Make Significant Differences to EconomicEquilibria?," American Economic Review, 1985, 75, 708–720.

Arellano, Manuel and Olympia Bover, "Another look at the instrumental variable estimation of error-components models," Journal of Econometrics, July 1995, 68 (1), 29–51.

Baqaee, David and Emmanuel Farhi, "Productivity and Misallocation in General Equilibrium," 2017.

Barrero, Jose Maria, "The Micro and Macro Implications of Managers' Subjective Beliefs," Technical Report, Stanford 2018.

Ben-David, Itzhak, John Graham, and Campbell Harvey, "Managerial Miscalibration," Quarterly Journal of Economics, 2013, 128, 1547–1584.

Bloom, Nicholas, Steven J. Davis, Lucia Foster, Brian Lucking, Scott Ohlmacher, and Itay Saporta-Eksten, "Business-Level Expectations and Uncertainty," Technical Report, Stanford December 2017.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Diagnostic expectations and credit cycles," Journal of Finance, 2017.

_ , _ , Rafael LaPorta, and Andrei Shleifer, "Diagnostic Expectations and Stock Returns," Technical Report 2017.

_ , _ , Yueran Ma, and Andrei Shleifer, "Overreaction in Macroeconomic Expectations," Technical Report 2018.

Bouchaud, Jean-Philippe, Philipp Krueger, Augustin Landier, and David Thesmar, "Sticky Expectations and the Profitability Anomaly," Forthcoming Journal of Finance, 2018.

Broda, Christian and David Weinstein, "Globalizations and the Gains from Variety," Quarterly Journal of Economics, 2006, pp. 541–582.

Catherine, Sylvain, Thomas Chaney, Zongbo Huang, David Alexandre Sraer, and David Thesmar, "Aggregate Effects of Collateral Constraints," 2017.

Coibion, Olivier and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," American Economic Review, 2015.

Cooper, Russell W. and John C. Haltiwanger, "On the Nature of Capital Adjustment Costs," The Review of Economic Studies, 2006, 73 (3), 611–633.

David, Joel, Hugo Hopenhayn, and Venky Venkateswaran, "Information, Misallocation, and Aggregate Productivity," Quarterly Journal of Economics, 2016, 131, 943–1005.

Degeorge, Francois, Jayendu Patel, and Richard Zeckhauser, "Earnings management to exceed thresholds," Journal of Business, 1999, 72, 1–33.

Dixit, Avinash K and Joseph E Stiglitz, "Monopolistic Competition and Optimum Product Diversity," American Economic Review, June 1977, 67 (3), 297–308.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer, "Expectations and investment," NBER Macroeconomics Annual, 2016, 30 (1), 379–431.

Hassan, Tarek and Thomas Mertens, "The Social Cost of Near-Rational Investment," American Economic Review, 2017, 107, 1059–1103.

Hirshleifer, David, Angie Low, and Siew-Hong Teoh, "Are Overconfident CEOs Better Innovators?," The Journal of Finance, 2012, 67 (4), 1457–1498.

Hsieh, Chang-Tai and Peter Klenow, "Misallocation and Manufacturing TFP in China and India," Quarterly Journal of Economics, 2009.

Landier, Augustin and David Thesmar, "Financial Contracting with Optimistic Entrepreneurs: Theory and Evidence," Review of Financial Studies, 2009, 22, 117–150.

Malmendier, Ulrike and Geoffrey Tate, "CEO Overconfidence and Corporate Investment," Journal of Finance, 2005, 60, 2661–2700.

_ and _ , "Who makes acquisitions? CEO overconfidence and the market's reaction," Journal of Finance, 2008, 89, 20–43.

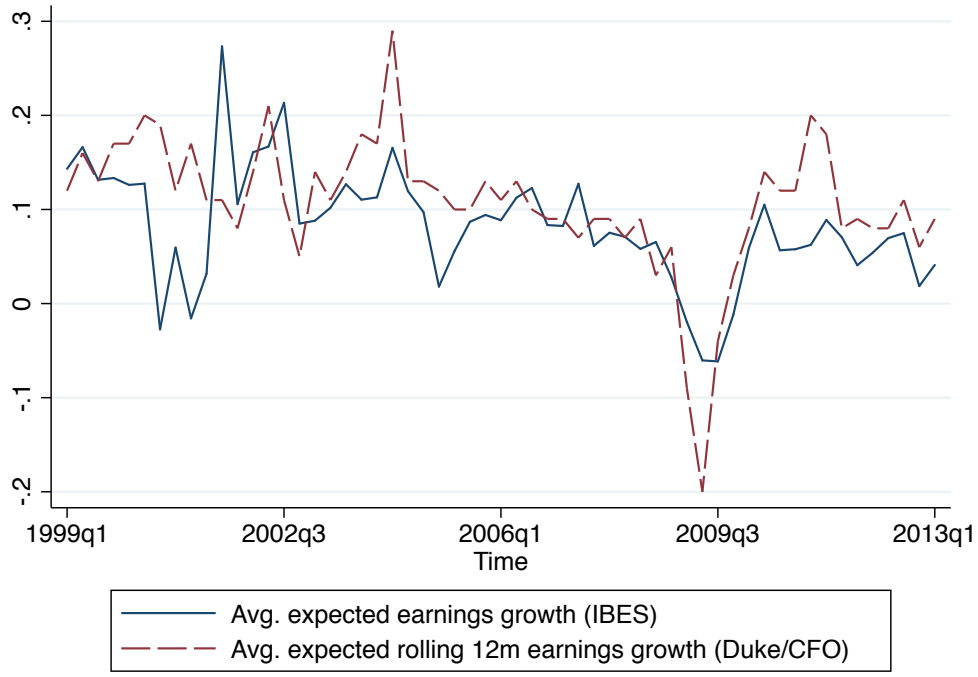
_ and Stefan Nagel, "Learning from Inflation Experiences," Quarterly Journal of Economics, 2016, 131, 53–87.

Nickell, Stephen, "Biases in Dynamic Models with Fixed Effects," Econometrica, 1981, 49 (6), 1417–26.

Sraer, David and David Thesmar, "A Sufficient Statistics Approach for Aggregating Firm-Level Experiments," NBER Working Paper 24208 2018.

FIGURES

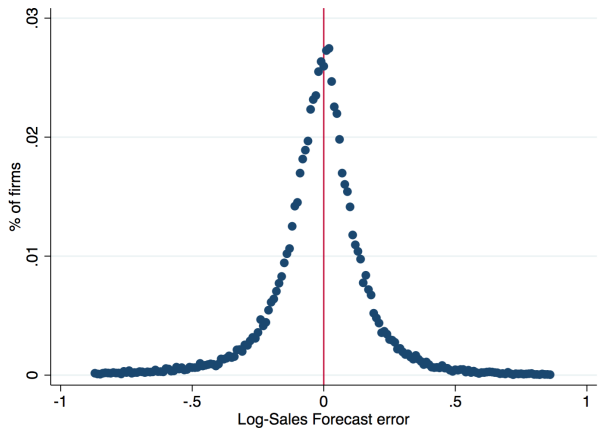
Figure 1: Sales Growth Forecasts: IBES Guidance and Duke CFO Survey



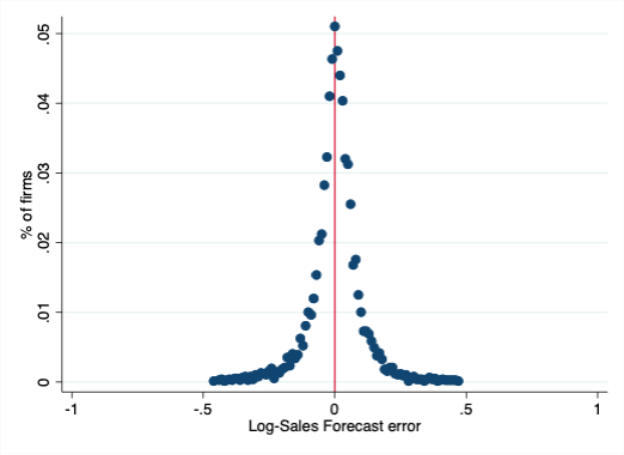
Note: This plot shows the (sales-weighted) average forecast of sales growth from IBES guidance data and from Duke CFO survey data. The IBES guidance data forecasts fiscal year sales; for each quarter, we keep forecasts that have horizons 2 to 6 quarters (i.e. between 2 to 6 quarters before the fiscal year end of the relevant fiscal year). The CFO survey forecasts are about rolling 12 months sales growth.

