A New Look at Uncertainty Shocks: Imperfect Information and Misallocation

Tatsuro Senga*

November 2016

Abstract

Uncertainty faced by individual firms appears to be heterogeneous. In this paper, I construct a new set of empirical measures of firm-level uncertainty using data such as the IBES and Compustat. The panel data that I construct reveals persistent differences in the degree of uncertainty facing individual firms not reflected by existing measures. Consistent with existing measures, I find that the average level of uncertainty across firms is countercyclical, and that it rose sharply at the start of the Great Recession. I next develop a heterogeneous firm model that embeds Jovanovics (1982) model of learning. Each firm gradually learns about its own productivity, and each occasionally experiences a shock forcing it to start learning afresh. In the model, uncertainty will be resolved gradually as firms operate longer and get better informed. The model can capture the cross-sectional and cyclical features of firm-level uncertainty well. When calibrated to reproduce the level and cyclicality of the measure of firm-level uncertainty, I show that an uncertainty shock explains 28 percent of the observed decline in GDP and 31 percent of the fall in investment during the Great Recession.

JEL: E22, E32, D8, D92

Keywords: uncertainty, learning, misallocation and business cycles.

---

*Queen Mary University of London (E-mail: t.senga@qmul.ac.uk). I am grateful to Aubhik Khan and Julia Thomas for the valuable advice they have provided throughout this project. I would also like to thank Nick Bloom for his advice and support. I thank Julieta Caunedo, Mitsuru Katagiri and Thomas Schelkle for their helpful discussion at conferences. I have benefited from Cheng Chen, Dean Corbae, Thorsten Drautzburg, Francois Gourio, Hugo Hopenhayn, Cosmin Ilut, Nir Jaimovich, Boyan Jovanovic, Nobuhiro Kiyotaki, Pok-sang Lam, Kehrig Matthias, Alex Monge-Naranjo, Hikaru Saijo, Yongseok Shin, Stephen Terry, Joseph Vavra, Venky Venkateswaran. I am grateful to participants at seminars at the Bank of Canada, Banque de France, Board of Governors, Bonn, CEMFI, EIEF, Florida State, Hitotsubashi, Hong Kong, Kent, Kyoto, Ohio State, Philadelphia Fed, Queen Mary, Stanford, Southern California, USC Marshall, Tinbergen Institute, Virginia and numerous conferences, for comments and suggestions. All remaining errors are mine.

1 Introduction

> "*Subjective uncertainty is about the "unknown unknowns". When, as today, the unknown unknowns dominate, and the economic environment is so complex as to appear nearly incomprehensible, the result is extreme prudence, [. . .], on the part of investors, consumers and firms.*" Olivier Blanchard (2009)

How large is the role of increased uncertainty in driving economic downturns? Is there a link between a rise in firm-level uncertainty and the subsequent pace of economic recovery? To explore these questions, I construct new empirical measures of firm-level uncertainty, and I show that the degree of uncertainty varies across firms and the average level of uncertainty, as well as its dispersion, across firms is countercyclical. To account for these regularities, I develop a heterogeneous firm model that incorporates learning at the firm level with uncertainty shocks. The model can explain my empirical findings well for both cross-sectional and cyclical features. The cross-sectional distribution of firm performance as well as firm-level uncertainty are well matched with data. In a recession, the distribution of firm-level uncertainty is skewed by having a large mass of firms with high uncertainty, and the average level of uncertainty across firms is also high. In addition, the model successfully reproduces a gradual recovery of the aggregate economy following uncertainty shocks.

A defining feature of this paper is that the uncertainty faced by firms not only varies over time but also varies across firms. One common approach in the uncertainty shock literature, following the seminal work of Bloom (2009), has been to study stochastic volatility models. I break with this tradition primarily because stochastic volatility models cannot deliver the heterogeneous uncertainty evident in microdata.[1] I integrate Jovanovic's (1982) model of learning into an otherwise standard heterogeneous firm business cycle framework. In this model, by contrast, uncertainty, defined as the conditional variance of forecasts of firm performance, varies across firms depending on the information each firm possesses. Firms are heterogeneous in both productivity and their confidence about that productivity; more informed firms have lower posterior variances of their beliefs. Two different firms can have the same posterior mean while differing in their posterior variances. Hence, uncertainty differs across firms. A second appealing feature of the model is the fact that the recession in response to an uncertainty shock is not followed by a sharp recovery, as happens in existing stochastic-volatility-based uncertainty shock

[1] In a common stochastic volatility approach, there is full information and all agents know the true distribution of shocks that they face, including its volatility, which varies over time. In uncertain times, the volatility that every agent faces rises equally. See, for example, Vavra (2014), Bloom et al. (2014) and Bachmann and Bayer (2013).

models.[2] Instead, my model with a non-trivial distribution of firms with learning drives a slow economic recovery as firms gradually regain information and confidence. Moreover, these results require no additional rigidity or frictions. In the absence of labor and capital adjustment costs, uncertainty shocks still cause recessions.[3]

I construct a new panel dataset of firm-level uncertainty based on data from the Institutional Brokers' Estimate System (I/B/E/S), Center for Research in Securities Prices (CRSP) and Compustat databases. By merging these data, I construct an annual panel of US public firms with uncertainty measures such as an ex-ante earnings forecast dispersion among market analysts, ex-post-realized forecast errors and stock price volatility measures. Appealing features of the dataset, particularly with regard to the inclusion of earnings forecast data, include the following: (1) it is disaggregated at the firm level, thereby allowing the examination of the cross-sectional characteristics of firm-level uncertainty, (2) it contains ex-ante information on firm profitability, which is arguably better suited than ex-post information for gauging the degree of uncertainty individual firms face and (3) the result obtained can be fairly directly mapped into the model that I build. In particular, I transform earnings data into the return on assets (ROA) data and use the latter to discipline the model.

The firm-level measures of uncertainty uncover the following new facts. First, the degree of uncertainty facing individual firms differs across firms; for example, Apple's measure of uncertainty was much lower than Ford's during the Great Recession in 2009, and vice versa during the dot-com recession in 2001. Second, the first and second moments of the distribution of firm-level uncertainty measures are countercyclical. Specifically, the median, mean and cross-sectional dispersion are all negatively correlated with real GDP series.

In light of the evidence above, I propose a new model that features heterogeneous uncertainty, and I study its role in propagating aggregate shocks. My model builds on a standard heterogeneous firm business cycle model, but I deviate from the standard model in three ways. First, idiosyncratic productivity has two components: an i.i.d. transitory component around a base component. These components cannot be observed separately, and therefore each firm must learn the true value of its base component in a Bayesian way.[4] Second, the base component is randomly reset. When this occurs, a new base

[2] See, for the discussion, Bachmann, Elstener and Sims (2013) and Bachmann and Bayer (2013).

[3] The large body of literature about the relation between uncertainty and investment studies the real options effect in models with adjustment costs, as in Bertola and Caballero (1994), Dixit and Pindyck (1994), Abel and Eberly (1996) and Caballero and Engel (1999).

[4] Bernanke (1983) develops a single-firm, partial equilibrium model with dynamic Bayesian inference specifications to study short-term fluctuations of irreversible investment under time-varying option values.

productivity is drawn from a common distribution known to all firms. Though a firm knows when its base has been reset, it does not know its new realization; thus, it must restart the process of learning its value. Otherwise, the firm maintains its current base component and continues its learning. In this way, I integrate learning into a model of heterogeneous firms that are subject to persistent shocks to idiosyncratic productivity, as in Hopenhayn (1992). Third, I assume that the common reset probability of firms' base productivity components is a stochastic, two-state Markov process. When the reset probability is high, many firms draw new base productivities, which leads to a larger variance of TFP growth rates across the distribution of firms. Thus, an uncertainty shock is associated with a rise in the variance of firm-level TFP growth rates. In this way, this model is consistent with an important empirical observation documented in previous work regarding uncertainty shocks (Bloom et al., 2014). The rise in the reset probability also implies that many firms lose information and restart learning. This additional effect increases a mass of firms with high uncertainty, leading to an increase in the average level of uncertainty as well as dispersion across firms.

My main findings are as follows. First, the model produces rapid downturns and slow recoveries in aggregate variables following uncertainty shocks. Second, aggregate productivity shocks deliver responses quite similar to those in conventional equilibrium business cycle models, and these shocks remain an important source of fluctuations in my model. For this reason, the model delivers familiar second moments for the cyclical component of aggregate quantities, as has been shown in the literature of the business cycle (e.g., Cooley and Prescott, 1995).

In my model, the recession following an uncertainty shock stems from two effects, one uncertainty and one distributional. The uncertainty effect arises as all firms anticipate a higher reset probability, implying an increased likelihood of large changes in their productivities. Given one-period time to build for capital, this leads them to change their target levels of capital. Firms that believe their current base component is higher than the unconditional mean reduce their capital targets. On the other hand, firms that believe their current base productivity is low relative to the unconditional mean raise their capital targets. Given that the distribution of the base component of productivity is symmetric and the production function has decreasing returns to scale, the net impact on aggregate investment is negative. This effect is immediately reversed when the shock ends, which on its own would deliver a quick expansion led by pent-up investment demand. However, this fails to spur aggregate investment because of the offsetting impact of the distributional effect. As an unusually large number of firms experience a reset of their base components, the economy becomes increasingly populated by uninformed firms as the uncertainty shock persists. These firms, having lost their information, must restart their learning. In early stages of the learning process, each firm puts more weight on the prior mean rather than the mean of its observations. Unless a firm is fully informed about its base component, its capital stock is either excessive or insufficient relative to the efficient level of capital stock based on full information about the true value of its base component and the interest rate. Thus, there is a misallocation of resources arising from over- or undercapacity. In particular, since uninformed firms are cautious and have low confidence, while their population share rises over the course of an uncertainty driven recession, aggregate investment, employment and GDP fall. This cannot be quickly reversed when the uncertainty shock ends, as it takes time for the distribution of firms to recover their knowledge about their productivity. Thus, the negative impact of an uncertainty shock persists beyond the shock itself.

At the start of a recession, the uncertainty and distributional effects reinforce each other, and this leads to a rapid drop in investment, GDP and other aggregate variables. However, in the recovery phase, the uncertainty and distributional effects offset each other. Their relative strengths must be quantitatively assessed. In my calibrated model, the distributional effect dominates. This leads to a sluggish recovery, a finding that stands in sharp contrast to other models in the literature.

Related Literature The idea that links uncertainty to business cycles and especially to the slow rate of recovery after slumps dates back to Keynes (1936) and was further formulated by Bernanke (1983) in his study of investment fluctuations.[5] In the recent equilibrium business cycle literature, the seminal contribution of Bloom (2009) studies a business cycle model in which individual firms face time-varying volatility shocks to their own productivities. He shows that uncertainty shocks, defined as a shock to the variance of the idiosyncratic productivity process, generate bust-boom cycles. A rise in stochastic volatility, in a setting where firms face nonlinear costs of factor adjustment, deters investment as firms adopt a "wait and see" policy in response to the shock. In this class of models with exogenous shocks to volatility, the aggregate effects tend to be short-lived. However, Bachmann et al. (2013) argue that the quick recovery following the wait-and-see effect is not consistent with U.S. data. In particular, they document persistent and prolonged dynamics following a rise in their measure of uncertainty. I

[5] As stated in The General Theory, Ch. 22, *"it might be possible to achieve a recovery without the elapse of any considerable interval of time [. . .]. But, in fact, this is not usually the case [. . .]. It is the return of confidence, to speak in ordinary language, which is so insusceptible to control in an economy of individualistic capitalism. This is the aspect of the slump which bankers and business men have been right in emphasizing..."*

contribute to this literature by developing a tight link between uncertainty at the start of a recession and the gradualism of the subsequent recovery.

My paper contributes to the literature that examines firm-level uncertainty using microdata and its aggregate implications by simulating a quantitative model. For example, Vavra (2014) showed that the real effect of a monetary shock decreases as the average level of uncertainty across firms increases in a price-setting model with stochastic volatility in firm-level productivity. Baley and Blanco (2016) adapt a Bayesian approach as in my paper for a price-setting model and show that the dispersion of firm-level uncertainty matters for the real effect of a monetary shock. The direction of these studies is also shared by Ilut and Saijo (2016), who investigate a business cycle amplification mechanism with ambiguity. My paper differs as it studies the dynamics of capital misallocation due to time-varying uncertainty.

Orlik and Veldkamp (2014) study an alternative origin of uncertainty fluctuations in a model of Bayesian learning. Uncertainty is associated with doubt about the true model of the economy. In particular, they argue that small increases in the awareness of tail risk is important in driving fluctuations of uncertainty.[6] In contrast, the distribution of outcomes is known to firms in my paper. What is unknown is their own actual realizations of outcomes.

In recent years, interest in uncertainty and learning over the business cycle has increased.[7] For example, Fajgelbaum et al. (2013) show a mechanism by which recessions increase uncertainty in a model of irreversible investment. Saijo (2014) builds a model with nominal rigidities and proposes a mechanism for endogenous fluctuations in uncertainty. Both papers analyze fluctuations in the amount of information available to agents. In recessions, economic activity contracts, and this reduces the flow of information and increases uncertainty. Neither this feedback nor real and nominal rigidities are necessary in my model for uncertainty shocks to produce recessions. Furthermore, unlike these papers, my model has time-varying distribution of firms, which is part of the aggregate state. Following uncertainty shocks, it delivers endogenous fluctuations in TFP through changes in the degree of misallocation of capital and labor, leading to a sluggish economic recovery in the presence of learning.

