Uncertainty Shocks in a Model with Mean-Variance Frontiers and Endogenous Technology Choices*

M. Saif Mehkari†

May 2016

Abstract

This paper builds a model to show how increases in aggregate uncertainty - an uncertainty shock - can generate recessions. Uncertainty shocks in the model are able to both account for a significant portion of business cycle fluctuations observed in data and generate positive comovements between output, consumption, investment, and hours. The key assumption of the model is that firm managers endogenously choose what projects to undertake and that the menu of these projects lies on a positively sloped mean-variance frontier – high-return projects are also high-risk projects. In times of high aggregate uncertainty, managers choose to undertake low-risk projects, and thus low-return projects, which in turn leads to a recession. Moreover, the model also matches various stylized facts about time series and cross-sectional variations in TFP and suggests shortcomings in using TFP data to calculate exogenous TFP shocks.

J.E.L. Classification: E3.
Keywords: Business Cycles, Uncertainty Shocks.

*I would like to thank Chadwick Curtis, Bill Dupor, Aubhik Khan, Nan Li, Daniel Mackay, two anonymous referees, along with various conference and seminar participants for comments and suggestions.
†Department of Economics, Robins School of Business, University of Richmond, USA. Email: smehkari@richmond.edu

1 Introduction

There is increasing evidence for the negative correlation between measures of economic activity and measures of uncertainty. For example, the correlation between the second moment of aggregate TFP, a measure of uncertainty, and real GDP is negative. This paper constructs a model where exogenous changes in the second moment of firm-level TFP - uncertainty shocks - endogenously generate business cycle fluctuations that account for some of these negative correlations.

In the model, risk-averse managers, in addition to making capital and labor hiring decisions, also decide on what projects to undertake. A manager's project choice determines the firm's stochastic TFP process. The critical assumption of this paper is that the menu of project choices, and thus TFP choices, available to the manager lie on a positively sloped mean-variance frontier. High-return projects are also high-risk projects. This generates a risk-return tradeoff for the manager. Whereas choosing a high-return project would, on average, generate more output, and thus profits, it does expose the manager to a higher amount of risk.

The presence of risk-return tradeoffs in the model cause exogenous changes in aggregate uncertainty - uncertainty shocks - to have important effects. Uncertainty shocks alter the riskiness of the project choices available to the manager causing them to reoptimize and adjust their production decisions. For example, as illustrated in Figure 1, risk-averse managers in the model find that a sudden increase in aggregate uncertainty makes their current project choice too risky, and they thus reoptimize by choosing a low-risk project. The low-risk project is also a low-return/low-TFP project. Consequently, during times of high aggregate uncertainty firm-level production falls, which in turn causes aggregate production to fall. A recession ensues. This mechanism not only qualitatively generates the observed negative correlation between the second moment of TFP and GDP, but is quantitatively also able to explain a significant portion of the variability observed in GDP data. Moreover, this mechanism also generates comovement between consumption and investment; a feature that eludes many real business cycle models driven by second moment shocks.

In addition to generating sizable business cycle fluctuations in real variables, the model in this paper also provides a rich set of results with regard to TFP. First, as described above an increase in aggregate uncertainty causes managers to endogenously choose low-return/low-TFP projects which in turn results in an endogenous drop in the mean level of aggregate TFP. This result replicates the observed negative correlation between the first and second moments of TFP in the data. Second, in the model, an assumption of heterogeneous risk preferences among the managers generates important heterogeneity in the amount of reoptimization each manager performs in response to changes in aggregate uncertainty. Specifically, as illustrated in Figure 1, during times of high aggregate uncertainty, low risk-averse managers reoptimize their choices very little along the frontier (L to L'), and thus, the average return/TFP level at firms run by low risk-averse managers falls very little during these periods. On the other hand, high risk-averse managers reoptimize significantly in response to high aggregate uncertainty (H to H'), and thus, in high aggregate uncertainty periods the average return/TFP level at their firms falls by a significant amount. This heterogeneity in responses causes the cross-sectional variance of firm-level TFP levels - cross-sectional dispersion of TFP - to rise and cross-sectional skewness to fall during recessions; two observations that are also true in the data.

The results discussed above when taken together, point to a broader contribution of the model, in that not only does the model explain how uncertainty shocks can lead to economically significant fluctuations in real variables, such as output, consumption, labor hours, and investment, it also provides a structural framework to explain how the various moments of TFP may be related. In particular, the model shows how a shock to one moment of the aggregate TFP series can propagate through firm-level decisions causing changes to other moments of both the firm-specific and aggregate TFP series.[1]

Understanding how exogenous shocks affect the different moments of the TFP series is important, because if such shocks cause endogenous movements in TFP, then it becomes difficult to disentangle the magnitude of TFP movements in the data that are exogenous vs. endogenous. For example, as explained above, in the model an exogenous increase in aggregate uncertainty results in an endogenous drop in the first moment of the aggregate TFP process. This suggests that in the data part of the fluctuations in the level of TFP may be purely endogenous and not indicative of the presence of independent exogenous TFP shocks. In my baseline calibration, exogenous changes in uncertainty are able to endogenously explain roughly a quarter of the variability in the mean level of TFP in the data. Further, in the model the true magnitude of the exogenous increase in the second moment is endogenously dampened by managers choosing relatively lower risk projects. This dampening effect illustrates how the second moment of TFP in the data can be a systematically biased estimate of the true underlying exogenous uncertainty shock.