Figure 2: Forecast Errors: histogram

Panel A: Italian data



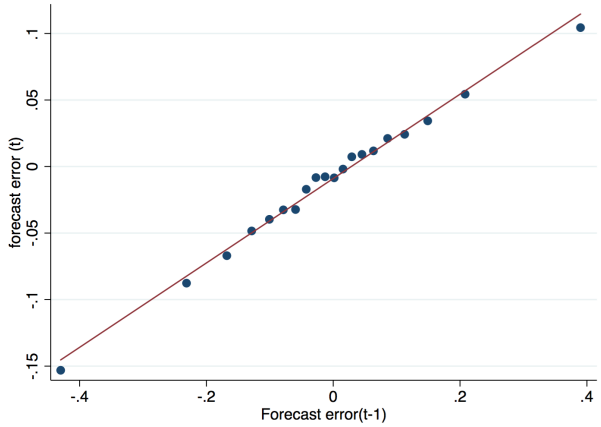
Panel B: U.S. data



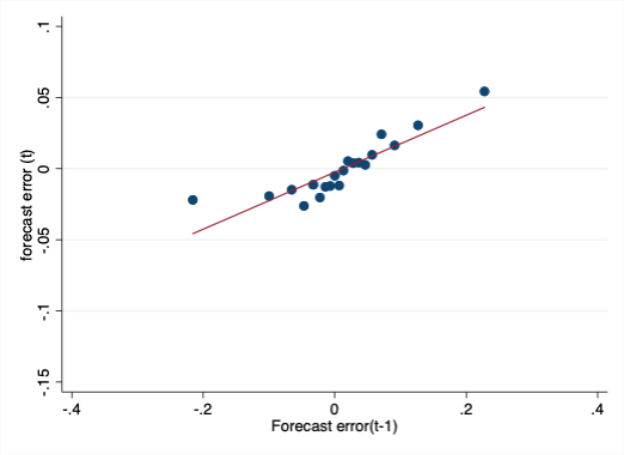
Note: This figure plots the distribution of log-sales forecast error in our two samples. The log-sales forecast error is computed as the difference between realized sales in fiscal year t and sales guidance for year t issued at the beginning of year t. Panel A uses one-year ahead sales forecasts from the business survey run on a representative sample of Italian firms by the Bank of Italy. Panel B uses one-year ahead sales forecasts coming from managerial guidances of a subset of U.S. publicly listed firms. In both panels, the scale on the x axis is the same.

Figure 3: Forecast Error Persistence: Binned Scatter Plot

Panel A: Italian data



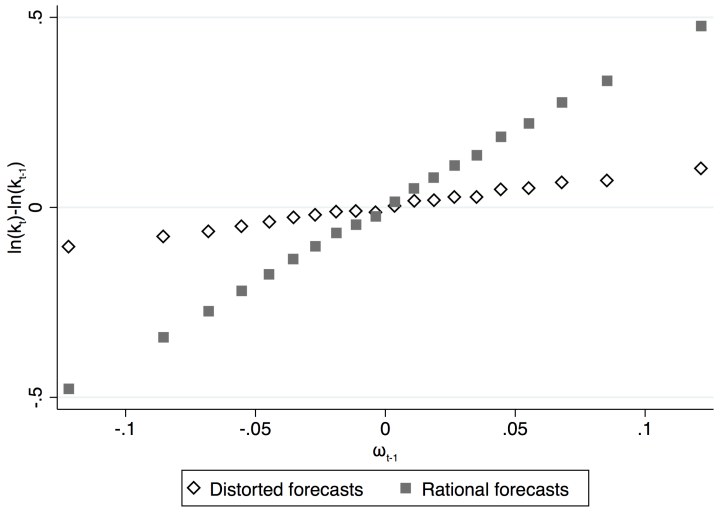
Panel B: U.S. data



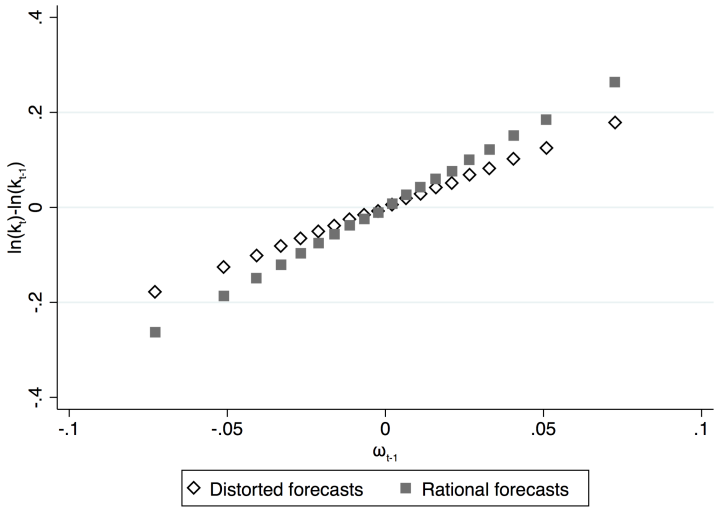
Note: This figure is a binscatter plot of year t log-sales forecast error on year t-1 log-sales forecast error. The log-sales forecast error is computed as the difference between realized sales in fiscal year t and sales guidance for year t issued at the beginning of year t. Each panel corresponds to one of our datasets. To simplify comparison across samples, the scales on the x and the y axis in both panels are the same.

Figure 4: Firm Investment with Rational vs. Distorted Forecasts

Panel A: Italian data



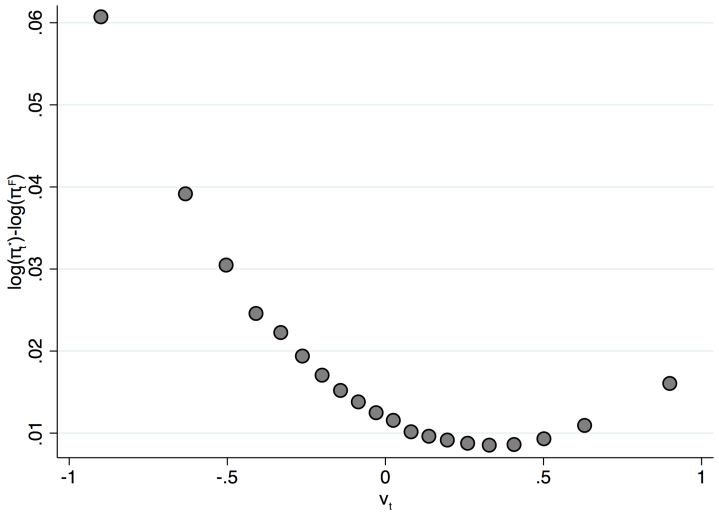
Panel B: U.S. data



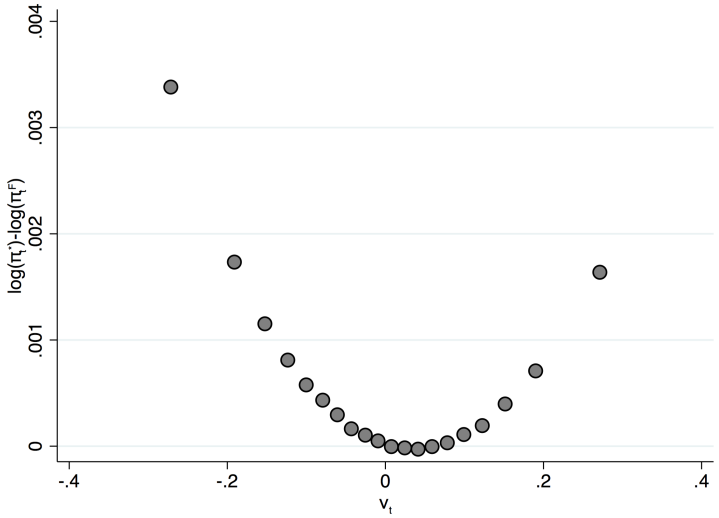
Note: This figure uses the estimation of Table 5 to simulate investment for managers using rational forecast (dark squares) or distorted forecasts (white diamonds). We construct 20 buckets of date t-1 TFP surprise (ω_{t-1}) and compute, for each of these buckets, average log capital growth for both managers with rational forecasts ($\gamma = \lambda = 0$) and managers with distorted forecasts (γ and λ at their estimated values). In Panel A, we use results from Italian data. In Panel B, we use results from U.S. data.

Figure 5: Firm Profits with Rational vs. Distorted forecasts

Panel A: Italian data



Panel B: U.S. data



Note: This figure uses the estimation of Table 5 to simulate expected profit losses for managers using distorted forecasts relative to managers using fully rational forecasts (i.e. with $\gamma = \lambda = 0$). We construct 20 buckets of date t innovations (ω_t) and compute average profit losses, for each of these buckets. In Panel A, we use Italian estimates. In Panel B, we use U.S. estimates.