This work is also related to existing papers that study the role of the allocation of resources across heterogeneous agents and its impact on aggregate productivity (e.g., Restuccia and Rogerson, 2008). Hsieh and Klenow (2009) argue that misallocation of

[6] See also Kozlowski, Veldkamp and Venkateswaran (2016).

[7] There are papers that examine economic environments wherein agents learn from market outcomes. For example, Van Nieuwerburgh and Veldkamp (2006) and Caplin and Leahy (1993) study the relation between the flow of information and economic activity in models without uncertainty shocks.

resources has a substantial impact on aggregate TFP in India and China. More recently, the role of financial frictions in generating capital misallocation and its aggregate implications have been studied in several quantitative environments (Khan and Thomas, 2013; Buera and Moll, 2013; Buera et al., 2011). Instead of financial frictions, I study the role of information frictions in causing a loss in aggregate productivity through the misallocation of resources. David et al. (2013) also study misallocation in a model of learning at the firm level. However, my paper looks at the implications of misallocation over business cycles, while they focus on a stationary equilibrium.

I also contribute to the empirical literature on uncertainty. Several proxies have been developed within the literature, ranging from the volatility of GDP or stock prices to disagreement and forecast errors in survey data, as uncertainty is difficult to identify. For example, Leahy and Whited (1996) construct a measure of uncertainty from the volatility of stock returns for individual firms. Guiso and Parigi (1999) use survey data on demand forecasts by Italian firms to infer the level of uncertainty facing these individual firms. Bond et al. (2005) consider several measures including volatility in monthly consensus earnings forecasts, the variance of forecast errors for consensus forecasts and the dispersion in earnings forecasts across market analysts. To estimate the impact of uncertainty on investment, they use panel data and look at the cross-sectional features of firm uncertainty and the investment behavior of individual firms, rather than the uncertainty distribution's cyclical properties as in Bloom et al. (2014), Kehrig (2016) and Vavra (2014). In this paper, I use data on earnings forecasts by individual analysts as in Bond (2005); however, I examine not only the average cross-sectional distribution but also the cyclical changes of this uncertainty measure. Bachmann et al. (2013) use survey data from the IFO Business Climate Survey, which asks forecasters about their own future prospects rather than about macroeconomic variables such as GDP, to extensively study various measures of uncertainty. I also use forecast disagreement to measure uncertainty.

My model builds on Jovanovic's (1982) learning model, which has been applied to study a broad range of topics such as the disparate response of heterogeneous firms to aggregate shocks Li and Weinberg, 2003; Alti, 2003) and the differential sensitivity of product switching behavior among exporters learning about their demand (Timoshenko, 2013).

The rest of the paper is organized as follows. Section 2 reports empirical results. In Section 3, the model of heterogeneous firms with learning is developed. Section 4 presents my quantitative results, both stationary equilibrium results matched against a variety of micro-level moments and the business cycle results in the presence of aggregate uncertainty. Section 5 concludes.

2 Empirics

In this section, I first build an annual panel dataset of firms' ex-ante earnings forecast dispersion and ex-post forecast errors using data from the I/B/E/S and Compustat. I then use this panel to construct new empirical measures of firm-level uncertainty. These new measures reveal persistent differences in the degree of uncertainty facing individual firms. Consistent with existing measures, these new measures show that the average level of uncertainty across firms is countercyclical. In particular, there was a sharp rise at the start of the Great Recession.

2.1 Data

Among many papers that use survey data to measure uncertainty, I follow the literature in using earnings forecasts to build a proxy for firm-level uncertainty, as in Johnson (2004), Bond at al. (2005), and Janunts (2010).[8]

The first data source that I use is the I/B/E/S, which contains a point forecast of earnings per share (EPS) made by an individual analyst. For each firm at each month, a researcher can calculate the cross-analyst dispersion of earnings forecasts. The I/B/E/S also contains actual earnings records.[9] By comparing earnings forecasts with the actual earnings records, one can also calculate forecast errors for each firm at each month.

As the second data source, I use Compustat, from which I take accounting data from the balance sheet, profit and loss, and cash flow statements. By adding each firm's accounting fundamentals data, I can implement a necessary data transformation that allows me to tightly map earnings forecast data into the model-generated moments. In particular, by using total assets from Compustat, I transform both earnings forecasts and records, which are in dollars, into ROA expressed as a percentage.

The result is a panel containing forecast dispersion and forecast errors about ROA for individual firms. One of the attractive features is that I can elicit both an ex-ante and ex-post uncertainty measure at the firm-level. Another important attractive feature is that earnings data coupled with total assets data allows me to use the panel data for my

[8] To proxy uncertainty about macroeconomic variables, surveys such as the Survey of Professional Forecasters (SPF) have been widely used (see, for example, Zarnowitz and Lambros, 1987; Giordani and Sderlind, 2003; and Rich and Tracy, 2010). For example, Zarnowitz and Lambros (1987) show a positive relationship between forecast dispersion and uncertainty, while Rich and Tracy (2010) find little evidence in support of using disagreement to measure uncertainty.

[9] Earnings that can be obtained from the I/B/E/S are so-called street earnings, which are different from earnings that can be obtained from Compustat using the generally accepted accounting principles (GAAP).

quantitative analysis. To the best of my knowledge, this paper is the first to achieve a tight mapping of firm-level uncertainty measures into a quantitative business cycle model with heterogeneous firms.

These features are distinct among other commonly used measures of firm-level uncertainty such as realized stock-return volatility or implied volatility.[10] For instance, stock-return volatility may reflect all possible sources of risk rather than the firm's fundamentals, which I intend to isolate in this paper. In addition to earnings forecast data, I have also used stock returns data from the CRSP database. The construction of uncertainty measures are explained in the next section.

2.2 Firm Uncertainty Panel

For firm i in each month m during year t, I can observe analyst j's point forecast of EPS for this current year t, denoted $feps_{i,j,m,t}$. Earnings forecasts are transformed into ROA forecasts, $froa_{i,j,m,t}$, by using data on the number of outstanding shares during year t, $sharenum_{i,t}$, and total assets data at the end of the previous year $(t-1)$, $at_{i,t-1}$:

$$froa_{i,j,m,t} = \frac{feps_{i,j,m,t} * sharenum_{i,t}}{at_{i,t-1}}. \tag{1}$$

Then, I define the forecast dispersion-based uncertainty measure, $fdisp_{i,m,t}$, defined as the coefficient of variation of $froa_{i,j,m,t}$.[11]

For firm i in each month m during year t, I can obtain the median forecast among analysts, $fmedian_{i,m,t}$. With the realized ROA, $roa_{i,t}$, I define the forecast error-based uncertainty measure, $ferror_{i,m,t}$, as the absolute value difference between $roa_{i,t}$ and $fmedian_{i,m,t}$. If the median ROA forecast is 8 percent and the actual ROA turns out to be 9 percent, then the forecast error is 1 percent.

So far, these uncertainty measures are monthly based and thus I collapse this monthly panel to a yearly panel in order to combine it with other yearly-based items, such as sales. To this end, I focus on forecasts made 8 months before the end of each firm's accounting year. This allows me to match the annual model frequency and ensures that I have enough forecasts as many analysts report their first forecast around this time.[12] This results in

[10] See, for example, Leahy and Whited (1996) and Bloom, Bond and Van Reenen (2007) for a stock-returns-based uncertainty measure. Another measure of uncertainty recently used in the literature relies on options price data as in Stein and Stone (2013).

[11] An earning per share (EPS) cross-analyst disagreement measure requires more than one analyst coverage; thus data with only one forecast is dropped.

[12] For example, for a firm with its fiscal year ending at March, the number of forecasts starts to increase in February and by April it reaches the number close to the final figure. The results are not qualitatively

$fdisp_{i,t}$ and $ferror_{i,t}$ for firm i in year t.

I also use stock return data to construct a measure of realized stock price volatility on a yearly basis. Specifically, I take daily stock returns data from the Center for Research in Securities Prices (CRSP) database and calculate the annualized volatility measure of stock returns.

Linking the Compustat, I/B/E/S and CRSP databases into an annual firm-by-year panel, the resulting dataset is an unbalanced panel of 46,271 data for 6,453 firms with 8 data points on average, spanning from 1977 to 2014.[13] This panel contains uncertainty measures, performance measures such as ROA and various firm characteristics, including size, age and analyst coverage. Table 1 reports descriptive statistics of the panel.

2.3 Cross-Sectional Properties: Uncertainty Varies Across Firms

From Table 1, it is evident that there is heterogeneity across firms. For example, firm size (proxied by sales, total assets and the number of employees), age and performance (e.g., ROA) exhibit a substantial heterogeneity. Furthermore, the distribution of forecast dispersion and forecast errors are highly skewed to the right and there is a large variability across firms. The forecast error for 5th percentile firm is about zero, while that of 95th percentile firm is more than 11 percent.

Table 2 shows the sample mean of key variables for the subsamples of firms with low and high uncertainty. Firms with low uncertainty tend to be larger in size (sales, total assets or the number of employees), to be older and more likely to survive longer and to have greater analyst coverage (measured by the number of analysts who report forecasts).

Uncertainty varies across not only firms but also time. Figure 1 plots how the distribution of uncertainty measures moved during the Great Recession. In the next section, I construct time series indexes from firm-level uncertainty measures in this panel data and show the cyclical properties of these measures.

2.4 Cyclical Properties: Uncertainty Varies Across Time

Using the panel data, I define the following variables for each year t: (1) $fdisp\_mean_t$ as the mean of $fdisp_{it}$; (2) $ferror\_mean_t$ as the mean of $ferror_{it}$; (3) $fdisp\_sd_t$ as the standard deviation of $fdisp_{it}$; (4) $ferror\_sd_t$ as the standard deviation of $ferror_{it}$;

affected even if I use alternative options to collapse a monthly panel to a yearly panel, either by using a different month or taking a yearly average.

[13] I exclude financial and utilities firms from the sample by dropping firms with Standard Industrial Classification (SIC) between 4900 and 4999 and then between 6000 and 6999.

Table 1: Panel: Descriptive Statistics

	mean	sd	p5	p25	p50	p75	p95
Sales	2945.8	7543.4	33.9	185.4	568.8	1913.9	14224.0
Total assets	3152.2	8256.5	50.8	189.1	563.3	1984.0	15444.1
Employment	15.0	51.1	0	1	3	11	62
Age	9.3	7.6	2	4	7	13	26
Life	16.6	10.5	3	8	15	23	38
Analyst coverage	9.3	7.2	2	4	7	13	24
Leverage Ratio	0.216	0.193	0.000	0.041	0.191	0.325	0.580
ROA	0.030	0.111	-0.148	0.010	0.038	0.075	0.161
fdisp	0.193	0.449	0.007	0.027	0.063	0.156	0.769
ferror	0.026	0.049	0.000	0.003	0.009	0.026	0.110
vol	0.432	0.219	0.182	0.275	0.377	0.532	0.873

Note: The table above shows the cross-sectional moments of the firm-by-year panel. The panel data is constructed by merging data from Compustat, CRSP, and I/B/E/S, resulting in an unbalanced panel of $6,453$ firms between 1977 and 2014, consisting of $46,271$ firm-year observations. Sales and total assets are in millions of dollars. Age is the number of years calculated from the first year of observation. Life is the number of years during which observations can be found. Analyst coverage is the number of analysts who reported earnings forecasts. ROA is calculated as earnings (= street earnings per share (EPS) multiplied by the number of outstanding shares) divided by total assets. Forecast dispersion ($=fdisp$) is the coefficient of variation of ROA forecasts across analysts. Forecast error ($=ferror$) is calculated as the absolute value of a percentage gap between the realized ROA and the median ROA forecast. *vol* is the annualized stock returns volatility constructed using daily data. All data are winsorized at the 1 percent level.

Table 2: Subsamples: Descriptive Statistics

	Low uncertainty firms		High uncertainty firms	
	mean	sd	mean	sd
Sales	3580.0	8314.9	1069.7	4000.3
Total assets	3815.0	9092.8	1191.4	4484.2
Employment	18.3	58.0	5.0	15.9
Age	10.1	7.9	7.1	5.9
Life	18.4	11.3	7.8	4
Analyst coverage	10.1	7.5	7.2	5.8
Leverage ratio	0.222	0.183	0.198	0.218
Numbers of firms	34581		11690	

Note: The table above shows the mean and standard deviation of the subsamples of the firm-by-year panel. The samples are split into each for firms with low and high uncertainty, being categorized into each sub-sample whether $fdisp$ is above or below the mean value for the full sample. Explanations of each variable in the above table can be found in the notes of Table 1.

(5) roa_mean_t as the mean of roa_{it}; (6) roa_sd_t as the standard deviation of roa_{it}; (7) vol_mean_t as the mean of vol_{it}; and (8) vol_sd_t as the standard deviation of vol_{it}. Further, I take (9) BOS_t as the forecast disagreement index from Bachmann et al. (2013). I also take (10) TFP_{it} as the cross-sectional standard deviation of TFP shocks among U.S. plants from Bloom et al. (2014).