Throughout this paper I motivate changes in the mean-variance frontier as resulting from an aggregate second moment shock; however, at this stage it is important to point out that identical changes can also result from a non-traditional aggregate first moment shock that shifts the mean-variance frontier. This is best seen by looking at Figure 1. A positive aggregate uncertainty shock works to move the mean-variance frontier *up*. This upward movement can alternatively be generated by a negative aggregate first moment shock that would move the frontier *left* (i.e. up), or can be generated by a combination of negative first and positive second moment shocks.

On the empirical end, the paper uses data from the NBER-CES Manufacturing Industry Database and Compustat to document the presence of a positively sloped TFP mean-variance frontier; firms and industries that experience the greatest variation in their TFP growth rates across time, on average are also the ones that have the highest mean TFP growth rates. There exists a *positive* relationship between the first and second moment of the long-run TFP process. This

[1] This result is similar to Curdia and Reis (2012) who empirically make a more general case for correlated shocks. I add on to this case here by providing a model that explains how optimization by economic agents can lead to a subset of macroeconomic shocks being correlated.

is an interesting observation, because at the business cycle frequency the first and second moments of TFP are *negatively* correlated. This paper takes as exogenous the long positive relationship, by allowing firms to choose projects along a positively sloped mean-variance frontier, and provide a theory to explain how this positive relationship in the long run can manifest itself as a negative relationship between the moments of TFP at business cycle frequencies.

I primarily motivate movements along the empirically observed mean-variance frontier as resulting from managers making different project choices. Alternatively, the reader may also view movements along the frontier as a reduced-form equivalent to other within firm decisions that involve a mean-variance tradeoff. Examples of such micro-level decisions include Narita (2011) who constructs a model where higher uncertainty increases the agency problem. The worsened agency problem leads to less risk taking and the subsequent abandonment of risky projects at the firm-level. Arellano et. al. (2012) show that higher uncertainty can cause firms' to reduce the size of their projects at the micro-level. Panousi & Papanikolaou (2012) empirically document that managers tend to underinvest when faced with increased uncertainty. As a result, the movements along the frontier in this paper can be interpreted in many different ways. The main results of this paper are independent of the interpretation of exactly how the movements along the frontier occur. The results only depend on the presence of a positively sloped mean-variance frontier and the risk-return tradeoff movements along it induce.

With respect to the literature, my paper is related to a recent set of papers that aim to understand the relationship between economic aggregates and changes in uncertainty. These papers can be roughly divided into two strands. First, a growing literature, studies how exogenous second moment shocks can have first moment effects on economic aggregates.[2] The basic neoclassical model is not particularly conducive to studying the effects of second moment shocks. For small values of household risk aversion the effects of second moment shocks are negligible in the basic neoclassical model, while for high levels of risk aversion (as in Basu & Bundick (2012)) the effects are not negligible, but then investment and consumption do not comove in response to these shocks - a result which is at odds with the data where there exists strong comovement between consumption and investment. Basu & Bundick (2012) provide an explanation for why the incorrect comovement occurs in the basic neoclassical model. It should be noted here that many models in this literature inherit the negative comovement problem and are either unable to solve it or have to appeal to nominal rigidities or additional real rigidities to generate comovements between consumption and investment. The model in this paper generates this comovement without appealing to any addi-

[2]Papers in this category include: Basu & Bundick (2012), Bloom (2009), Bloom et. al. (2012), Leduc & Liu (2012), Mericle (2011), Narita (2011), and Schaal (2012) who all look at the role of various nominal and real rigidities in transmitting second moment shocks to business cycles. Arellano et. al. (2012), Christiano et. al. (2013), Dorofeenko et. al. (2008), Gilchrist et. al (2010), and Fernandez-Villaverde et. al. (2011) who focus primarily on rigidities and frictions in the financial channel as a way for second moment shocks to have first moment effects. Bernanke(1983), Bond & Cummin (2004), Guiso (1999), and Lee (2013), who model how uncertainty affects the level of capital and investment.

tional rigidities; the comovment between consumption and investment is a natural byproduct of the mechanism in this paper.

The timing of factor input decisions also matters in the basic neoclassical model. If factor input decisions are made after the uncertainty is resolved, for example in the basic model labor decisions are made after observing the current period TFP level, then higher uncertainty results in an expansion - a result which is once again at odds with the data where higher uncertainty is associated with a recession. The reason for this is that for a decreasing return to scale Cobb-Douglas production, function factor input choices, if made after the uncertainty is resolved, are convex in the TFP level. This result is often referred to as the Oi (1961), Hartman (1972) and Abel (1983) effect. Bloom et. al. (2012) and Bloom (2014) provide a more thorough explanation of this effect.