TABLES

Table 1: Summary Statistics

Variable	mean	p25	p50	p75	sd	*N*
Panel A1. Italian merged survey and accounts						
Firm forecast error (log actual sales(t) - log forecast(t-1|t))	-0.01	-0.09	-0.01	0.07	0.18	37,789
Log sales	10.03	8.93	9.84	10.98	1.55	37,789
Log assets	10.09	9.00	9.88	11.01	1.56	37,789
Sales/l.assets	1.14	0.70	1.01	1.41	0.66	37,395
Capex/l.assets	0.03	0.01	0.02	0.04	0.04	36,776
Cash flow/l.assets	0.07	0.03	0.06	0.10	0.07	36,947
Panel A2. Italian accounting data						
logsales	6.16	5.00	6.20	7.36	1.90	7,350,380
logassets	6.09	4.93	6.10	7.28	1.86	8,261,704
sales/lassets	1.24	0.34	1.00	1.73	1.21	8,164,493
capex/lassets	0.02	0.00	0.01	0.03	0.03	7,650,925
cashflow/lassets	0.05	0.00	0.04	0.10	0.12	7,953,679
Panel B1. U.S. guidance sample						
Firm forecast error (log actual sales - log forecast)	0.00	-0.04	0.00	0.05	0.09	5,278
Analyst forecast error	0.00	-0.04	0.00	0.05	0.11	4,878
Log sales	7.09	5.92	7.03	8.23	1.69	5,278
Log assets	7.18	5.96	7.07	8.29	1.64	5,278
Sales/l.assets	1.18	0.70	1.02	1.43	0.72	5,270
Capex/l.assets	0.04	0.02	0.03	0.06	0.04	5,264
Cash flow/l.assets	0.19	0.05	0.13	0.28	0.19	5,278
Q	2.13	1.32	1.78	2.57	1.17	5,205
Book-to-market	0.44	0.22	0.36	0.59	0.39	5,261
Panel B2. U.S. Compustat Sample						
Log sales	4.82	3.11	5.20	6.99	3.04	49,568
Log assets	4.82	3.05	5.05	7.02	3.09	54,076
Sales/l.assets	1.07	0.35	0.82	1.48	1.02	50,022
Capex/l.assets	0.05	0.01	0.03	0.06	0.06	48,939
Cash flow/l.assets	0.24	0.03	0.12	0.35	0.27	54,069
Q	1.88	1.12	1.47	2.20	1.20	33,370
Book-to-market	0.56	0.25	0.46	0.76	0.54	33,455

Note: Summary statistics. Panels A1 and A2 correspond to Italian data. Panel A1 focuses on Italian firms for which sales forecasts were available from the Bank of Italy. Panel A2 produces the same descriptive statistics for the universe of Italian firms present in the Company Accounts Data System. Panels B1 and B2 are U.S. data. Panel B1 focuses on COMPUSTAT (i.e. listed) firms for which managerial guidances – at least 5 of them – and analyst forecasts were available. Panel B2 produces the same descriptive statistics for all COMPUSTAT firms. Firm forecast error is the logarithm of actual sales for the fiscal year minus the log of the guidance issued in the first quarter for the coming fiscal year. Analyst forecast error is the logarithm of actual sales minus the log of the analysts consensus forecast in the first quarter of the coming fiscal year. Q is the ratio of market value to total assets. Book-to-market is the ratio of the book value of equity to market value of equity. Market based variables are unavailable for Italian firms, most of which are not publicly listed. All variables are winsorized at the median +/- 5 times the interquartile range.

Table 2: Predicting Actual Sales

	Actual Sales							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Italian data								
Firm forecast	0.974***		0.862***		0.854***	0.836***	0.822***	0.664***
	(0.007)		(0.017)		(0.016)	(0.018)	(0.016)	(0.020)
Log(l.Sales)			0.139***		0.147***	0.159***	0.175***	0.141***
			(0.018)		(0.017)	(0.021)	(0.017)	(0.028)
Log(l.Assets)			-0.136***		-0.145***	-0.158***	-0.173***	-0.289***
			(0.018)		(0.016)	(0.021)	(0.017)	(0.023)
Constant	0.027***		0.136***					
	(0.006)		(0.023)					
Fixed effects	No		No		Year	Industry	Ind-Year	Firm&Year
Observations	37,317		37,315		37,315	37,315	37,153	37,261
Adj R^2	0.90		0.90		0.91	0.91	0.91	0.93
Panel B: U.S. data								
Firm forecast	1.005***	1.076***	1.003***	1.052***	0.998***	1.003***	0.997***	0.948***
	(0.005)	(0.034)	(0.009)	(0.029)	(0.009)	(0.009)	(0.009)	(0.021)
Analyst forecast		-0.072**		-0.056				
		(0.035)		(0.035)				
l.Sales			-0.002	0.006	0.005	-0.003	0.005	-0.019
			(0.011)	(0.013)	(0.010)	(0.010)	(0.009)	(0.017)
Q			0.006**	0.006***	0.006**	0.007**	0.006**	-0.001
			(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.004)
Log(l.Assets)			-0.001	-0.009	-0.006	0.001	-0.006	-0.074***
			(0.011)	(0.013)	(0.011)	(0.010)	(0.009)	(0.015)
Constant	0.002	0.004	0.005	0.020				
	(0.005)	(0.005)	(0.019)	(0.019)				
Fixed effects	No	No	No	No	Year	Industry	Ind-Year	Firm&Year
Observations	5,257	4,862	5,173	4,786	5,173	5,173	5,171	5,172
Adj R^2	0.97	0.97	0.97	0.97	0.97	0.97	0.97	0.98

Note: Panel A focuses on Italian firms for which sales forecasts were available from the Bank of Italy. Panel B focuses on COMPUSTAT (i.e. listed) firms for which managerial guidances and analyst forecasts were available – at least 5 of them. We use OLS to estimate $sales_{it}/assets_{i,t-1} = \alpha + F_{i,t-1}sales_{it}/assets_{i,t-1} + controls_{i,t-1}$, where $sales_{it}$ is sales of firm i in fiscal year t, $F_{i,t-1}sales_{it}$ is beginning-of-year forecast (ealiest forecast in first fiscal quarter) of fiscal year t sales (normalized by lagged assets). Industry dummies are the 12 Fama-French industries in the U.S. data and 2-digit industry dummies in Italian data. Standard errors are clustered by both firm and time. ***, ** and * means statistically significant at the 1%, 5% and 10% confidence level.

Table 3: Explaining Capital Expenditures

	Capital Expenditures							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Italian data								
Firm forecast	0.004***		0.002***		0.002***	0.004***	0.003***	0.010***
	(0.001)		(0.000)		(0.000)	(0.001)	(0.001)	(0.001)
l.capx			0.435***		0.431***	0.411***	0.408***	
			(0.014)		(0.014)	(0.015)	(0.015)	
Cash flow			0.074***		0.072***	0.071***	0.069***	0.049***
			(0.006)		(0.005)	(0.006)	(0.005)	(0.006)
l.Cash flow			0.024***		0.023***	0.025***	0.023***	0.046***
			(0.005)		(0.004)	(0.005)	(0.004)	(0.004)
Log(l.Assets)			0.000*		0.000**	0.000	0.000**	-0.008***
			(0.000)		(0.000)	(0.000)	(0.000)	(0.001)
Constant	0.025***		0.004**					
	(0.002)		(0.002)					
Fixed effects	No		No		Year	Industry	Ind-Year	Firm&Year
Observations	36,569		30,259		30,259	30,259	30,105	30,725
Adj R^2	0.01		0.27		0.28	0.29	0.29	0.37
Panel B: U.S. data								
Firm forecast	0.016***	0.017**	0.004***	0.023***	0.004***	0.004***	0.004***	0.020***
	(0.002)	(0.008)	(0.001)	(0.006)	(0.001)	(0.001)	(0.001)	(0.004)
Analyst forecast		-0.001		-0.019***				
		(0.008)		(0.007)				
l.capx			0.742***	0.745***	0.746***	0.727***	0.733***	
			(0.025)	(0.024)	(0.025)	(0.027)	(0.027)	
Q			0.002***	0.002***	0.002***	0.003***	0.002***	0.006***
			(0.000)	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)
Cash flow			-0.041***	-0.041***	-0.037***	-0.040***	-0.037***	-0.043***
			(0.008)	(0.008)	(0.007)	(0.008)	(0.007)	(0.006)
l.Cash flow			0.030***	0.031***	0.026***	0.032***	0.028***	0.014**
			(0.007)	(0.008)	(0.007)	(0.007)	(0.006)	(0.006)
Log(l.Assets)			-0.001**	-0.001***	-0.001**	-0.001*	-0.001*	-0.009***
			(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.003)
Constant	0.025***	0.025***	0.007***	0.007***				
	(0.003)	(0.003)	(0.003)	(0.002)				
Fixed effects	No	No	No	No	Year	Industry	Ind-Year	Firm&Year
Observations	5,241	4,853	5,147	4,770	5,147	5,147	5,145	5,157
Adj R^2	0.07	0.07	0.66	0.67	0.67	0.67	0.68	0.70

Note: Panel A focuses on Italian firms for which sales forecasts were available from the Bank of Italy. Panel B focuses on COMPUSTAT (i.e. listed) firms for which managerial guidances and analyst forecasts were available – at least 5 of them. We use OLS to estimate $capx_{it}/assets_{i,t-1} = \alpha + F_{i,t-1}sales_{it}/assets_{i,t-1} + controls_{it}$, where $capx_{it}$ is capital expenditures of firm i in fiscal year t, $F_{i,t-1}sales_{it}$ is beginning-of-year forecast (ealiest forecast in first fiscal quarter) of fiscal year t sales (normalized by lagged assets). Industry dummies are the 12 Fama-French industries in the U.S. data and 2-digit industry dummies in Italian data. ***, ** and * means statistically significant at the 1%, 5% and 10% confidence level.