Table 3 shows a correlation matrix between the uncertainty measures from both my panel data and other studies together with real GDP series. First, earnings forecast-based uncertainty measures ($fdisp\_mean$, $fdisp\_sd$, $ferror\_mean$, $ferror\_sd$) are negatively correlated with real GDP series, ranging between -0.528 and 0.043. Second, the correlation between forecast dispersion-based measures ($fdisp\_mean$ and $fdisp\_sd$) and forecast error-based measures ($ferror\_mean$ and $ferror\_sd$) is positive, ranging from 0.038 to 0.468, consistent with Bachmann et al. (2013). Third, earnings forecast-based uncertainty measures are also positively correlated with stock-return-based measures (vol_mean and vol_sd) with correlations between -0.028 and 0.604. Finally, earnings forecast-based uncertainty measures are positively correlated with the TFP dispersion based-measure (TFP) developed in Bloom et al. (2014) with correlations between 0.101 and 0.639. Figure 2 plots time series of some of these variables to highlight the patterns discussed here.

To sum up, by merging the I/B/E/S data and the Compustat data, I construct an annual panel of firms' ex-ante earnings forecast dispersion among market analysts and ex-post realized forecast errors. I then document the following stylized facts. First, the degree of uncertainty facing individual firms is heterogeneous; the level of uncertainty decreases as firms age and grow. Second, the first and second moments of the firm-level uncertainty distribution are countercyclical; the correlation with GDP growth rates are negative. Thus, forecast dispersion and its variance both fall with higher GDP growth. Finally, my measures of uncertainty are positively correlated with other common measures in the literature, including stock price volatility-based and balance sheet-based measures. In the following sections, I explore these cross-sectional and cyclical features of firm-level uncertainty in a model with Jovanovic's (1982) learning with uncertainty shocks.

Table 3: Correlation between uncertainty measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) fdisp_mean	1										
(2) ferror_mean	0.276	1									
(3) fdisp_sd	0.650	0.272	1								
(4) ferror_sd	0.0382	0.521	0.468	1							
(5) roa_mean	-0.462	-0.311	-0.595	-0.285	1						
(6) roa_sd	-0.366	0.439	-0.0112	0.578	0.0285	1					
(7) vol_mean	-0.0280	0.604	0.334	0.466	-0.514	0.485	1				
(8) vol_sd	0.0963	0.529	0.351	0.495	-0.675	0.411	0.901	1			
(9) BOS	-0.170	0.353	-0.152	0.411	0.0135	0.302	0.266	0.356	1		
(10) TFP	0.212	0.639	0.101	0.146	-0.0764	0.298	0.360	0.239	0.164	1	
(11) GDP	-0.528	-0.481	-0.142	0.0436	0.356	0.163	-0.259	-0.322	-0.187	-0.475	1

Note: The above table is a correlation matrix for the measures of uncertainty. $fdisp\_mean$ and $ferror\_mean$ are the cross-sectional mean of $fdisp_{it}$ and $ferror_{it}$, across firms i for each year t, respectively. $fdisp\_sd$, and $ferror\_sd$ are the cross-sectional standard deviation of $fdisp_{it}$ and $ferror_{it}$, across firms i for each year t, respectively. roa_mean is the cross-sectional mean of ROA across firms for each year t. roa_sd is the cross-sectional standard deviation of ROA across firms for each year t. vol_mean is the cross-sectional mean of the annualized stock returns volatility across firms i for each year t. vol_sd is the cross-sectional standard deviation of the annualized stock returns volatility across firms i for each year t. BOS is the forecast disagreement index from Bachmann et al. (2013). TFP is the cross-sectional standard deviation of TFP shocks for the U.S. establishment from Bloom et al. (2014). GDP is the growth rate of the real GDP series. All indexes are HP-filtered using a smoothing parameter of 100.

3 Model

Below, I take a standard equilibrium business cycle model with heterogeneous firms and extend it as follows. First, I assume that firms' idiosyncratic productivity has both a base and a temporary component, and these two components cannot be observed separately. The temporary component is i.i.d. while the base component is persistent and, as such, relevant for firms' investment decisions. Firms learn about their base components over time, by observing their total productivity and updating their beliefs as in Jovanovic (1982). Second, each firm is subject to exogenous shocks to this base component. In each period, a firm retains its current base component with probability $1 - \pi$, but loses the current level and draws a new one with probability π. The new base component is drawn from a time-invariant distribution and independent of last period's productivity level. Whenever a firm draws a new base component, it must restart the learning process. Third, I assume that π is time-varying. A rise in uncertainty in this model happens when π is high, which implies that an unusually large number of firms change their productivity level and begin the process of learning anew.

3.1 Production, learning

The model economy is perfectly competitive and has an infinite horizon. There are a large number of competitive firms producing a homogeneous good. Each firm uses capital stock k, and labor n, via an increasing and concave production function,

$$y = z\varepsilon F(k,n), \tag{2}$$

where $F(k,n) = (k^{\alpha}n^{1-\alpha})^{\nu}$, with $0 < \alpha < 1$ and $0 < \nu < 1$.

There are two productivity terms in the production function, one aggregate, z, and one idiosyncratic, ε. z represents an exogenous stochastic total factor productivity common across all firms: $z \in \{z_1, \ldots, z_{N_z}\}$, where $\Pr(z' = z_m \mid z = z_l) \equiv \pi^z_{lm} \geq 0$, and $\sum_{m=1}^{N_z} \pi^z_{lm} = 1$ for each $l = 1, \ldots, N_z$. For the firm-specific idiosyncratic counterpart, I assume that ε is the sum of two components: a persistent one, θ, and a transitory one, a;

$$\varepsilon = \theta + a. \tag{3}$$

The base component of firm specific productivity, θ, changes infrequently and the timing of such changes, though not their value, is known to the firm. As noted above, with probability $1 - \pi$, the current base component is maintained. With probability π, the current base component is lost and a new value is drawn. This is independent of the firm's state. The transitory component, a, is independently and identically distributed over time. The distribution of both θ and a are known to all firms: $\theta \sim N(\overline{\theta}, \sigma^2_{\theta})$ and $a \sim N(0, \sigma^2_{\varepsilon})$.

Firms observe ε, but θ and a are not observed separately. Firms can extract information about their θ by accumulating observations of ε. While these observations are affected by the i.i.d. draws of a every period, repeatedly observing ε, firms learn about their θ.

We formalize this learning process as follows. Consider a firm with $\overline{\varepsilon}$ —the mean of the observations of idiosyncratic shocks ε_i for $i = 1, \ldots, t$, where t is the number of observations. To form a belief about their base component θ, $(\overline{\varepsilon}, t)$ is sufficient information. Therefore, a firm with $(\overline{\varepsilon}, t)$ infers the posterior distribution: $\theta \sim N(A, B)$ with

$$A = \frac{\sigma_a^2}{\sigma_a^2 + t\sigma_\theta^2}\overline{\theta} + \frac{t\sigma_\theta^2}{\sigma_a^2 + t\sigma_\theta^2}\overline{\varepsilon} \tag{4}$$

$$B = \frac{\sigma_a^2\sigma_\theta^2}{\sigma_a^2 + t\sigma_\theta^2} \tag{5}$$

where $\overline{\varepsilon} = (\sum_{i=1}^{t} \varepsilon_{i=1})/t$ and t is the number of observations. Each period after observing ε, the posterior distribution of θ is updated, and over time it converges to the true value of θ as t becomes large enough.

3.2 Distribution of firms

The exogenous aggregate state is summarized by $s = (z, \pi)$. In addition, a non-trivial, time-varying distribution of firms is a part of the aggregate state. As shown in the last section, firms form expectations over their productivity next period. Starting with the last period when their base component is reset, firms observe their productivity over time, and the mean of these observations and the number of observations are a part of each firm's state. This number of observations corresponds to the time-since-reset. Thus, firms at the beginning of each period are identified by the mean of their observations of idiosyncratic shocks, $\overline{\varepsilon}$, the number of these observations, t, and their current productivity draw, ε, alongside their predetermined capital stock, k. I summarize the distribution of firms over $(\overline{\varepsilon}, t, \varepsilon, k)$ using the probability measure μ defined on the Borel algebra, $\mathcal{S}$, generated by the open subsets of the product space, $\mathcal{S}= \mathbb{R}_+ \times \mathbb{Z} \times \mathbb{R}_+ \times \mathbb{R}_+$.

Given the distribution of firms, the aggregate state of the economy is fully summarized by (s, μ), and the distribution of firms evolves over time according to a mapping, Γ, from the current aggregate state; $\mu' = \Gamma(s, \mu)$.

3.3 Firm's problem

Firms solve the following problem given their firm-level state together with the aggregate state. The problem consists of choosing the capital stock for the following period, k', and the labor input for current period, n. Let $V(\overline{\varepsilon}, t, \varepsilon, k; s, \mu)$ be the value function of a firm,

$$
\begin{aligned}
V(\overline{\varepsilon}, t, \varepsilon, k; s, \mu) \quad &= \quad \max_{n,k'} \Big[z\varepsilon (k^{\alpha} n^{1-\alpha})^{\nu} - \omega n + (1-\delta)k - k' \\
& \quad\quad + (1-\pi) E_{s'|s} d\left(s', s, \mu\right) E_{\overline{\varepsilon}'|\overline{\varepsilon},t} V(\overline{\varepsilon}', t+1, \varepsilon', k'; s', \mu') \\
& \quad\quad + \pi E_{s'|s} d\left(s', s, \mu\right) E_{\varepsilon'} V(\varepsilon', 1, \varepsilon', k'; s', \mu') \Big] \qquad (6)
\end{aligned}
$$

$$
\text{subject to} \quad : \quad \overline{\varepsilon}' = \frac{t\overline{\varepsilon} + \varepsilon'}{t+1}, \qquad (7)
$$

$$
\text{and} \quad : \quad \mu' = \Gamma(s, \mu). \qquad (8)
$$

Each firm's profits are its output less wage payments and investment. With probability $1-\pi$, the current base component is maintained and hence their expectation over ε' and thus $\overline{\varepsilon}'$ are conditional on $(\overline{\varepsilon}, t)$. Furthermore, they discount next period's value by the state contingent discount factor, $d\left(s', s, \mu\right)$. With probability π, the current base component is lost and a new one is drawn, independent of the current state. In the first period after any reset of the base component, firms take an average of the mean value of $\overline{\theta}$ and the first draw of ε'. The state contingent discount factor is determined by households decision rules as explained below.

3.4 Households

There is a large number of identical households in this economy, formally a unit measure. Households choose consumption, supply labor, and hold their wealth in firm shares to maximize lifetime expected utility as follows.

$$
\begin{aligned}
V^h(\lambda; s, \mu) &= \max_{c, n^h, \lambda'} \left[U\left(c, 1-n^h\right) + \beta E_{s'|s} V^h(\lambda'; s', \mu') \right] && (9) \\
\text{subject to:} \quad & c + \int_{\mathbf{S}} \rho_1 \left(\overline{\varepsilon}', t+1, \varepsilon', k'; s, \mu\right) \lambda' \left(d \left[\overline{\varepsilon}' \times t+1 \times \varepsilon' \times k' \right] \right) \leq \\
& w(s, \mu) n^h + \int_{\mathbf{S}} \rho_0 (\overline{\varepsilon}, t, \varepsilon, k; s, \mu) \lambda (d [\overline{\varepsilon} \times t \times \varepsilon \times k]) && (10) \\
: \quad & \mu' = \Gamma(s, \mu) && (11)
\end{aligned}
$$

Households hold one-period shares in firms, which is denoted by the measure λ. Given the prices—the real wage, $w(s, \mu)$, and the prices of shares, $\rho_0(\overline{\varepsilon}, t, \varepsilon, k; s, \mu)$ and $\rho_1\left(\overline{\varepsilon}', t+1, \varepsilon', k'; s, \mu\right)$, households choose their current consumption, c, hours worked, n^h, and the numbers of new shares, $\lambda'\left(\overline{\varepsilon}' \times t+1 \times \varepsilon' \times k'\right)$.

Let $C^h(\lambda; s, \mu)$ and $N^h(\lambda; s, \mu)$ represent the household decision rules for consumption, hours worked, and let $\Lambda^h(\overline{\varepsilon}', t+1, \varepsilon', k', \lambda; s, \mu)$ be the household decision rule for shares purchased in firms that will begin the next period with $(\overline{\varepsilon}', t+1, \varepsilon', k')$.

3.5 Recursive equilibrium

A recursive competitive equilibrium is a set of functions

$$
\begin{aligned}
\text{prices} &: (\omega, d, \rho_0, \rho_1) \\
\text{quantities} &: (N, K, C, N^h, \Lambda^h) \\
\text{values} &: (V, V^h)
\end{aligned}
$$

that solve firm and household problems and clear the markets for assets, labor, and output:

1. V satisfies (6) - (8), and (N, K) are the associated policy functions for firms.

2. V^h satisfies (9) - (11), and (C, N^h, Λ^h) are the associated policy functions for households.

3. $\Lambda^h(\overline{\varepsilon}, t, \varepsilon, k, \mu; s, \mu) = \mu(\overline{\varepsilon}, t, \varepsilon, k)$ for each $(\overline{\varepsilon}, t, \varepsilon, k) \in \mathcal{S}$.