Bloom (2014) provides a survey of papers in this first strand of the literature. In his survey, he further subdivides the first strand into papers that use the real option effect vs. those that use risk aversion and changes in risk-premia to explain how second moments can matter.[3] My paper falls in the second sub-category. In my model, preference-level risk aversion plays a crucial role in movements along the mean-variance frontier. This is consistent with Bachmann & Bayer (2013) who find evidence that the real option effect may not be an important source of propagation for second moment shocks.

A few papers in this first strand of the literature also link uncertainty shocks to fluctuations in the first moment of TFP. For example, Gilchrist et. al (2010) and Bloom et. al (2012) in their models generate a negative correlation between the first and second moment of aggregate TFP. In both cases the negative correlation is due to a reallocation of inputs across firms with heterogeneous productivity levels. This is in contrast to my model, where the negative correlation is mainly due to firm level project/TFP choices. On the empirical end, Bachmann & Bayer (2013) find VAR evidence that high uncertainty predicts low future TFP. This is closely related to evidence in my model where due to endogenous TFP choices there exists a negative correlation between uncertainty and *current* TFP, and due to the persistence of these processes there also exists a negative correlation between uncertainty and *future* TFP.

The second major strand of the literature includes papers that show how exogenous shocks can endogenously generate changes in uncertainty. Specifically for endogenously generated TFP uncertainty, both Kehrig (2012) and Tian (2011) construct models where exogenous first moment shocks can lead to changes in across-firm TFP dispersion.[4] This paper contributes to this strand of the literature as well. In my model, aggregate second moment shocks endogenously generate

[3] Bloom (2014) in addition to these two categories also talks about the effects of uncertainty on long-run growth through the "growth option" effect and "Oi-Hartman-Abel" effect.

[4] Other examples of endogenously generated countercylical uncertainty include Bachmann & Moscarini (2012) who show how price volatility can increase in bad times, Fostel and Geanakopolos (2012) who explain why firms invest in more volatile technologies during recessions, Van Nieuwerburgh and Veldkamp (2006) who model the endogenous countercyclicality of dispersion in learning, and Decker, D'Erasmo, and Boedoet (2014) who show how market entry/exit decisions can generate output dispersion within a firm.

changes in across-firm TFP dispersion.

The rest of this paper is organized as follows: Section 2 presents empirical evidence for the existence of a mean-variance frontier and document some stylized facts. In section 3, I present my model and explain my calibration. Section 4 presents my results. Section 5 concludes.

2 Data

This section first presents three stylized facts that the quantitative model will explain. It will then document the existence of an upward sloping mean-variance frontier.

2.1 Stylized Facts

This subsection documents 3 key stylized facts about TFP data:

I. Real GDP and the second moment of the aggregate TFP process are negatively correlated.

Panel (a) of Figure 2 plots the second moment of the TFP process vs. HP-filtered Real GDP. TFP here is measured as the Solow residual and the second moment of TFP is measured as the 8 quarter rolling standard deviation of TFP growth.[5] As can be seen GDP is negatively correlated with the second moment of TFP. As robustness I also calculate an alternative measure of the second moment of TFP, the conditional heteroskedasticity derived from a GARCH(1,1) model of TFP growth, and find that this measure is also negatively correlated with GDP. The correlation for this alternative measure is -0.41 with statistical significance at the 1% level.[6] Finally, it should be noted that this negative correlation, along with the countercyclical nature of other measures of macroeconomic uncertainty, have been previously documented in the literature.[7] The negative correlation between GDP and the second moment of TFP is often taken as evidence of second moment shocks - uncertainty shocks - as being responsible for a fraction of business cycle fluctuations.

II. The first and second moments of the aggregate TFP process are negatively correlated.

Panel (b) of Figure 2 plots the 8 quarter moving average of TFP growth, a measure of the first moment of the TFP process, against the 8 quarter rolling standard deviation of TFP growth. The two are negatively correlated. Additionally, similar to before, as robustness I also calculate the correlation between the conditional heteroskedasticity derived from a GARCH(1,1) model of TFP growth, as my measure of the second moment of TFP, and the 8 quarter moving average of TFP

[5]The TFP Series is taken from John Fernald's Website. The full data sources and explanation can be found in the footnote to the figure.

[6]This alternative measure of the second moment has been previously used by Bloom et. al. (2012).

[7]Bloom (2014) provides a survey of this evidence.

growth and find that these two are also negatively correlated. The correlation is -0.465 with statistical significance at the 1% level. Times of high TFP uncertainty in the data are also times of low TFP growth. Taken together with the observation of a positively sloped TFP mean-variance frontier this observation paints an interesting picture of the relationship between the first and second moments of TFP. In the time series, a higher second moment of TFP is correlated with a lower-mean level of TFP. However, in the long run cross-section, industries that exhibit higher second moments of TFP also have higher mean levels of TFP. The correlation between the first and second moments of TFP is negative in the time series but positive in the long run cross-section.