Table 4: Persistence of Forecast Errors

	Forecast error			
	(1)	(2)	(3)	(4)
	Manager	Analyst	Manager	Analyst
Panel A: Italian business survey data				
Forecast error(t-1)	0.324***		0.172***	
	(0.017)		(0.014)	
Year FE	Yes		Yes	
Firm FE	No		Yes	
Observations	32,383		18,628	
Adj R^2	0.12			
Panel B: U.S. data				
Forecast error(t-1)	0.179***	0.172***	0.115***	0.119***
	(0.041)	(0.037)	(0.029)	(0.017)
Year FE	Yes	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Observations	4,251	18,552	1,892	12,248
Adj R^2	0.08	0.11		

Note: Panel A focuses on Italian firms for which sales forecasts were available from the Bank of Italy business survey. Panel B focuses on COMPUSTAT (i.e. listed) firms for which managerial guidances and analyst forecasts were available – at least 5 of them. In all specifications, we regress the log-sales forecast error (log actual sales minus log sales forecast made in first quarter of fiscal year t) on log-sales forecast error of year $t-1$. Columns (1) and (2) include year fixed-effects and are estimated using OLS on the sample of firms with at least 5 sales forecasts. Columns (3) and (4) include both year and firm fixed-effects, and are estimated using dynamic panel GMM (Arellano and Bover (1995)). In the U.S., forecasts are measured using managerial guidances (cols (1) and (3)) or analyst forecast consensus (cols (2) and (4)). Standard errors are clustered by both firm and time. ***, ** and * means statistically significant at the 1%, 5% and 10% confidence level.

Table 5: Structural Estimation of Baseline Investment Model with Distorted Beliefs

Panel A: Italian business survey data

Moments:

$\hat{\chi}$	$\hat{\sigma}_\tau^2$	$\widehat{\mathbb{V}\mathrm{ar}}[\widehat{FE}_{it}]$	$\hat{\kappa}_1$	$\hat{\kappa}_2$	$\hat{\kappa}_3$
0.789	0.072	0.031	0.321	0.660	0.571
(0.056)	(0.004)	(0.001)	(0.011)	(0.023)	(0.030)

Estimates:

$\hat{\gamma}$	$\hat{\lambda}$	$\hat{\rho}$	$\hat{\sigma}_\omega$	$\hat{\sigma}_\psi$	$\hat{\sigma}_\zeta$
-0.781	-0.087	0.789	0.059	0.262	0.115
(0.092)	(0.006)	(0.056)	(0.003)	(0.007)	(0.004)

Panel B: U.S. Managerial guidance data

Moments:

$\hat{\chi}$	$\hat{\sigma}_\tau^2$	$\widehat{\mathbb{V}\mathrm{ar}}[\widehat{FE}_{it}]$	$\hat{\kappa}_1$	$\hat{\kappa}_2$	$\hat{\kappa}_3$
0.727	0.008	0.008	0.175	0.406	0.787
(0.031)	(0.001)	(0.000)	(0.024)	(0.022)	(0.072)

Estimates

$\hat{\gamma}$	$\hat{\lambda}$	$\hat{\rho}$	$\hat{\sigma}_\omega$	$\hat{\sigma}_\psi$	$\hat{\sigma}_\zeta$
-0.324	-0.032	0.727	0.035	0.083	0.041
(0.069)	(0.016)	(0.031)	(0.002)	(0.005)	(0.007)

Note: Panel A focuses on Italian firms for which sales forecasts were available from the Bank of Italy business survey. Panel B focuses on COMPUSTAT (i.e. listed) firms for which managerial guidances and analyst forecasts were available – at least 5 of them. Moments are defined as follows. We first define $\nu_{it} = \left(\frac{\alpha\theta}{\Phi}\right)(\ln(p_{it}y_{it}) - \Phi\ln(k_{it}))$. We estimate $\hat{\nu_{it}} = \delta_i + \delta_t + \chi\hat{\nu_{it}} + \tau_{it}$ using dynamic panel GMM. $\hat{\chi}$ and $\hat{\sigma}_\tau^2$ are the estimated persistence and variance of residuals. $\widehat{\mathbb{V}\mathrm{ar}}[\widehat{FE}_{it}]$ is the variance of residuals from a regression of log-sales forecast errors on year fixed-effects. $\hat{\kappa}_1$ is the estimated coefficient from a regression of log-sales forecast error on lagged log-sales forecast error, controlling for year fixed-effects. $\hat{\kappa}_2$ is the estimated coefficient of a regression of date t productivity residual, $\widehat{\tau_{it}}$ on date-t reported log-sales forecast error, controlling for firm and year fixed-effects. $\hat{\kappa}_3$ is the estimated coefficient of a regression of the log sales-to-capital ratio, $\ln(p_{it}y_{it}) - \ln(k_{it})$, on date-t reported log-sales forecast error, controlling for firm and year fixed-effects. $\hat{\gamma}$ and $\hat{\lambda}$ are the estimated coefficient characterizing distorted expectations. $\hat{\rho}$ is the estimated persistence of TFP. $\hat{\sigma_\omega}$ is the estimated volatility of TFP innovations. $\hat{\sigma}_\psi$ is the volatility of private information. $\hat{\sigma}_\zeta$ is the volatility of noise introduced by managers in reported forecasts. Standard error are obtained by bootstrapping on the estimation sample using a block bootstrap at the firm-level.

Table 6: Structural Estimation of Investment Model with Distorted Beliefs and Time-to-build in Labor

Panel A: Italian business survey data

TO BE DONE

Panel B: U.S. Managerial guidance data

Moments

$\hat{\chi}$	$\hat{\sigma}_\tau^2$	$\widehat{\mathbb{V}\mathrm{ar}}[\widehat{FE}_{it}]$	$\hat{\kappa}_1$	$\hat{\kappa}_2$	$\hat{\kappa}_3$
0.757	0.051	0.008	0.175	0.951	0.787
(0.030)	(0.004)	(0.000)	(0.024)	(0.048)	(0.072)

Estimates

$\hat{\gamma}$	$\hat{\lambda}$	$\hat{\rho}$	$\hat{\sigma}_\omega$	$\hat{\sigma}_\psi$	$\hat{\sigma}_\zeta$
-0.313	-0.037	0.757	0.076	0.212	0.041
(0.066)	(0.012)	(0.030)	(0.004)	(0.010)	(0.007)

Note: Panel A focuses on Italian firms for which sales forecasts were available from the Bank of Italy business survey. Panel B focuses on COMPUSTAT (i.e. listed) firms for which managerial guidances and analyst forecasts were available – at least 5 of them. Moments are defined as follows. We first define $\nu_{it} = \ln(p_{it}y_{it}) - \theta \ln(k_{it})$. We estimate $\hat{\nu_{it}} = \delta_i + \delta_t + \chi\hat{\nu_{it}} + \tau_{it}$ using dynamic panel GMM. $\hat{\chi}$ and $\hat{\sigma}_\tau^2$ are the estimated persistence and variance of residuals. $\widehat{\mathbb{V}\mathrm{ar}}[\widehat{FE}_{it}]$ is the variance of residuals from a regression of log-sales forecast errors on year fixed-effects. $\hat{\kappa}_1$ is the estimated coefficient from a regression of log-sales forecast error on lagged log-sales forecast error, controlling for year fixed-effects. $\hat{\kappa}_2$ is the estimated coefficient of a regression of date t productivity residual, $\widehat{\tau_{it}}$ on date-t reported log-sales forecast error, controlling for firm and year fixed-effects. $\hat{\kappa}_3$ is the estimated coefficient of a regression of the log sales-to-capital ratio, $\ln(p_{it}y_{it}) - \ln(k_{it})$, on date-t reported log-sales forecast error, controlling for firm and year fixed-effects. $\hat{\gamma}$ and $\hat{\lambda}$ are the estimated coefficient characterizing distorted expectations. $\hat{\rho}$ is the estimated persistence of TFP. $\hat{\sigma_\omega}$ is the estimated volatility of TFP innovations. $\hat{\sigma}_\psi$ is the volatility of private information. $\hat{\sigma}_\zeta$ is the volatility of noise introduced by managers in reported forecasts. Standard error are obtained by bootstrapping on the estimation sample using a block bootstrap at the firm-level.