4. The labor and goods market clear.

$$
\begin{aligned}
N^h(\mu; s, \mu) &= \int_{\mathbf{S}} [N(\overline{\varepsilon}, t, \varepsilon, k)] \cdot \mu(d[\overline{\varepsilon} \times t \times \varepsilon \times k]) \\
C(\mu; s, \mu) &= \int_{\mathbf{S}} [z\varepsilon F\left(k, N(\overline{\varepsilon}, t, \varepsilon, k)\right) - (K(k, b, \varepsilon; z, \mu) - (1-\delta)k)] \cdot \mu(d[\overline{\varepsilon} \times t \times \varepsilon \times k])
\end{aligned}
$$

5. the resulting individual decision rules for firms and households are consistent with the aggregate law of motion, Γ, where Γ defines the mapping from μ to μ'.

Using $C(s, \mu)$ and $N(s, \mu)$ to describe the market-clearing values of household consumption and hours worked, it is straightforward to show that market-clearing requires that (a) the real wage equal the household marginal rate of substitution between leisure and consumption:

$$
w\left(s, \mu\right) = D_2 U\Big(C(s, \mu), 1 - N(s, \mu)\Big) / D_1 U\Big(C(s, \mu), 1 - N(s, \mu)\Big),
$$

that (b) firms' state-contingent discount factors are consistent with the household marginal rate of substitution between consumption across states:

$$
d\left(s', s, \mu\right) = \beta D_1 U\Big(C(s', \mu'), 1 - N(s', \mu')\Big) / D_1 U\Big(C(s, \mu), 1 - N(s, \mu)\Big).
$$

4 Quantitative Analysis

In this section, I present my calibration strategy to match both micro and macro data. I set the length of a period for this model to 1 year and solve the model using a non-linear method, which involves value function iterations over the state space described in the model section. In this paper, there is a non-trivial time-varying distribution of firms, which is a part of the aggregate state in this economy. I take the Krusell Smith (1997) approach that is implemented by Khan and Thomas (2003, 2008) in a heterogeneous firm model.[14]

[14] Terry (2014) compares a variety of alternative approaches to solve heterogeneous firm models with aggregate uncertainty.

4.1 Steady State and Calibration

Functional forms and stochastic processes

I assume that the representative household's period utility is $u(c, L) = \log c + \eta L$, as in models of indivisible labor (e.g., Hansen, 1985; Rogerson, 1988). As seen in the previous sections, I assume that each heterogeneous firm undertakes production via the Cobb-Douglas production function: $z\varepsilon(k^{\alpha}n^{1-\alpha})^{\nu}$, where α determines capital and labor's share of income and ν governs returns to scale in this economy. For aggregate and idiosyncratic productivity processes: z and $\varepsilon = \theta + a$, I assume

$$\begin{aligned}
\log z' &= \rho_z \log z + \eta_z' \text{ with } \eta_z' \sim N\left(0, \sigma_{\eta_z}^2\right) \quad \text{and} \qquad (12)\\
\varepsilon &= \theta + a \\
&: \theta \sim N(\overline{\theta}, \sigma_{\theta}^2) \text{ and} \\
&: a \sim N(0, \sigma_a^2). \qquad (13)
\end{aligned}$$

$\overline{\theta}$ is the mean, and σ_{θ}^2 is the variance of the base component of idiosyncratic productivity, and σ_a^2 is the variance of the temporary component of idiosyncratic productivity.[15]

For time-varying π, I assume that π follows a two-state Markov chain with π_L and π_H. The transition matrix is $\Pi = \begin{bmatrix} \rho_L & 1-\rho_L \\ 1-\rho_H & \rho_H \end{bmatrix}$.

Common parameters

I calibrate the following five parameters against aggregate moments for the U.S. economy: (1) $1-\alpha$: labor's income share, (2) ν: returns to scale, (3) β: the household discount factor, (4) δ: the depreciation rate and (5) η: the leisure preference. First, I set ν to imply an average private capital-to-output ratio of 2.3, given the value of $1-\alpha$ determining the average labor share of income at 0.6 (Cooley and Prescott, 1995). Next, the depreciation rate, δ, is taken so that the model matches an average investment-to-capital ratio at 0.07. The preference parameter, η, is set to imply an average hours worked of one-third. Finally, I set the household discount factor to match an average real interest rate of 4 percent as in Gomme, Ravikumar and Rupert (2011).[16]

[15] I discretize these productivity processes and ensure all grid values are positive in calibrated version of the model.

[16] The average private capital-to-output ratio and the average investment-to-capital ratio are calculated from the U.S. National Income and Product Accounts Tables and Fixed Assets Accounts Tables for postwar periods as in Khan and Thomas (2013).

Firm-level parameters and aggregate shocks

Given the common parameters calibrated as above, I jointly calibrate the following firm-level parameters and then set the parameters that govern exogenous aggregate shock processes. First, (1) the mean of base components of idiosyncratic productivity, $\overline{\theta}$, (2) the variance of base components of idiosyncratic productivity, σ_{θ}^2, (3) the variance of temporary components of idiosyncratic productivity, σ_a^2, and (4) the steady-state level of the reset probability, π_L, are calibrated to match the following moments from my panel.

First, I aim to obtain the moments relating to ROA consistent with the empirical data, particularly the cross-sectional moments regarding the actual ROA and ROA forecast errors, defined as the percentage difference between the actual ROA and the conditional expectation made in the previous period. The left top panel of Figure 3 shows the distribution of ROA and the left bottom panel of Figure 3 shows forecast errors about ROA from the model-simulated data. The model captures the cross-sectional moments of ROA well. As in the data, the distribution of ROA is skewed to the left with the mean (model: 0.29, data: 0.30) and standard deviation (model: 0.095, data: 0.111). In addition to these actual ROA moments, the model is also able to match with the moments of forecast errors about ROA. Consistent with the data, the distribution is skewed to the right with the mean (model: 0.028, data: 0.026), standard deviation (model: 0.034, data: 0.049) and interquartile range (model: 0.023, data: 0.025).

Second, as becomes clear later, the distribution of investment rates has an important implication for the aggregate results of the model; thus, I intend to achieve a tight correspondence between the model and the data. In particular, the fraction of firms that undertake negative investment is matched with the data (model: 0.102, data: 0.161).[17]

Third, I calibrate the process for the two aggregate shocks as follows. I set the high reset probability, π_H, to reproduce the size of changes in the cross-sectional average of forecast errors between low- to high-uncertainty periods in the panel data (38%).[18] The transition probabilities are estimated to match the transition patterns between low- to high-uncertainty periods during the same years in the panel data. Finally, the stochastic process of aggregate productivity is calibrated by setting ρ_z to 0.909 and σ_{η_z} to 0.014 as in Khan and Thomas (2013). All parameters are summarized in Table 4.

[17] I obtain the fraction of firms with negative investment using the measurement of investment as capital expenditure net of retirement from the Compustat data. Cooper and Haltiwanger (2006) report the fraction of observations with negative investment is 10.4 percent. While the former covers listed firms, the latter covers U.S. plants. Coverage of both factors does not extend recent years as data on retirement is not available; I err on the side of safety rather than overstating negative investment in the model.

[18] Low-uncertainty periods correspond to years of ferror_mean below the average over the sample periods, and high-uncertainty periods correspond to years of ferror_mean above the sample average.

Table 4: Parameter values

ν	α	δ	β	η	ϕ	σ_θ^2	σ_a^2	π_L	π_H	ρ_L	ρ_H	ρ_z	σ_z
0.80	0.25	0.069	0.96	2.0	0.081	0.195	0.0515	0.12	0.45	0.86	0.50	0.904	0.014

4.1.1 Untargeted Moments

The model predicts that forecast errors become smaller as learning proceeds. Forecast errors involve intrinsic errors due to σ_a, which is independent of learning. On the contrary, as firms become better informed as learning proceeds, the distance between the true value of θ and the conditional expectation $E[\theta|(\bar{\varepsilon}, t)]$ diminishes. This leads to the negative relationship between forecast errors and time-since-reset (TSR), defined as the number of periods elapsed since the last change of θ. This can be seen in Figure 4, which plots forecast errors from the data generated by the simulation. The right panel of Figure 4 shows the negative relationship between forecast errors and TSR. Further, the left panel of Figure 4 shows the negative relationship between forecast errors and firm size, measured by total assets. This negative relationship is due to the fact that the capital choice is convex in estimates of productivity due to decreasing returns to scale technology. As learning proceeds, capital accumulation by firms that revise their estimate upward is larger than capital deccumulation by firms that revise their estimate downward. Thus, as TSR increases, the cross-sectional average of capital stock increases, as in Figure 5.

Figure 6 compares the model against the data. The left panel of Figure 6 plots the mean of forecast errors and the right panel of Figure 6 plots the standard deviation of forecast errors against TSR from the model and firm age from the data. As Figure 6 shows, the model captures well the features of the negative relationship in the mean of forecast errors against learning duration, although the model overall understates the standard deviation.[19]

4.1.2 Two-Sided Capital Misallocation

Imperfect information about total factor productivity across firms causes a misallocation of capital and labor. Firms operating with imperfect information deviate from the optimal allocation of resources and exhibit both over- as well as undercapacity. This pattern of misallocation is distinct from that which appears with financial frictions such as lending

[19] Note, though, that learning duration in the model corresponds to TSR but in data corresponds to firm age, measured by the number of periods appearing in the data.

subject to default risk (e.g. Khan et al., 2014) or a collateral constraint (e.g. Buera and Moll, 2013).

When a firm believes its base productivity is higher (lower) than the prior, its capital stock tends to be lower (higher) than the efficient level based on full information. This undercapacity (overcapacity) persists over time until the posterior mean converges to the true base productivity. Overall, unless firms are fully informed about their base components, their capital stock is either excessive or insufficient relative to the efficient level consistent with the true value of its base component and the interest rate. The longer it takes a firm to learn, the more severe the resource misallocation problem.

Figure 7 provides an example of learning and capital accumulation patterns from one individual firm. The top panel shows the idiosyncratic productivity, which is the sum of the true base component and the i.i.d. component. The middle panel shows the uncertainty proxies: the forecast errors and the posterior variance of forecasts of the base productivity. The bottom panel plots the capital stock chosen by the firm and compares it with the frictionless capital stock choice for this firm.

In this example, the firm experiences a resetting of its base productivity twice, at period 7 and at period 12. Following these resettings, the posterior variance of forecast of the base productivity jumps and then gradually decreases again as learning proceeds. In the early period of learning, not only the posterior variance but also the forecast errors tend to be large. Furthermore, the capital stock tends to deviate from the frictionless levels. For example, the firm gradually scales up its capital stock following a rise in its base productivity at period 7. Over this period, its capital stock is inefficiently low. In contrast, the firm maintains its capital stock above its frictionless level for a while after period 14. From period 14, the firm's observation of productivity is higher than the base level due to the i.i.d. components; however, because the firm is in the early stage of learning, it incorporates these i.i.d. components into their posterior belief, thus choosing its capital stock at the higher level relative to the frictionless level. As such, every time base components are reset and there is a change in base productivity, firms adjust their capital stocks slowly and a misallocation of capital persists. The time-varying reset probabilities have important cyclical implications as they change the degree of misallocation of capital and labor. In the next section, we explore dynamic implications in this environment.

4.2 Business Cycles

Table 5 presents the business cycle moments for a 2,500-period unconditional simulation with both aggregate productivity shocks and uncertainty shocks. Some of the features of the model business cycle are summarized as follows. Most second-moment statistics

Table 5: Unconditional business cycle moments

	Y	C	I	N	K
Mean	0.591	0.502	0.090	0.330	1.299
Standard deviation relative to Y	1.000	0.460	5.112	0.710	0.569
Correlation with Y	1.000	0.769	0.932	0.910	0.122

Note: The table above presents business cycle moments from a 2,500-period unconditional simulation. All series are HP-filtered in logs with a smoothing parameter of 100. The first row reports the standard deviation of the HP-filtered log series. The second row reports the relative size of the standard deviation of the HP-filtered log series to the standard deviation of the output series. The third row reports the contemporaneous correlation with the output series.

generated from the simulation are standard and familiar when evaluated against the business cycle literature. Specifically, consumption, investment and hours co-move with output. Consumption is less volatile than output, while investment is more volatile than output and, indeed more than its empirical counterpart.

4.2.1 The Great Recession Simulation

In this section, I explore the mechanism that propagates uncertainty shocks in the model. To accomplish this, I study an impulse response following an uncertainty shock, which involves a rise in the reset probability π. The size of the shock is set to reproduce the observed increase in the cross-sectional average of forecast errors, $ferror\_mean_t$, during the Great Recession. Specifically, I set π to 0.35 for four periods so that the model reproduces that the average size of forecast errors rises by 31% from the average between 2005 and 2006 to the average between 2008 and 2009. In figure 9, I show the model economy's response for two cases: the benchmark model with imperfect information (learning model) and a model without imperfect information (no learning model). The second model has an observable base component and is otherwise identical to the benchmark. The solid line is the response of the benchmark model to the uncertainty shock. The dashed line is the response of the no learning model to the uncertainty shock. For the benchmark model, the uncertainty shock alone reduces measured TFP by 0.73%, which is 33% of the observed reduction, and GDP by 1.55%, which 28% of the observed reduction. Investment falls by 5.89%, which is 31% when compared to the data.[20]

[20] In the data, the size of the recession is measured by the percentage change in each variable from the peak to the trough, 2007Q4 to 2009Q2.