III. The cross-sectional standard deviation of TFP across firms (i.e. cross-sectional dispersion of TFP) increases during recessions.

There is a large literature that documents the fact that the cross-sectional dispersion (or variance) of TFP across firms is countercylical. This literature includes results by Bloom et. al. (2012) who use the Census of Manufacturers and the Annual Survey of Manufacturers to show that dispersion of TFP growth rates for U.S. firms is countercylical. Bachmann and Bayer (2013) show the same countercylical dispersion result is true for German firms using the USTAN data base from Deutsche Bundesbank. Kehrig (2011), similar to Bloom et. al (2012), also uses the Census of Manufacturers and the Annual Survey of Manufacturers and shows the presence of countercyclical dispersion for the level of TFP (vs. TFP growth). Tian (2012) uses the NBER-CES Manufacturing Industry Database to show that the countercyclicality of TFP dispersion is also present at higher levels of industry aggregation. Eisfeldt and Rampini (2006) use multifactor productivity data from the Bureau of Labor Statistics and show that dispersion in this measure of TFP is also countercyclical. Table 1 I summarizes the various estimates of the correlation between TFP dispersion and output from these papers.[8]

A more recent literature, which includes Kehrig (2011) and Salgado et. al. (2015), further investigate the countercyclical nature of cross-sectional dispersion. They find that during recessions an expanding lower tail of the distribution of a number of firm-level variables - such as TFP, employment, and profits - is partly responsible for the increased cross-sectional dispersion. These papers thus conclude, that while cross-sectional dispersion is countercyclical, the cross-sectional skewness of these firm-level variables is procyclical, and that these two facts are related to each other.

[8]A series of other papers also document the presence of countercyclical dispersion for variables that are closely related to TFP. Some of these variables include industry growth rates in Bloom (2014) and Davis (2007), productivity growth in Gourio (2008), forecast uncertainty in Bachmann et. al. (2013), profitability in Chugh (2014), growth rate of sales in Bloom et. al (2012), Doepke & Weber (2006) and Doepke et. al (2005), growth rates of GDP in Higson (2002) and Higson (2004), price in Berger and Vavra (2011), and equity prices in Gilchrist et. al (2010).

2.2 The Mean-Variance Frontier

The key assumption of this paper is that there exists a mean-variance frontier in TFP. This section presents data and regressions to provide evidence for the existence of this frontier.

Figure 3 plots the mean and standard deviation of TFP growth across the years 1958 to 2005 for 86 different industries in the U.S. manufacturing sector.[9] TFP growth is measured as the growth in the Solow residual. Panel (a) presents the data for all 86 industries and panel (b) presents the data without the outliers. Each of these two plots illustrates a positive relationship between the mean and standard deviation of TFP growth at the industry level; industries with higher TFP growth rates on average also experience higher variance in TFP growth.

To statistically verify the observations in Figure 3, I run a regression of mean TFP growth on the standard deviation of TFP growth. Column (1) in Table 2 presents the results of this regression. The coefficient on the standard deviation of TFP growth is positive and statistically significant at the 5 percent level. A 1 percent increase in the standard deviation of TFP growth in a 4-digit manufacturing industry is associated with an approximately 0.57 percent higher mean TFP growth in the same industry.

Next, to ensure robustness of the result in column (1), in columns (2)-(9) I run a number of regressions with alternative specifications.[10] The baseline regression in column (1) included an intercept term to capture the presence of any aggregate uncertainty that cannot be reduced by moving along the mean-variance frontier. Column (2) drops the intercept term and assumes that all uncertainty can be eliminated by moving along the mean-variance frontier. The coefficient for the standard deviation of TFP growth in this new regression, without the intercept, remains positive and statistically significant.

The regression in column (1) includes all 86 industries and weighs them equally. As Figure 3 illustrates the data does include outliers and the industries vary greatly in size. Column (3) presents a variant of the baseline regression where each observation is weighted by the size of the industry, as measured by the industry's value added and outliers are dropped. For this alternative specification the standard deviation of TFP growth remains positive and statistically significant.

The main motivation for using the 4-digit classification for my baseline regression is that the SIC to NAICS transition in 1997 has led to a discontinuation in data availability for some industries at the 5- and 6-digit level before or after 1997. Column (4) presents the main regression with data at the 6-digit level (industries with missing data are dropped). The coefficient for the standard

[9]TFP data is from the NBER-CES Manufacturing Database (Variable: dtfp4). The 86 manufacturing sector industries represent 4-digit NAICS industries with the first two digits being 31-33. TFP growth is aggregated to the 4-digit level by calculating a weighted average of TFP at the 6-digit level. The value added's for each industry at the 6-digit level are used as weights.