APPENDIX – FOR ONLINE PUBLICATION

A Proofs

A.1 Proof of Proposition 1

At date t, the firm hires employees after observing date t revenue-based productivity to maximize profits:

$$\Pi_{it} = \max_{l_{it}} \left\{ A e^{\nu_{it}} k_{it}^{\alpha\theta} l_{it}^{(1-\alpha)\theta} - w l_{it} \right\} = \Omega e^{\frac{\Phi}{\alpha\theta}\nu_{it}} k_{it}^{\Phi},$$

where $\Phi = \frac{\alpha\theta}{1-(1-\alpha)\theta}$ and $\Omega = (1-(1-\alpha)\theta)\left(\frac{(1-\alpha)\theta}{w}\right)^{\frac{1-\alpha}{\alpha}\Phi} A^{\frac{\Phi}{\alpha\theta}}$. Let $\mathbb{F}_{it-1}[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}]$ be firm i's forecast at date $t-1$. At date $t-1$, firm i's capital stock is purchased to maximize expected profits:

$$\max_{k_{it}} \left\{ \Omega \mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right] k_{it}^{\Phi} - R k_{it} \right\} \Rightarrow k_{it} = \left(\frac{\Phi}{R}\right)^{\frac{1}{1-\Phi}} \Omega^{\frac{1}{1-\Phi}} \left(\mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right]\right)^{\frac{1}{1-\Phi}}$$

With our formulation of distorted expectations:

$$\ln\left(\mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right]\right) = \frac{\Phi}{\alpha\theta}\left(\rho(\nu_{it-1} + \gamma\omega_{it-1}) + (1+\lambda)\psi_{it}\right) + \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2 \sigma_\omega^2$$

Since k_{it} is purchased at date $t-1$, the date t-1 true forecast for date-t sales is:

$$\mathbb{F}_{it-1}\left[p_{it}y_{it}\right] = \frac{\Omega}{1-(1-\alpha)\theta} \mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right] k_{it}^{\Phi}$$

So that the log-sales forecast error at date t is:

$$\begin{aligned} FE_{it} &= \ln(p_{it}y_{it}) - \ln\left(\mathbb{F}_{it-1}\left[p_{it}y_{it}\right]\right) \\ &= \frac{\Phi}{\alpha\theta}\nu_{it} - \ln\left(\mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right]\right) \\ &= -\frac{\Phi}{\alpha\theta}\left(\gamma\rho\omega_{it-1} + \lambda\psi_{it}\right) + \frac{\Phi}{\alpha\theta}\omega_{it} - \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2\sigma_\omega^2 \end{aligned}$$

$\frac{\Phi}{\alpha\theta}\omega_{it} - \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2\sigma_\omega^2$ corresponds to rational expectation errors.$-\frac{\Phi}{\alpha\theta}\left(\gamma\rho\omega_{it-1} + \lambda\psi_{it}\right)$ corresponds to expectation errors due to managers' distorted forecasts.

Accounting for mis-reporting of true forecasts, observed forecast errors are given by:

$$\begin{aligned}\widehat{FE}_{it} &= \ln(p_{it}y_{it}) - \ln\left(\widehat{\mathbb{F}}_{it-1}\left[p_{it}y_{it}\right]\right) \\ &= \ln(p_{it}y_{it}) - \ln\left(\mathbb{F}_{it-1}\left[p_{it}y_{it}\right]\right) - \zeta_{it} \\ &= -\frac{\Phi}{\alpha\theta}\left(\gamma\rho\omega_{it-1} + \lambda\psi_{it}\right) + \frac{\Phi}{\alpha\theta}\omega_{it} - \frac{1}{2}\left(\frac{\Phi}{\alpha\theta}\right)^2\sigma_\omega^2 - \zeta_{it},\end{aligned}$$

where ζ_{it} is the "noise" introduced by managers in their reported forecasts.

The variance of log-sales forecast errors in the data is therefore given by:

$$\mathbb{V}\text{ar}[\widehat{FE}_{it}] = \sigma_\zeta^2 + \left(\frac{\Phi}{\alpha\theta}\right)^2\left((1+\gamma^2\rho^2)\sigma_\omega^2 + \lambda^2\sigma_\psi^2\right)$$

The covariance of date-t and date-t-1 reported log-sales forecast errors writes:

$$\mathbb{C}\text{ov}\left[\widehat{FE}_{it}, \widehat{FE}_{it-1}\right] = -\left(\frac{\Phi}{\alpha\theta}\right)^2\gamma\rho\sigma_\omega^2$$

Distorted beliefs lead to persistence in forecast errors. An unusually large innovation ω_{it-1} implies a positive forecast error today. For an agent over-weighting such unusually large realization (i.e. $\gamma > 0$), this large innovation means a high forecast for date t sales, which leads, on average, to a negative forecast error at date t.

A regression of reported log-sales forecast errors at date t on reported log-sales forecast errors at date $t-1$ leads to a regression coefficient κ_1:

$$\kappa_1 = -\frac{\left(\frac{\Phi}{\alpha\theta}\right)^2\gamma\rho\sigma_\omega^2}{\sigma_\zeta^2 + \left(\frac{\Phi}{\alpha\theta}\right)^2\left((1+\gamma^2\rho^2)\sigma_\omega^2 + \lambda^2\sigma_\psi^2\right)}$$

The covariance of log-productivity innovations (as measured by the econometrician) and reported log-sales forecast errors is:

$$\mathbb{C}\text{ov}\left[\omega_{it} + \psi_{it}, \widehat{FE}_{it}\right] = \left(\frac{\Phi}{\alpha\theta}\right)\left(\sigma_\omega^2 - \lambda\sigma_\psi^2\right)$$

A regression of date t log-productivity innovations on date-t reported log-sales forecast leads to a regression coefficient κ_2:

$$\kappa_2 = \frac{\left(\frac{\Phi}{\alpha\theta}\right)\left(\sigma_\omega^2 - \lambda\sigma_\psi^2\right)}{\sigma_\zeta^2 + \left(\frac{\Phi}{\alpha\theta}\right)^2\left((1+\gamma^2\rho^2)\sigma_\omega^2 + \lambda^2\sigma_\psi^2\right)}$$

Finally, from the formula for log-sales forecast, note that:

$$\mathbb{F}_{it-1}\left[p_{it}y_{it}\right] = \left(\frac{\Phi}{R}\right)^{\frac{\Phi}{1-\Phi}}\frac{\Omega^{\frac{1}{1-\Phi}}}{1-(1-\alpha)\theta}\left(\mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right]\right)^{\frac{1}{1-\Phi}} = \frac{R}{\alpha\theta}k_{it}$$

Therefore, the sales to capital ratio is related to the true log-sales forecast in the following way:

$$\ln(p_{it}y_{it}) - \ln(k_{it}) = \ln(p_{it}y_{it}) - \ln(\mathbb{F}_{it-1}\left[p_{it}y_{it}\right]) + \ln\left(\frac{\alpha\theta}{R}\right)$$

Since we observe only reported log-sales forecast:

$$\ln(p_{it}y_{it}) - \ln(k_{it}) = \ln(p_{it}y_{it}) - \ln(\widehat{\mathbb{F}}_{it-1}\left[p_{it}y_{it}\right]) + \zeta_{it} + \ln\left(\frac{\alpha\theta}{R}\right)$$

As a result, a regression of the log-sales-to-capital ratio on the reported log-sales forecast error should have a coefficient of:

$$\hat{\kappa}_3 = 1 - \frac{\sigma_\zeta^2}{\mathbb{V}\text{ar}\left[\widehat{FE}_{it}\right]}$$

A.1.1 Estimation of model with adjustment cost

A.2 Derivation of model with time-to-build in labor

With time-to-build in labor and capital, the firm maximizes:

$$\max_{k_{it},l_{it}} A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\left(k_{it}^{\alpha}l_{it}^{1-\alpha}\right)^{\theta} - wl_{it} - Rk_{it}$$

The first-order condition in labor implies:

$$l_{it} = \left(\frac{(1-\alpha)\theta}{w}A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right)^{\frac{\Phi}{\alpha\theta}} k_{it}^{\Phi}$$

And the first-order condition in capital leads writes:

$$\frac{\alpha\theta}{R}A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right] l_{it}^{(1-\alpha)\theta}k^{\alpha\theta-1} = 1$$

After injecting the labor FOC and simplifying, this leads to:

$$k_{it} = \left(\frac{\alpha\theta}{R}\right)^{\frac{1-(1-\alpha)\theta}{1-\theta}}\left(\frac{(1-\alpha)\theta}{w}\right)^{\frac{(1-\alpha)\theta}{1-\theta}}\left(A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right)^{\frac{1}{1-\theta}}$$

And labor demand is simply:

$$l_{it} = \left(\frac{\alpha\theta}{R}\right)^{\frac{\alpha\theta}{1-\theta}}\left(\frac{(1-\alpha)\theta}{w}\right)^{\frac{1-\alpha\theta}{1-\theta}}\left(A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right)^{\frac{1}{1-\theta}}$$

Therefore, the firm revenue is simply:

$$p_{it}y_{it} = Ae^{\nu_{it}}k_{it}^{\alpha\theta}l_{it}^{(1-\alpha)\theta} = A^{\frac{1}{1-\theta}}\left(\frac{\alpha\theta}{R}\right)^{\frac{\alpha\theta}{1-\theta}}\left(\frac{(1-\alpha)\theta}{w}\right)^{\frac{(1-\alpha)\theta}{1-\theta}} e^{\nu_{it}}\left(\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right)^{\frac{\theta}{1-\theta}}$$

Forecasted sales at date $t-1$ is just simply:

$$\mathbb{F}_{t-1}\left[p_{it}y_{it}\right] = A^{\frac{1}{1-\theta}}\left(\frac{\alpha\theta}{R}\right)^{\frac{\alpha\theta}{1-\theta}}\left(\frac{(1-\alpha)\theta}{w}\right)^{\frac{(1-\alpha)\theta}{1-\theta}}\left(\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right)^{\frac{1}{1-\theta}}$$

And the log-sales forecast error is given by:

$$\ln(p_{it}y_{it}) - \ln\left(\mathbb{F}_{t-1}\left[p_{it}y_{it}\right]\right) = \nu_{it} - \ln\left(\mathbb{F}_{t-1}\left[e^{\nu_{it}}\right]\right) = -\left(\gamma\rho\omega_{it-1} + \lambda\psi_{it}\right) + \omega_{it} - \frac{1}{2}\sigma_\omega^2$$