In the following sections, I further investigate how uncertainty shocks produce recessions in the model. I first explain the difference between the mechanism through which conventional uncertainty shocks in stochastic volatility models operate and the distinct mechanism in my model. I then focus on the role of learning in shaping aggregate fluctuations by comparing the model responses with learning and without learning. This will reveal an important mechanism through which the process of learning prolongs recessions in the model.

Relationship with Conventional Uncertainty Shocks

The conventional framework used in business cycle studies of uncertainty shocks assumes stochastic volatility where the variance of a stochastic processes is allowed to be time-varying. Bloom (2009), for example, introduces shocks to both the level of productivity and its variance.[21] In stochastic volatility models, a shock to the variance can lead to a recession. One mechanism that has been emphasized by the literature is the real option value associated with factor adjustment in the presence of nonlinear adjustment costs.

For any given firm, there are two effects that work in opposite directions following a shock to the variance of productivity. On the one hand, the firm might increase its investment due to Jensen's inequality effects (Oi-Hartman-Abel effects).[22] This is because the optimal choice of capital is convex in productivity. On the other hand, the firm might pause its investment completely and wait for the resolution of uncertainty. With more volatile productivity shocks, the value of an option to wait to see future outcomes in the following period increases. Therefore, the firm may undertake no investment. In economies with heterogeneous firms, some firms pause investment while other firms increase investment, depending on their levels of individual productivity. Quantitative studies in the literature have shown that the latter effect dominates and due to this extensive driven-mechanism the aggregate economy falls into recession with higher uncertainty.

[21] See, for example, Fernandez-Villaverde, Guerron-Quintana, Rubio-Ramirez and Uribe (2011), Basu and Bundick (2012), Arellano, Bai and Kehoe (2012), Gilchrist, Sim and Zakrajsek (2014), Christiano, Motto, Rostagno (2014) and Schaal (2015), Born and Pfeifer (2014) and Fernndez-Villaverde, Guerrn-Quintana, Kuester and Rubio-Ramrez (2015), among others.

[22] The positive impact of uncertainty shocks on investment is known as the Oi-Hartman-Abel effect (Oi, 1961; Hartman, 1972; Abel, 1983).

Recessions with Bayesian Uncertainty Shocks

In contrast to the extensive margin-driven mechanism with conventional uncertainty shocks, Bayesian uncertainty shocks in this paper operate through an intensive margin.

Uncertainty shocks lead to more churning of firm productivity, while the distribution of firm productivity remains unchanged. Each firm anticipates a higher likelihood of changing its productivity. The changes in the conditional expectation of next period productivity vary its size and sign across firms, depending on the current location on the productivity distribution. For firms with a posterior mean of productivity higher than the prior mean productivity, a higher probability of changing productivity implies a downward shift in the conditional expectation, putting more weight on the unconditional prior distribution instead of on the conditional posterior distribution. For firms with a posterior mean of productivity lower than the prior mean productivity, a higher probability of changing productivity implies an upward shift in the conditional expectation. Furthermore, the larger the distance between the conditional posterior mean and the unconditional prior mean, the larger the shift in the conditional expectation after the shock. However, all these effects on the expectation offset each other as the distribution of productivity is symmetrical. Therefore, there is no expectation shift in the first moment of productivity from the aggregate perspective.

Even though these changes in expectation offset each other in the aggregate, each firm undertakes capital stock adjustments based on the newly revised conditional expectation. Importantly, these capital stock adjustments do not offset each other across firms due to decreasing returns to scale technologies; thus, the real effect emerges in the aggregate from these capital stock adjustments following Bayesian uncertainty shocks.

Figure 8 demonstrates these disparate responses in the capital stock adjustment across firms. Firms take expectations about their future productivity levels by looking at two different distributions of productivity simultaneously: one is their own posterior distribution that has been updated by learning, and the other is the unconditional distribution, which is the common prior known to all firms. Facing a higher reset probability, firms put more weight on the unconditional distribution than their posterior distribution. Since the variance of the unconditional distribution is larger than that of the posterior, firms effectively infer a larger variance of their future productivity distribution.

The direction of the shift in the posterior mean depends on the current posterior mean. For a firm with a current posterior that is higher than the prior mean, a higher reset probability leads to a larger variance of productivity shocks and a fall in its mean. This leads the firm to scale down its capital stock, as seen in the right panel of Figure 8. In contrast, for a firm whose current posterior is lower than the prior mean, a higher reset probability results in a larger variance of productivity shocks with an upward shift in the mean, inducing it to increase its capital stock, as seen in the left panel of Figure 8.

Nonetheless, these two effects do not offset each other due to the fact that the production function exhibits decreasing returns to scale. Given that the productivity distribution of firms is symmetrical around its mean, the net impact on aggregate investment is negative, and the economy enters a recession despite the absence of adjustment costs.

4.2.2 The Role of Learning

In Figure 9, I show how the two models with and without learning shape the recovery patterns of the economy after recessionary periods. While investment, labor, and measured TFP series overshoot in the model without learning, the benchmark model with learning eliminates this rapid recovery. Instead, the benchmark model with learning exhibits a gradual recovery following a recession.

To gauge the importance of imperfect information following uncertainty shocks, I decompose the impact of uncertainty shocks into the two effects explained below. I first categorize firms into cohorts by their time-since-reset (TSR), and I look at the average investment for each cohort. This allows me to see the disaggregated investment response of firms. Furthermore, by comparing the mass of firms for each cohort, I can trace changes in the distribution of firms throughout the recession. This will prove useful for understanding the mechanism behind the rapid drops and slow recoveries in my model. Figure 10 shows this exercise for the beginning of the recession and the recovery separately.

4.2.3 The Onset of a Recession

When firms anticipate a higher reset probability for their base components of productivity, an uncertainty effect leads them to change their target levels of capital. As argued above, for firms whose current posterior is higher than the prior mean, the higher reset probability implies a larger variance for productivity shocks and fall in their mean, leading to a downward adjustment of capital stocks. On the other hand, for firms whose current posterior mean is lower than the prior, uncertainty shocks imply a larger variance of productivity shocks with an upward shift in the mean, resulting in an upward adjustment of capital stocks. Given that the distribution of the base component of productivity is symmetric and the production function has decreasing returns to scale, the downward adjustment of capital stocks for the top 50% of firms tends to dominate the opposing force from the bottom 50% of firms. Thus, investment falls for cohorts 6 to 20, as shown in the top left panel of Figure 10. This shifts the average investment curve from the solid to the dashed line.[23]

The top right panel of Figure 10 highlights the distributional effects. While uncertainty effects are now less pronounced and investment has largely reversed, there is a large inflow of firms into cohort 1 relative to the pre-recession level. Since the average investment level of cohort 1 is the lowest among all cohorts, this shift of the firm distribution leads to a substantial drop in aggregate investment.

4.2.4 A Slow Recovery

In this subsection, I examine how imperfect information eliminates an overshoot of investment. After the shock ends, firms raise their expectation of maintaining their current level of productivity. Now, if firms believe that their base component is higher than the mean, they are more confident in raising their scale of production. If firms believe that their base component is lower than the mean, they reduce their scale of production. For the reasons in the previous subsection, the pent-up demand of firms with higher productivity shifts the average investment curve from the dashed to the solid line, as seen in the bottom left of Figure 10.

As may be seen in there, this pent-up investment demand effect is strong in cohorts with large time-since-reset. While this potentially increases aggregate investment, the mass of firms in the relevant cohorts (3 and onward) is small compared to the pre-recession level; therefore, aggregate investment will not recover. Furthermore, despite the end of uncertainty shocks, the level of investment for cohorts 1 and 2 remain low due to a general equilibrium effect.[24] This also prevents a sudden recovery in aggregate investment after the shock ends.

Imperfect information not only eliminates an overshooting of investment but also slows the pace of recovery afterwards. The bottom right panel of Figure 10 explains how the model economy slowly recovers to its pre-recession level in the periods after the uncertainty shock. The key mechanism is misallocation. To achieve an efficient level of capital stock, firms need to have accurate information about their productivity. Thus, misallocation of capital and labor is more severe among cohorts with a smaller time-since-reset. As the figure shows, the mass of firms within cohorts 2 and 3 is larger than in the steady-state,

[23] This is easier to see in the top left panel of Figure 12, in which each line represents the percent change in investment from the steady state. The rise in the average investment in cohorts 1 to 4 is a general equilibrium effect. As shown in the top left panel of Figure 11, these rises are not present in partial equilibrium.

[24] See Figure 11 for the case in which investment increases for all cohorts in partial equilibrium.

while the mass of firms in cohort 4 and onward is smaller than before the recession. As time goes by, the mass of firms in cohorts 2 and 3 will gradually fill up the gap in the size of mass in cohort 5 and onward. Due to a slow-moving distribution of firms related to learning, the negative aggregate effect from misallocation persists until the distribution of information in the economy eventually returns to match that in the steady state.

5 Conclusion

This paper develops a heterogeneous firm model that incorporates Bayesian learning at the firm level with uncertainty shocks. The model can capture the cross-sectional and cyclical features of firm-level uncertainty well. In the model, uncertainty will be resolved gradually as firms operate longer and get better informed. Thus, it can replicate the negative relationship between the duration of operations and the size of forecast errors, consistent with data. Further, the model establishes a close link between the rise in firms' uncertainty at the start of a recession and the slow pace of subsequent recovery.

The approach taken in this paper to modeling uncertainty may be useful in other applications. For example, while this model has a very simple hiring and firing decision, economists have emphasized jobless recoveries and the mechanism proposed in this paper may offer important insights that link these to a rise in firms' uncertainty. Further, there have been attempts to link financial markets and aggregate fluctuations since the recent financial crisis. As stated by Bernanke (2008), *"The crisis we face in the financial markets has many novel aspects, [. . .] at the root of the problem is a loss of confidence by investors and the public in the strength of key financial institutions and markets."* Researchers may find it useful to examine model environments with type of time-varying uncertainty proposed here so as to study the link between a deterioration of trust in financial markets and recessions.

Furthermore, the way I discipline my model with learning via firm-level forecast errors may be useful in other exercises. Uncertainty cannot be observed directly and this makes it harder to render the quantitative analysis of such models using microdata. However, as shown in this paper, the model can be pin-downed by the time dependant features of measures of uncertainty such as forecast errors. Nonetheless, forecast errors that are constructed in this paper are still indirect as they reflect market views rather than manager's view. In this regard, more direct measurement of firm-level uncertainty can be explored more. Using forecasts by firm managers about future outcomes such as sales and product prices, we will be able to compare various measures of firm uncertainty more comprehensively, which is a fruitful area of studies of firm-level uncertainty going forward.

References

Abel, A. (1983). Optimal investment under uncertainty. *American Economic Review 73*(1), 228–233.

Abel, A. B. and J. C. Eberly (1996). Optimal investment with costly reversibility. *The Review of Economic Studies 63*(4), 581–593.

Alti, A. (2003). How sensitive is investment to cash flow when financing is frictionless? *The Journal of Finance 58*(2), 707–722.

Arellano, C., Y. Bai, and P. J. Kehoe (2012). Financial frictions and fluctuations in volatility. Staff Report 466, Federal Reserve Bank of Minneapolis.

Bachmann, R. and C. Bayer (2013). Wait-and-See business cycles? *Journal of Monetary Economics 60*(6), 704–719.

Bachmann, R., S. Elstner, and E. R. Sims (2013). Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics 5*(2), 217–249.

Baley, I. and J. A. Blanco (2016). Menu costs, uncertainty cycles, and the propagation of nominal shocks. Economics Working Papers 1532, Department of Economics and Business, Universitat Pompeu Fabra.

Basu, S. and B. Bundick (2012). Uncertainty shocks in a model of effective demand. NBER Working Papers 18420, National Bureau of Economic Research.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics 98*(1), 85–106.

Bertola, G. and R. Caballero (1994). Irreversibility and aggregate investment. *Review of Economic Studies 61*(2), 223–246.

Blanchard, O. (2009). (nearly) nothing to fear but fear itself. *The Economist*.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica 77*(3), 623–685.

Bloom, N., S. Bond, and J. Van Reenen (2007). Uncertainty and investment dynamics. *The Review of Economic Studies 74*(2), 391–415.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2014). Really Uncertain Business Cycles. Working Papers 14-18, Center for Economic Studies, U.S. Census Bureau.

Bond, S., R. Moessner, H. Mumtaz, and M. Syed (2005). Microeconometric evidence on uncertainty and investment. *Institute for Fiscal Studies, London, UK*.

Born, B. and J. Pfeifer (2014). Policy risk and the business cycle. *Journal of Monetary Economics 68*, 68–85.

Buera, F. J., J. P. Kaboski, and Y. Shin (2011). Finance and development: A tale of two sectors. *The American Economic Review 101*(5), 1964–2002.