[10]In addition to the robustness regressions presented here, I also run a number of other robustness regressions that include splitting the samples for before and after the great moderation period, and directly controlling for the value added of the industry in the NBER Sample. In each of these additional robustness regressions the coefficient for the standard deviation of TFP growth is always positive and statistically significant.

deviation of TFP growth continues to remain positive and statistically significant. This check provides evidence that the positively sloped mean-variance frontier is a robust feature at lower levels of aggregation as well.

To provide further evidence of the mean-variance frontier at lower levels of aggregation, in column (5) I report the results of the baseline regression using firm-level TFP from the Compustat Database.[11] Similar, to the baseline result in column (1), the coefficient on the standard deviation of TFP growth is positive and statistically significant. A 1 percent increase in the standard deviation of TFP growth at the firm-level is associated with an approximately 0.44 percent higher mean TFP growth in the same firm.

Unlike the NBER sample, the Compustat sample is not balanced and includes firms of all ages. In column (6) I only include firms with an age of at least 10 years (i.e. firms that have existed in the Compustat sample for 10+ years) and find that the positive sloped frontier is a robust feature of this alternative sample as well.

Finally, column (7) and (8) present regression specifications identical to (5) and (6), however, now I use the mean and standard deviation of the TFP level (vs. growth rates) from Compustat. As can be seen the mean-variance frontier continues to remain positively sloped in levels as well. [12]

3 Model Economy

The model economy consists of risk-averse managers and a representative household.

3.1 Risk-Averse Managers/Firms

Production

The model economy is populated by a unit mass of managers that are heterogeneous in their level of risk aversion. Each of these managers operates a small firm on behalf of the household. As managers, they have complete control over all production decisions for the firm which include choosing how much labor to hire, how much capital to rent, and what projects to undertake. These production decisions in turn determine the amount of output $y_{i,t}$, and thus profit $\pi_{i,t}$, that is generated by each manager-firm pair i:

$$\pi_{i,t} = y_{i,t} - w_t l_{i,t} - r_t k_{i,t} \tag{1}$$

[11] The firm level TFP measures are from Ayse Imrohoroglu's website and were used in Imrohoroglu and Tuzel (2013). These measures do not include any time or industry level effects. I convert their measures into TFP growth rates for my regression.

[12] I should note though that firm-level data includes industry or firm specific time trends which cannot be separately identified from firm or industry specific TFP. For the data used in columns (7) and (8) the process of detrending removed the effects of industry specific TFP making the results of these last two regressions biased.

with

$$y_{i,t} = \epsilon_{i,t} k_{i,t}^{\alpha} l_{i,t}^{\eta} \tag{2}$$

where $k_{i,t}$ and $l_{i,t}$ give the level of capital and labor employed by the i^{th} manager-firm pair, $\epsilon_{i,t}$ is the TFP level of the project choice, and w_t and r_t are the wage and rental rate respectively.

The manager's project choice determines the distribution from which the TFP level (or technology level), $\epsilon_{i,t}$, for his firm is drawn. The various project choices are given along a linear mean-variance frontier and are summarized by the managers choice of the variable $\xi_{i,t}$. A manager that chooses $\xi_{i,t}$, draws his TFP level, $\epsilon_{i,t}$, from a log normal distribution with mean level of $(\mu_t + \xi_{i,t})$ and a standard deviation of $(a_t + b\xi_{i,t})$. As a result, the manager faces a trade-off: choosing a high $\xi_{i,t}$ leads to a higher mean TFP level, and thus more output, but this increased output comes at the cost of higher risk. The managers individual risk preference determine his optimal $\xi_{i,t}$ choice.

Formally,

$$\epsilon_{i,t} = e^{\tilde{\mu} + \tilde{\sigma}\omega_t} \tag{3}$$

where $\tilde{\mu}$ and $\tilde{\sigma}$ are function of $\xi_{i,t}, \mu_t, a_t$ and b such that mean and standard deviation of $\epsilon_{i,t}$ are:[13]

$$E(\epsilon_{i,t}) = \mu_t + \xi_{i,t}$$

$$\text{Standard Deviation}(\epsilon_{i,t}) = a_t + b\xi_{i,t}$$

with

$$log(a_t) = \rho_a log(a_{t-1}) + (1 - \rho_a)\, log(\bar{a}) + \epsilon_a \text{ with } \epsilon_a \sim N(0, \sigma_a) \tag{4}$$

$$log(\mu_t) = \rho_\mu log(\mu_{t-1}) + (1 - \rho_\mu)\, log(1) + \epsilon_\mu \text{ with } \epsilon_\mu \sim N(0, \sigma_\mu) \tag{5}$$

$$\omega_t \sim i.i.d\ N(0, 1) \tag{6}$$

where b is the slope of the mean-variance frontier, a_t gives a measure of the intercept of the mean-variance frontier, μ_t is a traditional aggregate first moment shock, and ω_t determines the level of TFP shocks. It is instructive to think of both the mean and standard deviation of a firm's TFP process as each having two components: (1) An aggregate component, μ_t for the mean and a_t for the standard deviations; changes in which effect all manager-firm pairs equally. (2) A firm-specific component, or semi-aggregate component, $\xi_{i,t}$ for the mean and $b\xi_{i,t}$ for the standard deviations; managers can adjust this component but face a mean-variance tradeoff.