The variance of log-sales forecast errors in the data is therefore given by:

$$\mathbb{V}\text{ar}[\widehat{FE}_{it}] = \sigma_\zeta^2 + \left((1+\gamma^2\rho^2)\sigma_\omega^2 + \lambda^2\sigma_\psi^2\right)$$

A regression of reported log-sales forecast errors at date t on reported log-sales forecast errors at date

$t-1$ leads to a regression coefficient κ_1:

$$\kappa_1 = -\frac{\gamma\rho\sigma_\omega^2}{\sigma_\zeta^2 + (1+\gamma^2\rho^2)\sigma_\omega^2 + \lambda^2\sigma_\psi^2}$$

A regression of date t log-productivity innovations on date-t reported log-sales forecast leads to a regression coefficient κ_2:

$$\kappa_2 = \frac{\sigma_\omega^2 - \lambda\sigma_\psi^2}{\sigma_\zeta^2 + (1+\gamma^2\rho^2)\sigma_\omega^2 + \lambda^2\sigma_\psi^2}$$

Finally, from the formula for log-sales forecast, note that:

$$\mathbb{F}_{it-1}\left[p_{it}y_{it}\right] = \left(\frac{\Phi}{R}\right)^{\frac{\Phi}{1-\Phi}} \frac{\Omega^{\frac{1}{1-\Phi}}}{1-(1-\alpha)\theta} \left(\mathbb{F}_{it-1}\left[e^{\frac{\Phi}{\alpha\theta}\nu_{it}}\right]\right)^{\frac{1}{1-\Phi}} = \frac{R}{\alpha\theta}k_{it}$$

Therefore, the sales to capital ratio is related to the true log-sales forecast in the following way:

$$\ln(p_{it}y_{it}) - \ln(k_{it}) = \ln(p_{it}y_{it}) - \ln(\mathbb{F}_{it-1}\left[p_{it}y_{it}\right]) + \ln\left(\frac{\alpha\theta}{R}\right)$$

Since we observe only reported log-sales forecast:

$$\ln(p_{it}y_{it}) - \ln(k_{it}) = \ln(p_{it}y_{it}) - \ln(\widehat{\mathbb{F}}_{it-1}\left[p_{it}y_{it}\right]) + \zeta_{it} + \ln\left(\frac{\alpha\theta}{R}\right)$$

As a result, a regression of the log-sales-to-capital ratio on the reported log-sales forecast error should have a coefficient of:

$$\hat{\kappa}_3 = 1 - \frac{\sigma_\zeta^2}{\mathbb{V}\text{ar}\left[\widehat{FE}_{it}\right]}$$

We can therefore estimate this extended model as we did the baseline model. We start by first computing TFP residual. Note that with time-to-build in both labor and capital, the capital labor ratio is constant: $l_{it} = \left(\frac{(1-\alpha)\theta}{w}\right)\left(\frac{\alpha\theta}{R}\right)k_{it}$. Therefore, $p_{it}y_{it} = A\left(\frac{(1-\alpha)\theta}{w}\right)^{(1-\alpha)\theta}\left(\frac{\alpha\theta}{R}\right)^{(1-\alpha)\theta} e^{\nu_{it}}k_{it}^{\theta}$. We can thus construct TFP residuals as $\ln(p_{it}y_{it}) - \theta\ln(k_{it})$.

A.3 Derivation of model with CES production function

At date t, the firm maximizes profit by selecting labor:

$$\max_{l_{it}} Ae^{\nu_{it}}\left[\alpha k_{it}^{\frac{\xi-1}{\xi}} + (1-\alpha)l_{it}^{\frac{\xi-1}{\xi}}\right]^{\theta\frac{\xi}{\xi-1}} - wl_{it}$$

The first-order condition in labor implies:

$$Ae^{\nu_{it}}l^{-\frac{1}{\xi}}\left[\alpha k_{it}^{\frac{\xi-1}{\xi}} + (1-\alpha)l_{it}^{\frac{\xi-1}{\xi}}\right]^{\theta\frac{\xi}{\xi-1}-1} = \frac{w}{(1-\alpha)\theta}$$

This defines labor choice l_{it} as a function of k_{it}, the capital stock selected at date $t-1$ and ν_{it}, the realization of the date t productivity shock: $l(k_{it}, \nu_{it})$

At date $t-1$, the firm maximizes date t revenue by selecting its capital stock

$$\max_{k_{it}} A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\left[\alpha k_{it}^{\frac{\xi-1}{\xi}}+(1-\alpha)l(k_{it},\nu_{it})^{\frac{\xi-1}{\xi}}\right]^{\theta\frac{\xi}{\xi-1}}-wl(k_{it},\nu_{it})\right]-Rk_{it}$$

This leads to the following first-order condition:

$$A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\theta\left(\alpha k_{it}^{-\frac{1}{\xi}}+(1-\alpha)\frac{\partial l}{\partial k}(k_{it},\nu_{it})l(k_{it},\nu_{it})^{-\frac{1}{\xi}}\right)\left[\alpha k_{it}^{\frac{\xi-1}{\xi}}+(1-\alpha)l(k_{it},\nu_{it})^{\frac{\xi-1}{\xi}}\right]^{\theta\frac{\xi}{\xi-1}-1}-w\frac{\partial l}{\partial k}(k_{it},\nu_{it})\right]=R$$

Injecting the labor FOC, this implies:

$$A\mathbb{F}_{t-1}\left[e^{\nu_{it}}\theta\alpha k_{it}^{-\frac{1}{\xi}}\left[\alpha k_{it}^{\frac{\xi-1}{\xi}}+(1-\alpha)l(k_{it},\nu_{it})^{\frac{\xi-1}{\xi}}\right]^{\theta\frac{\xi}{\xi-1}-1}\right]=R,$$

which can also be rewritten as:

$$\mathbb{F}_{t-1}\left[l(k_{it},\nu_{it})^{\frac{1}{\xi}}\right]=\left(\frac{R}{\alpha\theta}\frac{(1-\alpha)\theta}{w}\right)k_{it}^{\frac{1}{\xi}}$$

We assume that variations in ν_{it} are small around its forecasted mean: $\mathbb{F}_{t-1}[\nu_{it}]$. We consider a first-order approximation to the first-order condition:

$$\mathbb{F}_{t-1}\left[l(k_{it},\nu_{it})^{\frac{1}{\xi}}\right]=\mathbb{F}_{t-1}\left[l(k_{it},\mathbb{F}_{t-1}[\nu_{it}])^{\frac{1}{\xi}}+(\nu_{it}-\mathbb{F}_{t-1}[\nu_{it}])\left(\frac{1}{\xi}\right)\frac{\partial l}{\partial \nu}(k_{it},\mathbb{F}_{t-1}[\nu_{it}])\,l(k_{it},\mathbb{F}_{t-1}[\nu_{it}])^{\frac{1}{\xi}-1}\right]$$

This simplifies into:

$$l(k_{it},\mathbb{F}_{t-1}[\nu_{it}])^{\frac{1}{\xi}}=\left(\frac{R}{\alpha\theta}\frac{(1-\alpha)\theta}{w}\right)k_{it}^{\frac{1}{\xi}}$$

This is simply the first-order condition evaluated at the average forecast for ν_{it} (conditional on the information available at $t-1$). The first-order terms in $\nu_{it}-\nu_t^0$ are forecasted by the agents to be 0, and therefore do not appear in this first-order condition.

Using the labor FOC evaluated at $\nu_{it}=\mathbb{F}_{t-1}[\nu_{it}]$, we see that:

$$A\left(\frac{(1-\alpha)\theta}{w}\right)e^{\mathbb{F}_{t-1}[\nu_{it}]}\left[\alpha\left(\frac{k_{it}}{l(k_{it},\mathbb{F}_{t-1}[\nu_{it}])}\right)^{\frac{\xi-1}{\xi}}+(1-\alpha)\right]^{\theta\frac{\xi}{\xi-1}-1}=l(k_{it},\mathbb{F}_{t-1}[\nu_{it}])^{1-\theta}$$

Using the previous equation for the expected capital-labor ratio, we can solve for the optimal capital choice given belief $\mathbb{F}_{t-1}[\nu_{it}]$:

$$k_{it}=Ae^{\mathbb{F}_{t-1}[\nu_{it}]}\left(\frac{\alpha\theta}{R}\right)^{\xi}\left(\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1}+(1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right)^{\frac{1-(1-\theta)\xi}{(1-\theta)(\xi-1)}}$$

When the realized TFP ν_{it} is equal to its average, we also get the optimal labor:

$$l_{it}(k_{it},\mathbb{F}_{t-1}[\nu_{it}])=Ae^{\mathbb{F}_{t-1}[\nu_{it}]}\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi}\left(\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1}+(1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right)^{\frac{1-(1-\theta)\xi}{(1-\theta)(\xi-1)}}$$

For other realizations, we use a first-order expansion:

$$l_{it}(k_{it},\nu_{it}) = l_{it}(k_{it},\mathbb{F}_{t-1}[\nu_{it}]) + (\nu_{it} - \mathbb{F}_{t-1}[\nu_{it}])\frac{\partial l}{\partial \nu}(k_{it},\mathbb{F}_{t-1}[\nu_{it}])$$