Buera, F. J. and B. Moll (2015). Aggregate implications of a credit crunch: The importance of heterogeneity. *American Economic Journal: Macroeconomics 7*(3), 1–42.

Caballero, R. J. and E. M. Engel (1999). Explaining investment dynamics in us manufacturing: a generalized (s, s) approach. *Econometrica 67*(4), 783–826.

Caplin, A. and J. Leahy (1993). Sectoral shocks, learning, and aggregate fluctuations. *The Review of Economic Studies 60*(4), 777–794.

Christiano, L. J., R. Motto, and M. Rostagno (2014). Risk shocks. *American Economic Review 104*(1), 27–65.

Cooper, R. W. and J. C. Haltiwanger (2006). On the Nature of Capital Adjustment Costs. *Review of Economic Studies 73*(3), 611–633.

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

Dixit, A. K. and R. S. Pindyck (1994). *Investment under uncertainty.* Princeton university press.

Fajgelbaum, P., E. Schaal, and M. Taschereau-Dumouchel (2014). Uncertainty Traps. NBER Working Papers 19973, National Bureau of Economic Research.

Fernandez-Villaverde, J., P. Guerron-Quintana, J. F. Rubio-Ramirez, and M. Uribe (2011). Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review 101*(6), 2530–2561.

Fernndez-Villaverde, J., P. Guerrn-Quintana, K. Kuester, and J. Rubio-Ramrez (2015). Fiscal Volatility Shocks and Economic Activity. *American Economic Review 105*(11), 3352–3384.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2014). Uncertainty, financial frictions, and investment dynamics. NBER Working Papers 20038, National Bureau of Economic Research.

Giordani, P. and P. Soderlind (2003). Inflation forecast uncertainty. *European Economic Review 47*(6), 1037–1059.

Gomme, P., B. Ravikumar, and P. Rupert (2011). The Return to Capital and the Business Cycle. *Review of Economic Dynamics 14*(2), 262–278.

Guiso, L. and G. Parigi (1999). Investment and Demand Uncertainty. *The Quarterly Journal of Economics 114*(1), 185–227.

Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of monetary Economics 16*(3), 309–327.

Hartman, R. (1972). The effects of price and cost uncertainty on investment. *Journal of Economic Theory 5*(2), 258–266.

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

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly Journal of Economics 4*, 1403–1448.

Ilut, C. and H. Saijo (2016). Learning, confidence, and business cycles. Working paper.

Janunts, M. (2010). *Differences of opinion and stock returns*. dissertation, Universit de Neuchtel.

Johnson, T. C. (2004). Forecast dispersion and the cross section of expected returns. *The Journal of Finance 59*(5), 1957–1978.

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

Kehrig, M. (2016). The cyclicality of productivity dispersion. Working paper.

Keynes, J. M. (1936). The general theory of employment, interest, and money.

Khan, A., T. Senga, and J. K. Thomas (2014). Default risk and aggregate fluctuations in an economy with production heterogeneity. Working paper, The Ohio State University.

Khan, A. and J. K. Thomas (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica 76*(2), 395–436.

Khan, A. and J. K. Thomas (2013). Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity. *Journal of Political Economy 121*(6), 1055 – 1107.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2015). The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation. NBER Working Papers 21719, National Bureau of Economic Research.

Krusell, P. and A. A. Smith (1997). Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns. *Macroeconomic dynamics 1*(02), 387–422.

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy 106*(5), 867–896.

Leahy, J. and T. Whited (1996). The effect of uncertainty on investment: Some stylized facts. *Journal of Money, Credit and Banking 28*(1), 64–83.

Li, W. and J. Weinberg (2003). Firm-Specific Learning and the Investment Behavior of Large and Small Firms. *International Economic Review 44*(2), 599–625.

Oi, W. Y. (1961). The desirability of price instability under perfect competition. *Econometrica 29*(1), 58–64.

Orlik, A. and L. Veldkamp (2014). Understanding uncertainty shocks and the role of black swans. NBER Working Papers 20445, National Bureau of Economic Research.

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

Rich, R. and J. Tracy (2010). The Relationships among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts. *The Review of Economics and Statistics 92*(1), 200–207.

Rogerson, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics 21*(1), 3–16.

Saijo, H. (2014). The uncertainty multiplier and business cycles. *Tokyo Center for Economic Research (TCER) Paper No. E-67*.

Stein, L. C. and E. C. Stone (2013). The effect of uncertainty on investment, hiring, and r&d: Causal evidence from equity options. Working paper.

Terry, S. J. (2014). Alternative methods for solving heterogeneous firm models. Working paper.

Timoshenko, O. A. (2015). Product switching in a model of learning. *Journal of International Economics 95*(2), 233–249.

Van Nieuwerburgh, S. and L. Veldkamp (2006). Learning asymmetries in real business cycles. *Journal of monetary Economics 53*(4), 753–772.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics 129*(1), 215–258.

Zarnowitz, V. and L. A. Lambros (1987). Consensus and Uncertainty in Economic Prediction. *Journal of Political Economy 95*(3), 591–621.

Figure 1: The Great Recession

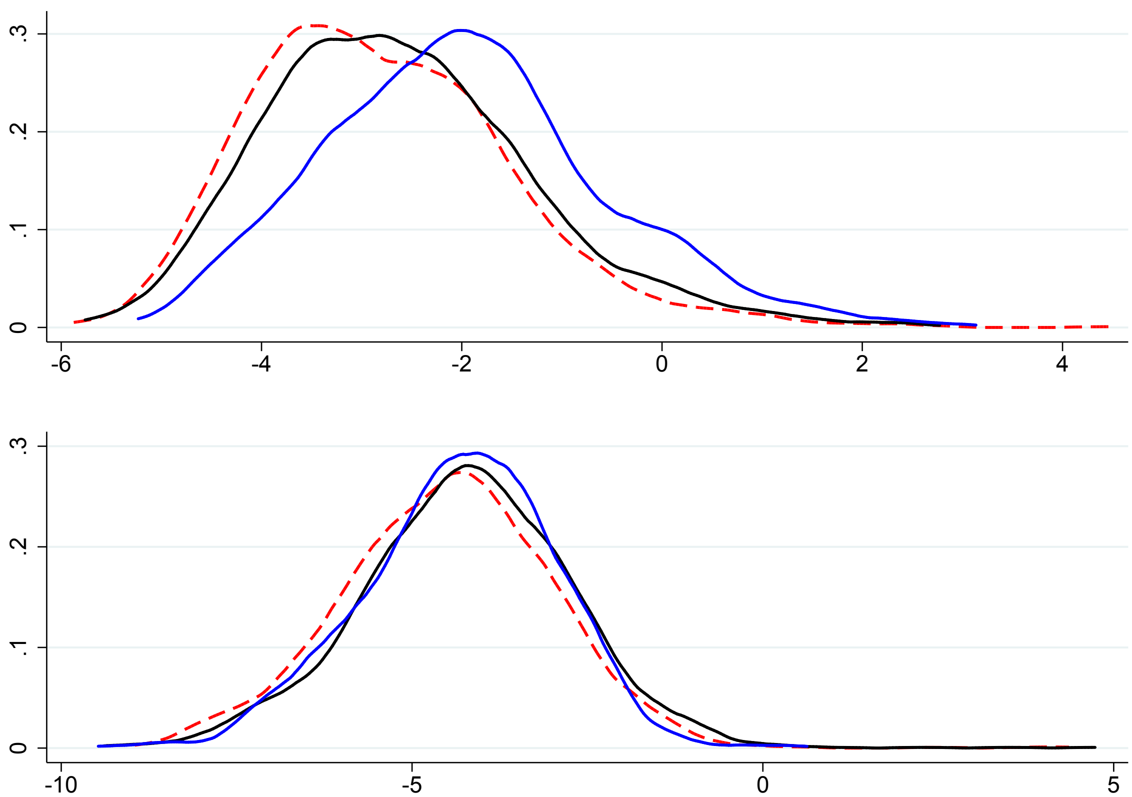

Note: The top panel plots the distributions of $fdisp$ across firms in 2007 (dashed line, red), 2008 (solid line, black) and 2009 (solid line, blue). The bottom panel plots the distributions of $ferror$ across firms in 2007 (dashed line, red), 2008 (solid line, black) and 2009 (solid line, blue).

Figure 2: Historical series

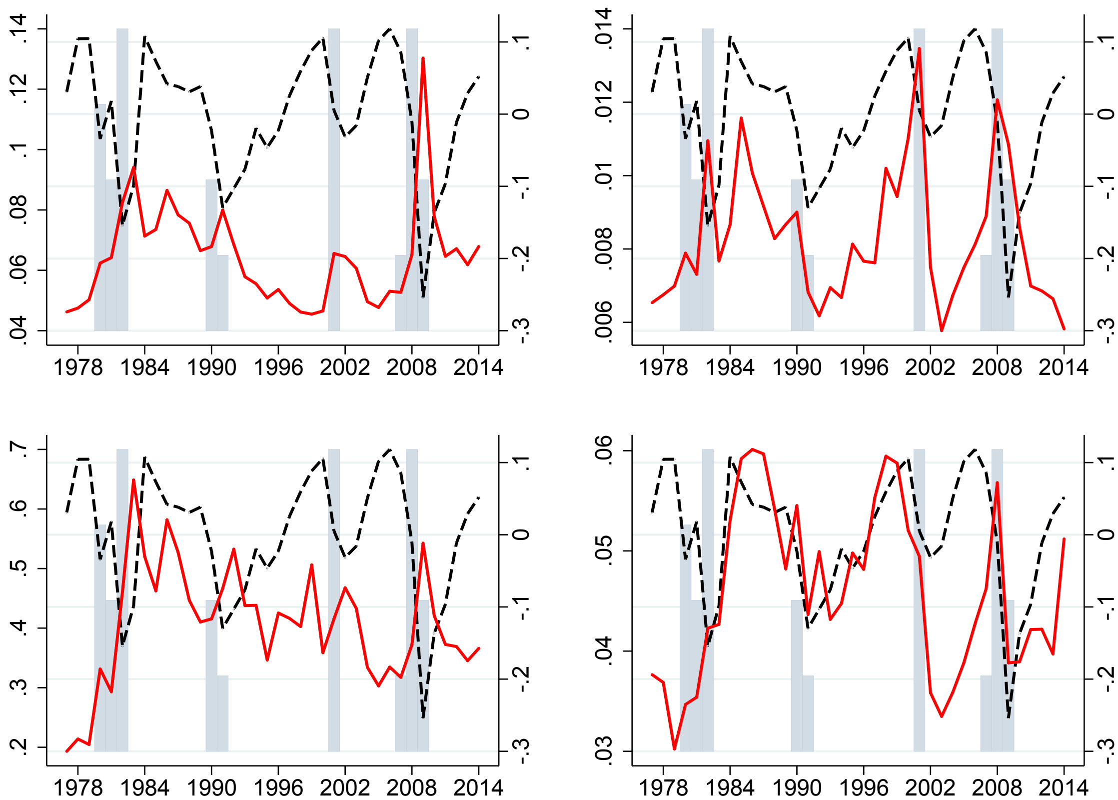

Note: The top left panel shows the mean of forecast dispersion (solid line, red). The top right panel shows the standard deviation of forecast dispersion (solid line, red). The bottom left panel shows the mean of forecast errors (solid line, red). The bottom right panel shows the standard deviation of forecast errors (solid line, red). In all panels, HP-filtered real GDP series are plotted (dashed line, black).

Figure 3: Return on assets, forecast errors and investment rate

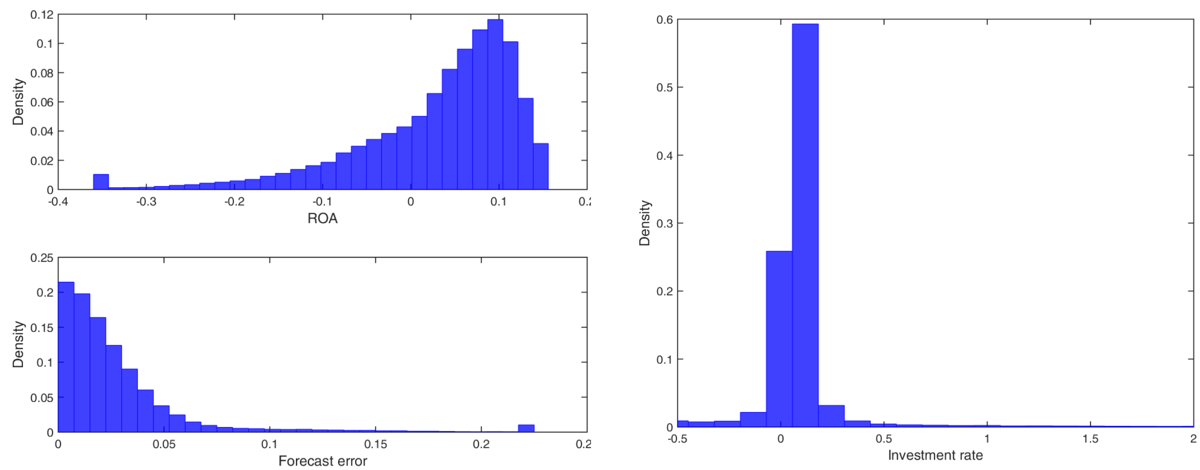

The upper left panel shows the distribution of return on assets (ROA), defined as a net profit divided by capital stock at the beginning of the period. The bottom left panel shows the distribution of forecast errors, which are defined as the percentage difference between the realized ROA and the conditional expectation of ROA made in the previous period. For ROA, moments of the distribution includes: mean (0.029) and standard deviation (0.095). For forecast errors, moments of the distribution includes: mean (0.028) and standard deviation (0.034). The right panel plots the distribution of investment rates. Moments of the distribution includes: mean (0.117), standard deviation (0.414), serial correlation (-0.012) and the fraction of negative investments (0.102). All data are generated from a simulation of 5,000 firms in the calibrated model, which was described in the text.