This paper will mainly be interested in how shocks to the aggregate component, a_t, affect the model economy. These shocks are often referred to as uncertainty shocks.[14] Consistent with the

[13]Note that $\tilde{\mu} = ln(\mu_t + \xi_{i,t}) - \frac{1}{2} ln\left(\left(\frac{a_t + b\xi_{i,t}}{\mu_t + \xi_{i,t}}\right)^2 + 1\right)$ and $\tilde{\sigma} = \sqrt{ln\left(\left(\frac{a_t + b\xi_{i,t}}{\mu_t + \xi_{i,t}}\right)^2 + 1\right)}$.

[14]To see why a_t is analogous to a traditional aggregate uncertainty shock fix the firm-specific choice $\xi_{i,t} = \xi$ and

literature, equation (4) models this uncertainty shock as an AR(1) process. For completeness, as discussed previously, it should be noted that shocks to a_t can also be viewed as non-traditional first moments shocks that alter the intercept of the mean-variance frontier. This is in contrast to a traditional first moment shock given by μ_t.

Manager

Managers in this economy run the firms on behalf of the household and as payment for these services they are entitled to a fraction, λ, of the profits, $\pi_{i,t}$, generated by the firm. The manager-firm pairs are single period lived with their period utility maximization problem given as follows:[15,16]

$$\max_{k_{i,t},l_{i,t},\xi_{i,t}} E_t U_i^M \left(C_{i,t}^M, L_{i,t}^M\right) \tag{7}$$

where consumption, $C_{i,t}^M$, is equal to the managers total compensation:

$$C_{i,t}^M = \lambda \pi_{i,t}, \tag{8}$$

and his labor, $L_{i,t}^M$, is given by:

$$L_{i,t}^M = \bar{L^M} + h(\xi_{i,t}, \bar{\xi_i}). \tag{9}$$

The manager spends a base fixed amount of time $\bar{L^M}$ providing managerial know-how to the firm. This base amount of time increases if he chooses to alter his firm's TFP choice, $\xi_{i,t}$, from its long-run steady state level, $\bar{\xi_i}$. This second effect on his labor supply is given by the function $h(\xi_{i,t}, \bar{\xi_i})$.[17] I choose $h(a,b) = [\tau(a-b)]^2$. Intuitively, I assume that manager-firm pairs specialize in a particular type of project, $\bar{\xi_i}$, and moving away from this type of project is costly. The magnitude of the cost depends on how different the project is with the difference being measured in how differently risky the project is, $\xi_{i,t} - \bar{\xi_i}$. This assumption ensures that the model does not generate large time-series swings in TFP choices. This adjustment cost is akin to the more familiar convex adjustment costs on factor inputs, such as investment adjustment costs and labor adjustment costs.[18]

Figure 4 illustrates the timeline of managerial decisions: The manager first observes the current state of the aggregate shocks, a_t and μ_t, he then makes a TFP choice, $\xi_{i,t}$, following which he

---

$\mu_t = 1$, then $\epsilon_{i,t} = \epsilon_t$ is a stochastic process with mean $1 + \xi$ and standard deviation $a_t + b\xi$. Changes in a_t now map directly to changes in the standard deviation of the aggregate TFP process.

[15] I assume the manager-firm pair is single period lived to make the problem tractable. This assumption makes the manager's problem static, which in turn makes the state space for the problem finite. The other positive aspect of this assumption is that when analyzing the model the within period risk effect is not be confounded by the across period precautionary saving motive.

[16] Alternatively, I could have assumed that managers are more impatient than the households causing their across-periods savings to be zero, or that managers do not have access to an intertemporal saving market. Both of these assumptions would lead to identical end results.

[17] Here, $h(x,x) = 0$, $h(x+y,x) > 0$ if $y \neq 0$, $\frac{\partial h(x+y,x)}{\partial y} > 0$, and $\frac{\partial^2 h(x+y,x)}{\partial y^2} \geq 0$.

[18] In section 4.3.1, I provide robustness exercises with respect to this assumption.

observes the TFP, $\epsilon_{i,t}$, and subsequently decides on the level of capital, $k_{i,t}$, and labor, $l_{i,t}$, to employ at his firm.