The two first-order conditions are simply:

$$\begin{cases} \theta(1-\alpha)A\mathbb{F}_{t-1}[e^{\nu_{it}}]\, l_{it}^{-\frac{1}{\xi}}\left[\alpha k_{it}^{\frac{\xi-1}{\xi}} + (1-\alpha)l_{it}^{\frac{\xi-1}{\xi}}\right]^{\frac{1-(1-\theta)\xi}{\xi-1}} = w \\ \theta\alpha A\mathbb{F}_{t-1}[e^{\nu_{it}}]\, k_{it}^{-\frac{1}{\xi}}\left[\alpha k_{it}^{\frac{\xi-1}{\xi}} + (1-\alpha)l_{it}^{\frac{\xi-1}{\xi}}\right]^{\frac{1-(1-\theta)\xi}{\xi-1}} = R \end{cases}$$

The labor / capital ratio depends only on factor prices:

$$\frac{k_{it}}{l_{it}} = \left(\frac{\alpha}{1-\alpha}\frac{w}{R}\right)^{\xi}$$

By injecting this constant ratio into the firm's FOC, we obtain:

$$\begin{cases} l_{it} = \left(\frac{(1-\alpha)\theta}{w}\right)^{\xi}(A\mathbb{F}_{t-1}[e^{\nu_{it}}])^{\frac{1}{1-\theta}}\left(\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right)^{\frac{1-(1-\theta)\xi}{(1-\theta)(\xi-1)}} \\ k_{it} = \left(\frac{\alpha\theta}{R}\right)^{\xi}(A\mathbb{F}_{t-1}[e^{\nu_{it}}])^{\frac{1}{1-\theta}}\left(\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right)^{\frac{1-(1-\theta)\xi}{(1-\theta)(\xi-1)}} \end{cases}$$

Firm i output is simply:

$$p_{it}y_{it} = Ae^{\nu_{it}}(A\mathbb{F}_{t-1}[e^{\nu_{it}}])^{\frac{\theta}{1-\theta}}\left[\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right]^{\frac{\theta}{(\xi-1)(1-\theta)}}$$

Since labor is proportional to capital, we can compute TFP residuals by remarking that:

$$\ln(p_{it}y_{it}) - \theta\ln(k_{it}) = \nu_{it} + \ln(A) + \theta\xi\ln\left(\frac{\alpha\theta}{R}\right) + \frac{\theta\xi}{\xi-1}\ln\left[\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right]$$

Note that the log-sales forecast errors is similar to the one derive in the model with time to build in labor:

$$\ln(p_{it}y_{it}) - \ln(\mathbb{F}_{t-1}[p_{it}y_{it}]) = \nu_{it} - \ln(\mathbb{F}_{t-1}[e^{\nu_{it}}]) = -(\gamma\rho\omega_{it-1} + \lambda\psi_{it}) + \omega_{it} - \frac{1}{2}\sigma_{\omega}^2$$

The variance of log-sales forecast errors in the data is therefore given by:

$$\mathbb{V}\text{ar}[\widehat{FE}_{it}] = \sigma_{\zeta}^2 + \left((1+\gamma^2\rho^2)\sigma_{\omega}^2 + \lambda^2\sigma_{\psi}^2\right)$$

A regression of reported log-sales forecast at date t on reported log-sales forecast at date $t-1$ leads to a regression coefficient κ_1:

$$\kappa_1 = -\frac{\gamma\rho\sigma_{\omega}^2}{\sigma_{\zeta}^2 + (1+\gamma^2\rho^2)\sigma_{\omega}^2 + \lambda^2\sigma_{\psi}^2}$$

A regression of date t log-productivity innovations on date-t reported log-sales forecast leads to a regression coefficient κ_2:

$$\kappa_2 = \frac{\sigma_{\omega}^2 - \lambda\sigma_{\psi}^2}{\sigma_{\zeta}^2 + (1+\gamma^2\rho^2)\sigma_{\omega}^2 + \lambda^2\sigma_{\psi}^2}$$

Finally, from the formula for log-sales forecast, note that:

$$\mathbb{F}_{it-1}\left[p_{it}y_{it}\right] = \left(\frac{\alpha\theta}{R}\right)^{-\xi}\left[\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right]k_{it}$$

Therefore, the sales to capital ratio is related to the true log-sales forecast in the following way:

$$\ln(p_{it}y_{it}) - \ln(k_{it}) = \ln(p_{it}y_{it}) - \ln(\mathbb{F}_{it-1}\left[p_{it}y_{it}\right]) + \xi\ln\left(\frac{\alpha\theta}{R}\right) - \ln\left[\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right]$$

Since we observe only reported log-sales forecast:

$$\ln(p_{it}y_{it}) - \ln(k_{it}) = \widehat{FE}_{it} + \zeta_{it} + \xi\ln\left(\frac{\alpha\theta}{R}\right) - \ln\left[\alpha\left(\frac{\alpha\theta}{R}\right)^{\xi-1} + (1-\alpha)\left(\frac{(1-\alpha)\theta}{w}\right)^{\xi-1}\right]$$

As a result, a regression of the log-sales-to-capital ratio on the reported log-sales forecast error should have a coefficient of:

$$\hat{\kappa}_3 = 1 - \frac{\sigma_\zeta^2}{\mathbb{V}\text{ar}\left[\widehat{FE}_{it}\right]}$$

We can therefore estimate this model exactly the same way we estimated the model with time-to-build in labor.

A.4 Proof of Proposition 2

We first take capital as given, and maximize profit with respect to labor given wage. We obtain:

$$\pi_{it} = (1-\theta(1-\alpha))Y^{1-\phi}e^{\frac{\phi z_{it}}{\alpha}}k^{\phi}\left(\frac{(1-\alpha)\theta}{w}\right)^{\frac{1-\alpha}{\alpha}\phi}$$

where $\phi = \frac{\theta\alpha}{1-\theta(1-\alpha)}$.

We then take the forecast F of the above expression, and maximize it with respect to capital k_{it}. We obtain the following formula for the revenue productivity of capital:

$$\alpha\theta\frac{p_{it}y_{it}}{k_{it}} = (r+\delta)\underbrace{\frac{e^{\frac{\Phi z_{it}}{\alpha}}}{F_{t-1}e^{\frac{\Phi z_{it}}{\alpha}}}}_{\equiv 1+\tau_{it}}$$

Time to build acts like a wedge τ_{it} between the effective cost of capital and the frictionless cost of capital. This wedge has a rational and bias component. Given that the mean of z is zero and that the innovation on z is $\ll 1$, we rewrite the log wedge as:

$$\ln(1+\tau_{it}) = \frac{\Phi}{\alpha}z_{it} - \ln\left(\mathbb{F}_{t-1}\left[e^{\frac{\Phi}{\alpha}z_{it}}\right]\right) = \ln(p_{it}y_{it}) - \ln\left(\mathbb{F}_{t-1}[p_{it}y_{it}]\right) = FE_{it}$$

In other words, the log-sales forecast error acts as a capital wedge for the firm. Based on this observation, we can use the formula in Sraer and Thesmar (2018) to calculate log TFP when the log-sales forecast is log-normally distributed (so that productivity and the log-sales forecast errors are jointly log-normally distributed), or alternatively, when variations in the log-sales forecast errors and productivity are small around their respective mean, so that we can consider a 2nd order Taylor expansion around these means. This directly provides the formula in Proposition 2.

A.5 Proof of Corollary 1

The variance of forecast errors when managers have perfect foresight is 0, which proves the first part of the Corollary.
The variance of forecast errors when managers have rational expectation is ≥ 0. Additionally, reported log-sales forecast error have larger variance than actual log-sales forecast errors: $\mathbb{V}\text{ar}\left[FE_{it}\right] \leq \mathbb{V}\text{ar}\left[\widehat{FE}_{it}\right]$. This proves the second part of the Corollary.

A.6 Proof of Corollary 2

When managers form expectation using Equation 4, their log-sales forecasts is given by: $\ln(\mathbb{F}_{t-1}\left[p_{it}y_{it}\right]) = C + \frac{\theta}{1-\theta}\left(\rho(z_{it-1} + \gamma\epsilon_{it-1}) + (1+\lambda)\eta_{it}\right)$, where C is constant across firms (see Appendix A.4 with $\nu = \theta z$, $\omega = \theta\epsilon$ and $\psi = \theta\eta$). Therefore, since z, ϵ and η are log-normally distributed, the log-sales forecast are log-normally distributed, and Proposition 2 applies. The variance of log-sales forecast errors is simply: $\mathbb{V}\text{ar}\left[FE_{it}\right] = \left(\frac{\Phi}{\alpha}\right)^2\left((1+\gamma^2\rho^2)\sigma_\epsilon^2 + \lambda^2\sigma_\eta^2\right)$ (see Appendix A.4). Therefore, TFP in the actual economy is: $\ln(TFP) = -\frac{1}{2}\left(\frac{\Phi}{1-\Phi}\right)\left(\frac{\Phi}{\alpha}\right)\left((1+\gamma^2\rho^2)\sigma_\epsilon^2 + \lambda^2\sigma_\eta^2\right)$. When managers have rational forecasts, $\gamma = \lambda = 0$ and $\ln(TFP^{\text{rational forecasts}}) = -\frac{1}{2}\left(\frac{\Phi}{1-\Phi}\right)\left(\frac{\Phi}{\alpha}\right)\sigma_\epsilon^2$. The difference between these two expression is the formula in the corollary.