Figure 4: Forecast errors decrease as learning proceeds

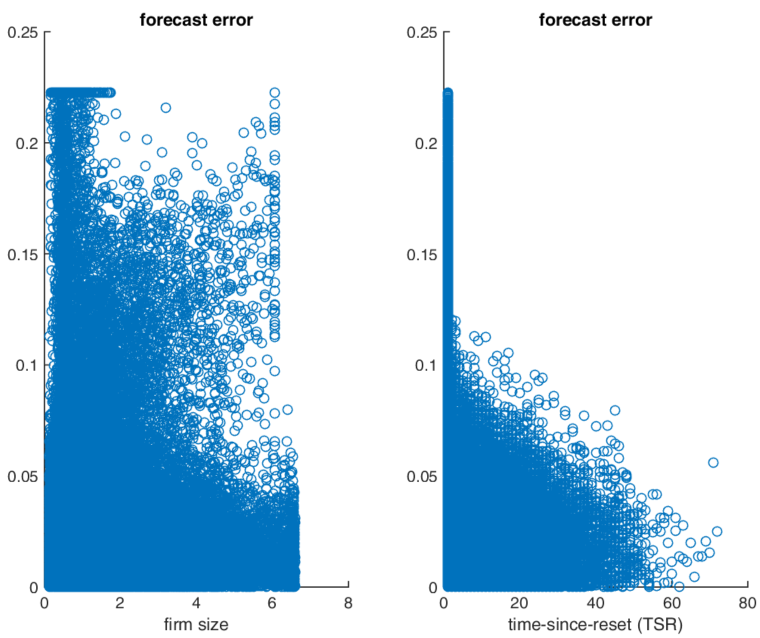

Note: The above plots show two scatter figures of forecast errors: one against capital stock in the left panel and one against time-since-reset in the right panel. Forecast errors are winsorized at one percent levels.

Figure 5: Capital choice

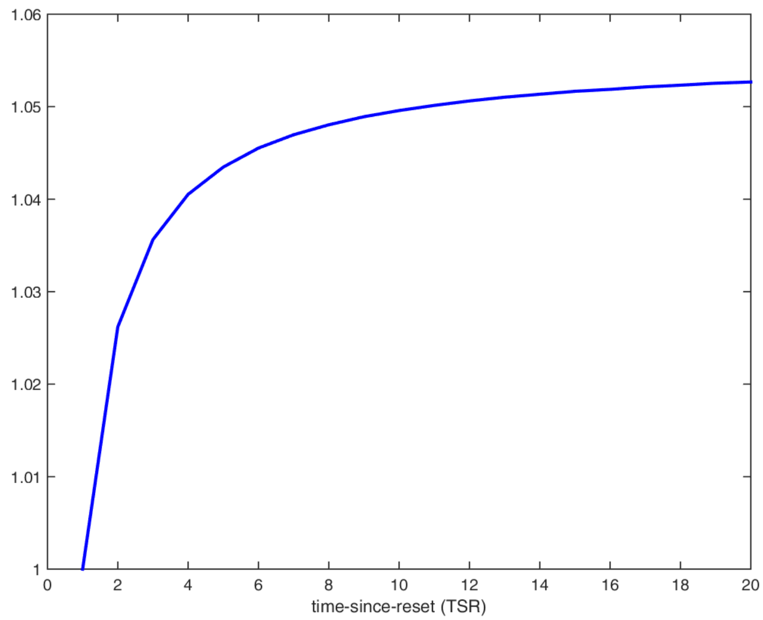

Note: The figure shows the average capital stock choice by TSR. All values are normalized by the initial level with TSR equal to 1.

Figure 6: Forecast errors and learning: Model and data

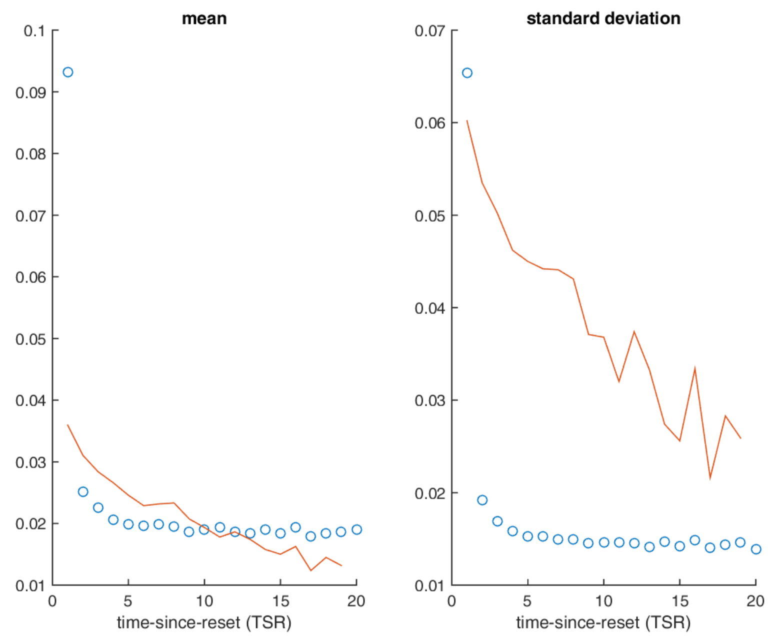

Note: The above figure shows the mean of forecast errors in the right panel and standard deviation of forecast errors in the left panel. In both panels, the dots represent the moment from the simulated model, and the lines represent the empirical moment obtained from the panel. Forecast errors are winsorized at one percent levels.

Figure 7: Learning cycles from simulation

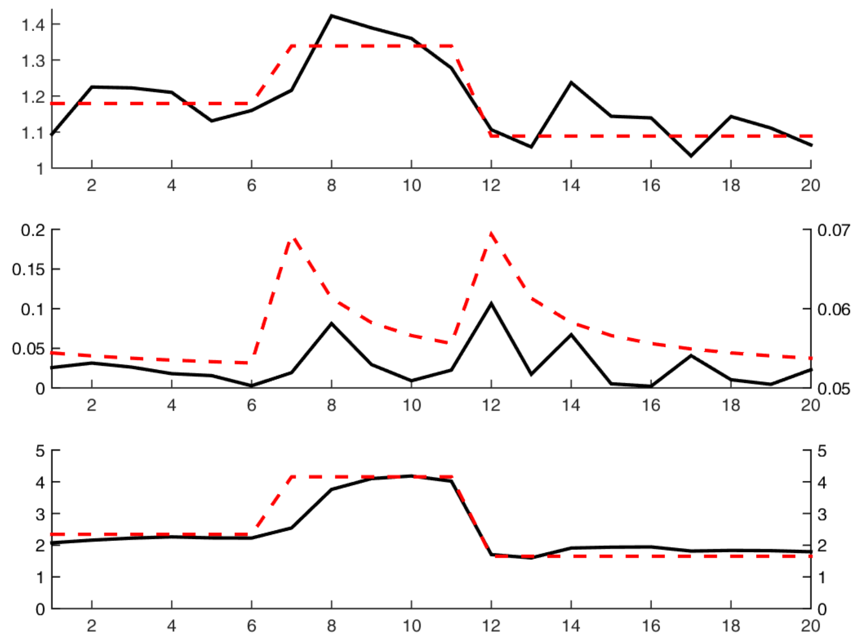

Note: This figure plots the patterns of the behavior of firms in the simulation without aggregate shocks. Five thousand firms are simulated for 5,000 periods, and a 20-period simulation result for one firm is taken here. The top panel shows a series of firm productivity (solid-line), which is observed by the firm, the sum of its base component, θ, and the i.i.d. component, a. The dot-line plots θ. The middle panel shows a series of forecast errors (solid-line) and the conditional variance of forecasts of productivity (dashed-line). The bottom panel shows a series of capital stocks (solid-line). The dashed-line corresponds to the frictionless level of capital stock in the absence of information frictions.

Figure 8: Disparate reactions to uncertainty shocks

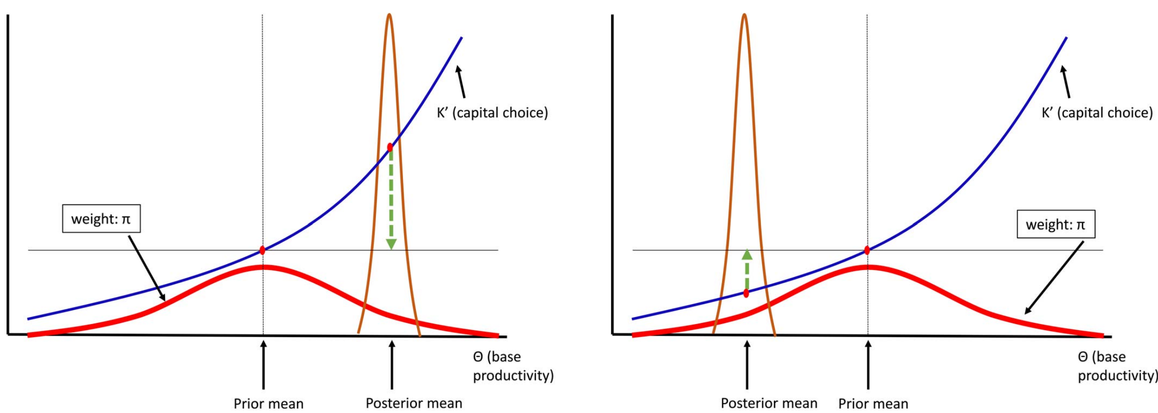

Figure 9: Uncertainty shock simulations with/without imperfect information

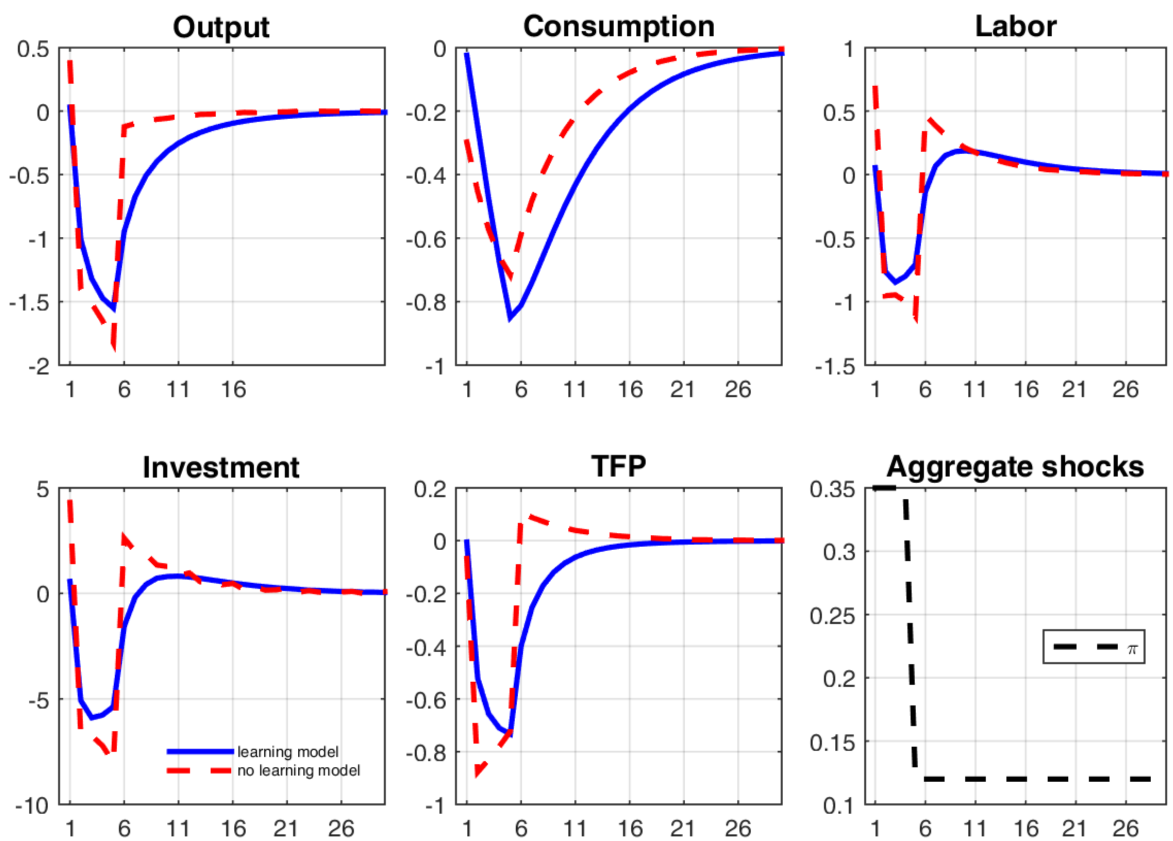

Note: Each panel except the lower right panel plots the aggregate economy's responses to uncertainty shocks. A solid blue line plots the responses of a model with learning, and a dashed red line plots the responses of a model without learning. The uncertainty shock is plotted in the lower right panel.