Given this timeline of events the manager's optimization decision yields familiar first order conditions with respect to his capital and labor decisions. The manager optimally sets the marginal product of each input equal to the spot price of that input:

$$w_t = \eta \epsilon_{i,t} k_{i,t}^{\alpha} l_{i,t}^{\eta-1} \tag{10}$$

$$r_t = \alpha \epsilon_{i,t} k_{i,t}^{\alpha-1} l_{i,t}^{\eta}. \tag{11}$$

These equations can be simplified further to calculate the capital and labor demand for each manager-firm pair:

$$k_{i,t}^{*} = \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t}{\eta}\right)^{\frac{-\eta}{1-\alpha-\eta}} \left(\frac{r_t}{\alpha}\right)^{\frac{\eta-1}{1-\alpha-\eta}} \tag{12}$$

$$l_{i,t}^{*} = \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t}{\eta}\right)^{\frac{\alpha-1}{1-\alpha-\eta}} \left(\frac{r_t}{\alpha}\right)^{\frac{-\alpha}{1-\alpha-\eta}} \tag{13}$$

Whereas the labor and capital decision are completely independent of the functional form of the utility function, the optimal project choice, $\xi_{i,t}$, is not. I must, thus, specify a functional form for utility. I assume that the managers' utility takes the form of standard King-Plosser-Rebelo (1988) preferences:

$$U_i^M(C,L) = \frac{\left[\left(Ce^{-\phi L}\right)^{1-\gamma_i}\right] - 1}{1-\gamma_i}$$

where ϕ is a scale parameter that compares the utility of consumption to the disutility of labor, and γ_i is a measure of risk aversion that varies across the different managers.

Given this utility function the optimal project choice problem can be greatly simplified and written as:

$$\max_{\xi_{i,t}} Be^{-\phi(1-\gamma_i)h(\xi_{i,t},\bar{\xi_i})} \frac{E_t\left[\left(\epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}}\right)^{1-\gamma_i}\right]}{1-\gamma_i} \tag{14}$$

where $\epsilon_{i,t}$ is given by equation (3). The B term includes the effects of wage and rental rate changes and is independent of the $\xi_{i,t}$ choice.[19]

I use numerical methods to solve for the optimal project choice in (14):[20,21]

$$\xi_{i,t}^{*} = \xi(\gamma_i, a_t, \mu_t). \tag{15}$$

[19] $B = \left(\lambda(1-\alpha-\eta)\left(\frac{\eta}{w_t}\right)^{\frac{\eta}{1-\alpha-\eta}}\left(\frac{\alpha}{r_t}\right)^{\frac{\alpha}{1-\alpha-\eta}} e^{-\phi \bar{L_M}}\right)^{1-\gamma_i}$

[20] The long-run steady state level, $\bar{\xi}_i$, is fully determined by γ_i and thus I do not include it as a separate argument of the optimal project choice function $\xi(\bullet,\bullet,\bullet)$

[21] I explain the numerical solution method in appendix A.

For the specific case of the baseline calibration, the optimal project choice function, $\xi(\gamma_i, a_t, \mu_t = 1)$, has three important properties:[22]

1. $\xi_{\gamma_i} \leq 0$

2. $\xi_{a_t} \leq 0$

3. $\xi_{a_t,\gamma_i} \leq 0$

$\xi_{\gamma_i} \leq 0$ indicates that a higher level of risk aversion leads to a lower $\xi_{i,t}$ choice; agents with a high level of risk aversion forgo drawing from a high *mean* distribution in order to be able to draw from a *low* risk distribution. $\xi_{a_t} \leq 0$ indicates that in times of high aggregate uncertainty the managers attempt to reduce risk by choosing a low $\xi_{i,t}$. A low $\xi_{i,t}$ choice reduces the firm-specific component of the risk at the cost of an on average lower TFP draw. At the extreme value of $\gamma_i = 0$, the manager is risk-neutral, and changes in uncertainty do not effect his choice with $\xi_{a_t}|_{\gamma_i=0} = 0$. Finally, the cross-partial $\xi_{a_t,\gamma_i} \leq 0$ indicates that the higher the level of risk aversion, γ_i, the more the manager adjusts $\xi_{i,t}$ downward in response to a shock to a_t.

To conclude, in this subsection I presented two key assumptions of the model: (1) managers have varying levels of risk aversion and (2) risk aversion plays a role in a manager's decision making. The finance literature on corporate governance provides empirical evidence for both these assumptions. For example, Graham et. al. (2013) in their paper find that there is heterogeniety in risk aversion across CEO's and that corporate decisions are affected by a CEO's risk preference.[23] Further, consistent with the results of this paper, Graham et. al. (2013) also find that firm's with higher rates of growth are on average run by low risk-averse CEO's (i.e. $\xi_{\gamma_i} \leq 0$).

3.2 Household

The model economy is also populated with a standard neoclassical representative household that maximizes its lifetime utility:

$$\sum_{t=0}^{\infty} \beta^t E_t U\left(C_t^H, L_t^H\right) \tag{16}$$

subject to the following period budget constraint

$$C_t^H + K_{t+1} = w_t L_t^H + r_t K_t + (1-\delta) K_t + \Pi_t \tag{17}$$

where C_t^H, K_t, and L_t^H give the levels of consumption, capital, and labor respectively, w_t and r_t are the wage and rental rate, δ is the depreciation rate, and Π_t is the total profits reimbursed by

[22] Note that these are numerical results and not analytical statements based on a proof

[23] Other papers that show that a managers attitude toward risk (either induce through a compensation package or inherent) matters for decision making include Guay (1999), Ross(2004), Low(2009), among others. Also, Buera et. al. (2011) and Moll (2014) are examples of papers in the economics literature that assume entrepreneurs are risk-averse.

the managers to the household.