B Appendix Tables

Table B.1: Forecast Errors on Forecast Revisions by Firm Size (Italian Data)

	Forecast error	
	(1)	(2)
Panel A: Firms with 20-49 employees		
Forecast error(t-1)	0.289***	0.143***
	(0.026)	(0.023)
Observations	10,712	5,537
Adj R^2	0.09	
Panel B: Firms with 50-99 employees		
Forecast error(t-1)	0.302***	0.153***
	(0.025)	(0.027)
Observations	7,106	4,142
Adj R^2	0.11	
Panel C: Firms with 100-199 employees		
Forecast error(t-1)	0.345***	0.211***
	(0.027)	(0.033)
Observations	5,739	3,486
Adj R^2	0.15	
Panel D: Firms with 200-500 employees		
Forecast error(t-1)	0.360***	0.192***
	(0.031)	(0.036)
Observations	4,968	2,995
Adj R^2	0.16	
Panel E: Firms with 501- employees		
Forecast error(t-1)	0.421***	0.180***
	(0.050)	(0.044)
Observations	3,858	2,468
Adj R^2	0.20	
Firm FE	No	Yes
Year FE	Yes	Yes

Note: This Table only use Italian data. It reproduces the regressions of Table 4, panel A, for different firm size groups. Each panel corresponds to regression results for one size group. ***, ** and * means statistically significant at the 1%, 5% and 10% confidence level.

Table B.2: Forecast Errors on Forecast Revisions (U.S. Data only)

	Forecast error	
	(1)	(2)
	Manager	Analyst
Forecast revision	0.33***	0.35***
	(0.04)	(0.035)
Constant	0.007**	-0.005*
	(0.0028)	(0.003)
Year FE	Yes	Yes
Obs	8,949	7,673
R^2	0.021	0.034

Note: This Table only use U.S. data, since Italian data do not offer multi-horizon forecasts. We regress here year t log-sales forecast error on the change in the log sales forecast since the last forecast for fiscal year t sales (forecast revision). Unfortunately U.S. guidance data – as opposed to analyst data – only rarely provide multiyear sales forecasts, so we go around this limitation by running the following regression:

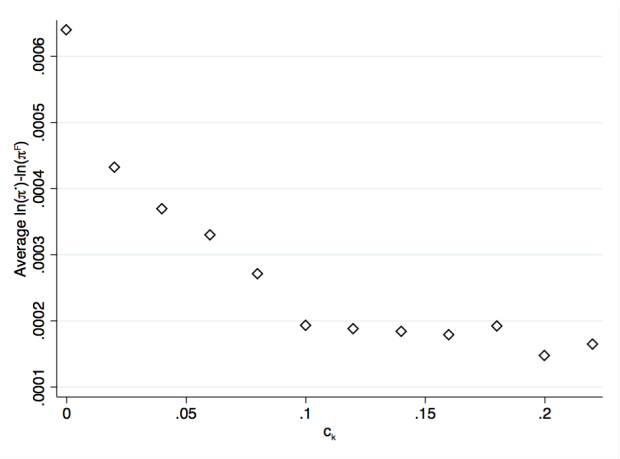
$$\log \text{sales}_{it} - \log F_{t-h}\text{sales}_{it} = cst + \gamma \left(\log F_{t-h} \log \text{sales}_{it} - \log F_{t-k}\text{sales}_{it}\right) + \epsilon_{iht}$$

where we lump together different horizons of revisions h (as long as h is no less than 2 quarters), and k is the time when the previous forecast of sales_{it} is made. Standard errors are clustered by both firm and time. ***, ** and * means statistically significant at the 1%, 5% and 10% confidence level.

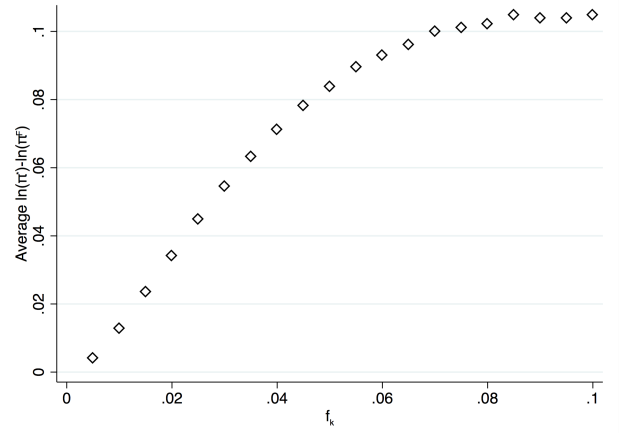
C Appendix Tables

Figure C.1: Average Profit Increase from Rationality: Adjutsment Costs (U.S. data.)

Panel A: Varying quadratic adjustment costs

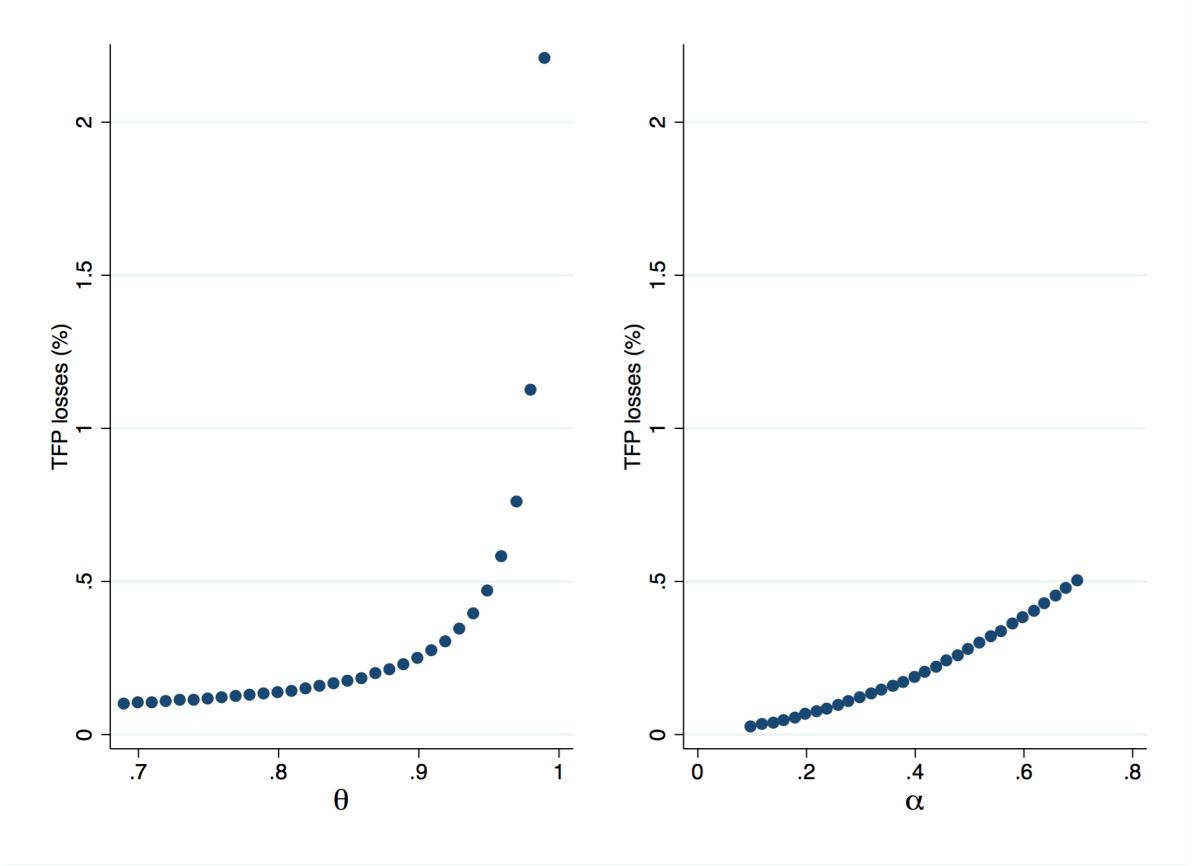


Panel B: Varying fixed adjustment costs



Note: This figure combines the estimation of Table 5, various degree of adjustment costs to simulate data for managers using rational forecast or distorted forecasts. It then computes, for each realization of state variables in the Bellman problem (6), the log-difference between rational and distorted profits. The figure shows the average log-difference in these profits for the different values of quadratic adjustment costs c_k in Panel A, and for different values of fixed costs f_k in Panel B.

Figure C.2: Sensitivity analysis: θ and α (U.S. data)



Note: This figure computes the TFP losses from distorted forecasts for different calibration of α and θ. On the left panel, we fix $\alpha = 1/3$ and let θ vary from .69 to .99. On the right panel, we fix $\theta = .8$ and let α vary from .1 to .7. For each calibrated value, we estimate the moments in Panel A, Table 5 and then estimate the model following the methodology developed in Section 3.2. We then compute the TFP losses using Equation 9.