Figure 10: Dynamics of firm distribution and investment following uncertainty shocks: Benchmark case

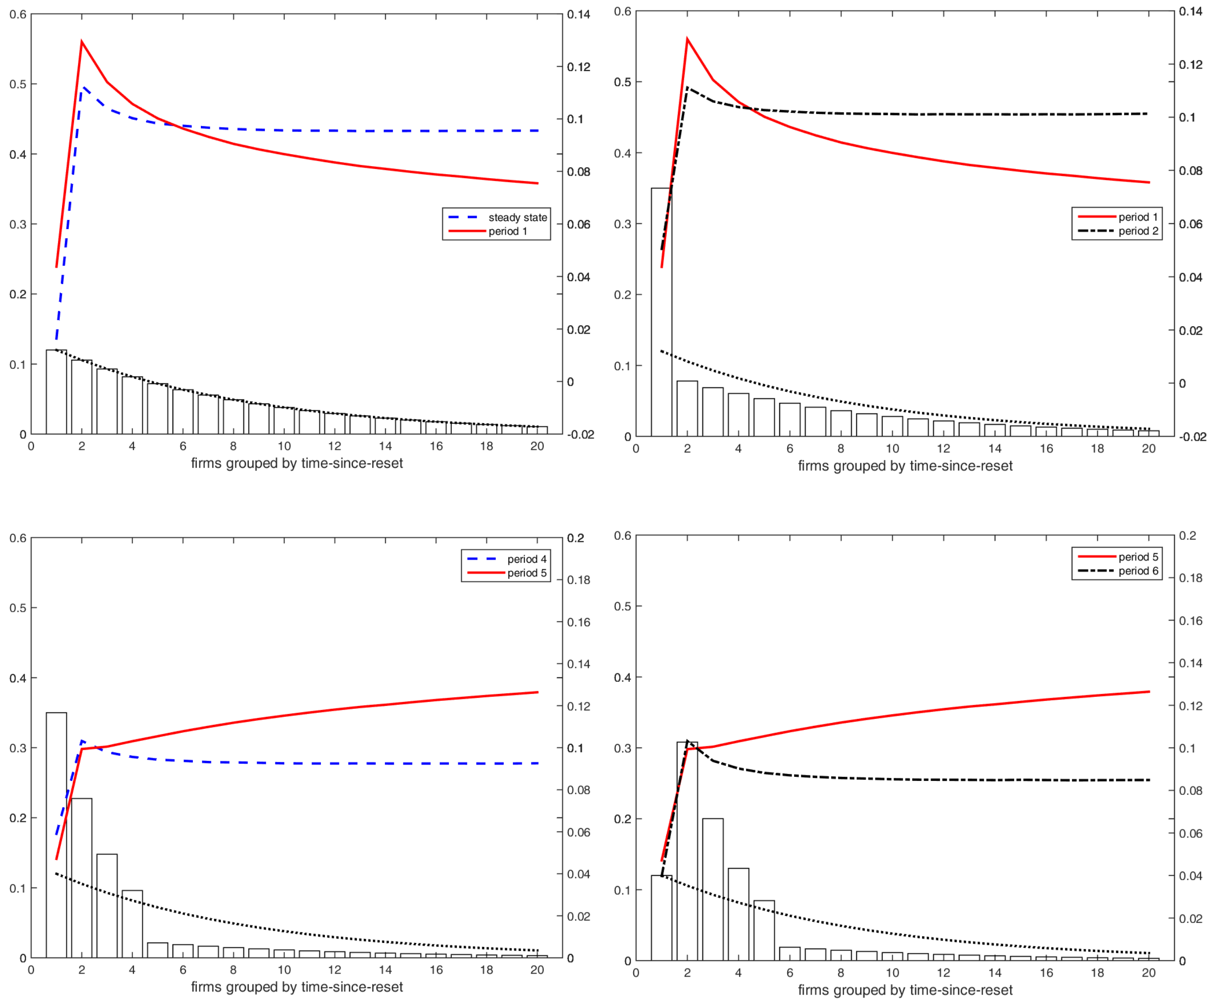

Note: Each bin represents the mass of firms in each cohort grouped by the time-since-reset (TSR) of their base component (left axis). A larger TSR implies that firms are more informed about their productivity levels. Each line plots the average level of investment for each cohort (right axis). Each dot shows the steady-state mass of firms in each time-since-reset bin.

Figure 11: Dynamics of firm distribution and investment following uncertainty shocks: Partial equilibrium

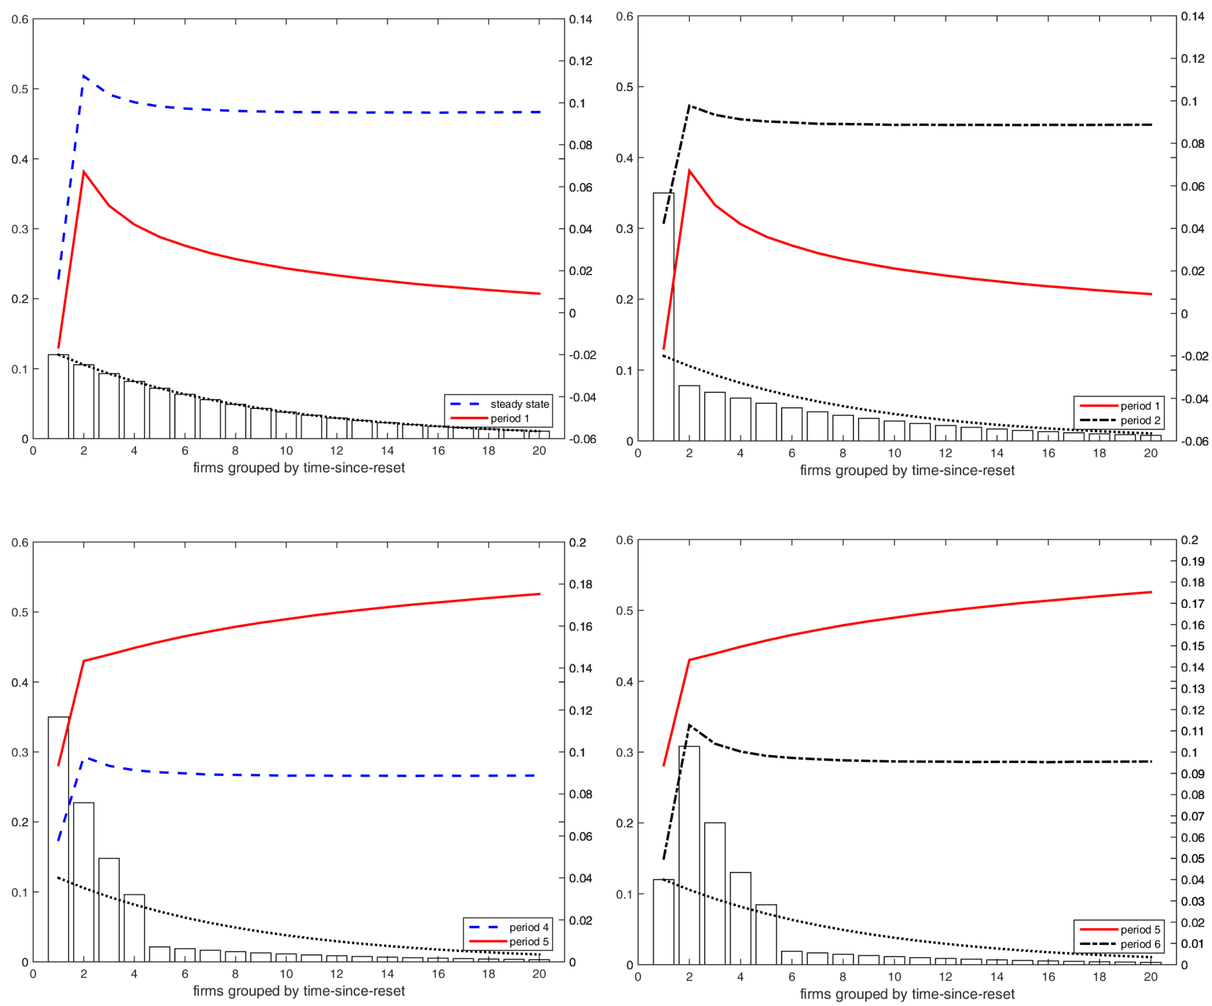

Note: Each bin represents the mass of firms in each cohort grouped by the time-since-reset (TSR) of their base component (left axis). A larger TSR implies that firms are more informed about their productivity levels. Each line plots the average level of investment for each cohort (right axis). Each dot shows the steady-state mass of firms in each time-since-reset bin.

Figure 12: Dynamics of firm distribution and investment following uncertainty shocks: Percentage change from the steady state

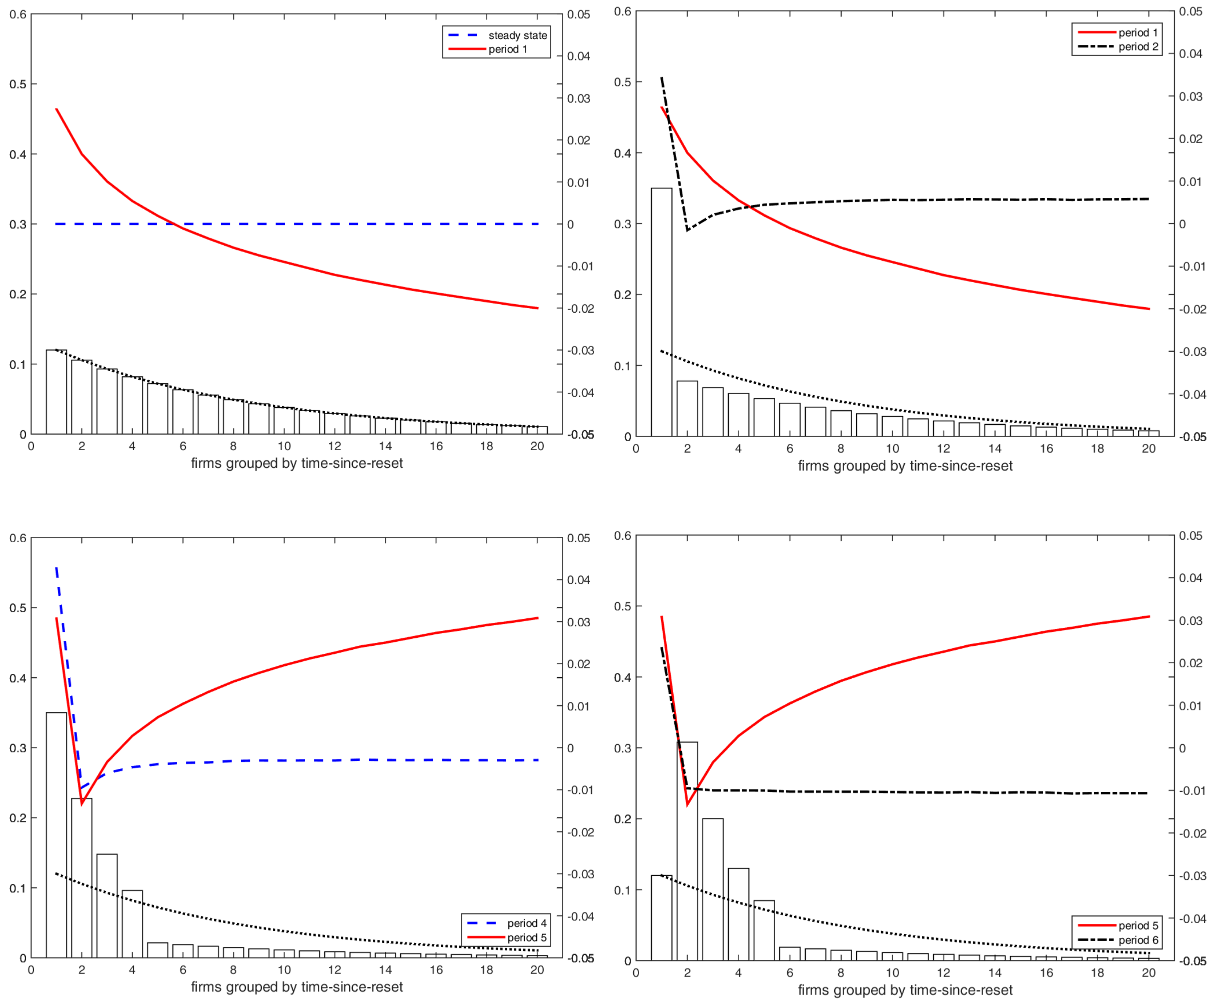

Note: Each bin represents the mass of firms in each cohort grouped by the time-since-reset (TSR) of their base component (left axis). A larger TSR implies that firms are more informed about their productivity levels. Each line plots the average level of investment for each cohort (right axis). Each dot shows the steady-state mass of firms in each time-since-reset bin.