The household, similar to the managers, also has King−Plosser−Rebelo (1988) preferences given by:

$$U(C,L) = \frac{\left[\left(Ce^{-\phi L}\right)^{1-\sigma}\right] - 1}{1-\sigma}$$

Here because the household makes intertemporal decisions, σ captures not only the risk preferences of the household, but the inverse of σ is also the intertemporal elasticity of substitution.

The resulting first order conditions for household are standard and can be simplified and written as:

$$U_{C_t^H} w_t = -U_{L_t^H} \tag{18}$$

$$U_{C_t^H} = \beta E_t \left[U_{C_{t+1}^H}\left(r_{t+1} + 1 - \delta\right)\right]. \tag{19}$$

3.3 Market Clearing and Aggregate Variables

In equilibrium, wages, w_t, and the rental rates, r_t, are such that the markets for labor and capital clear. These market clearing conditions are given as follows:

$$L_t^H = \int_{i\in[0,1]} l_{i,t}^* di = \int_{i\in[0,1]} \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t}{\eta}\right)^{\frac{\alpha-1}{1-\alpha-\eta}} \left(\frac{r_t}{\alpha}\right)^{\frac{-\alpha}{1-\alpha-\eta}} di \tag{20}$$

$$K_t = \int_{i\in[0,1]} k_{i,t}^* di = \int_{i\in[0,1]} \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t}{\eta}\right)^{\frac{-\eta}{1-\alpha-\eta}} \left(\frac{r_t}{\alpha}\right)^{\frac{\eta-1}{1-\alpha-\eta}} di \tag{21}$$

with $\epsilon_{i,t}$ is given by (3) and the market for goods automatically clears due to Walras' Law.

Also, in equilibrium the aggregate levels of capital and labor are given by K_t and $L_t = L_t^H$ in equations (20) and (21), and the aggregate levels of output, consumption, and profits reimbursed to the household, are calculated as follows:

$$Y_t = \int_{i\in[0,1]} y_{i,t} di = \int_{i\in[0,1]} \epsilon_{i,t} k_{i,t}^{*\alpha} l_{i,t}^{*\eta} di \tag{22}$$

$$C_t = C_t^H + C_t^M = C_t^H + \lambda \int_{i\in[0,1]} \pi_{i,t} di \tag{23}$$

$$\Pi_t = (1-\lambda) \int_{i\in[0,1]} \pi_{i,t} di \tag{24}$$

The level of aggregate TFP, the standard deviation of aggregate TFP, the cross-sectional dispersion of firm-level TFP, and the cross-sectional skewness of firm-level TFP are given as:

$$A_t = \frac{Y_t}{(K_t)^\alpha \left(L_t^H\right)^\eta} = \left[\int_{i\in[0,1]} \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} di\right]^{1-\alpha-\eta} \tag{25}$$

$$StDev(A_t) = \text{Standard Deviation of } A_t \text{ for } \omega_t \sim \ i.i.d.N(0,1) \tag{26}$$

$$Disp_t = \text{Coefficient of Variation of } \ (\mu_t + \xi_{i,t}) \ \text{ across all } i \in [0,1] \tag{27}$$

$$Skew_t = \text{Kelly's Skewness of } \ (\mu_t + \xi_{i,t}) \ \text{ across all } i \in [0,1] \tag{28}$$

3.4 Equilibrium

An equilibrium in this economy is described as a set of firm specific variables, $\left\{C^M_{i,t}, L^M_{i,t}, k^*_{i,t}, l^*_{i,t}, \mu^*_{i,t}, \epsilon_{i,t}\right\}$, a set of aggregate variables, $\left\{C_t, L_t, K_t, Y_t, \Pi_t, C^H_t, L^H_t, A_t, StDev(A_t), Disp_t, Skew_t\right\}$, and a set of prices $\{w_t, r_t\}$ such that for all i and t, equations (3), (8), (9), (12), (13), (15), (17), (18), (19), (20), (21), (22), (23), (24), (25), (26), (27), (28), and $L_t = L^H_t$ are satisfied.

3.5 Calibration

Table 3 presents the parameter values used in the baseline simulations.

The model is solved at a quarterly frequency. I set the household's subjective discount factor, $\beta = 0.985$, to generate an annual subjective discount factor of roughly 0.94. The intertemporal elasticity of substitution is set to a standard value of $\sigma = 1$.[24] I internally calibrate ϕ such that the household spends 0.25 of its time working in steady state. Finally, the depreciation rate is set at 0.025 to generate an annual capital depreciation rate of 10%.

For the firm I set the capital and labor shares to $\alpha = 0.25$ and $\eta = 0.5$. This is consistent with the observation that the ratio of labor to capital share in data is approximately 2 and the average marginal cost markup in the model economy, consistent with estimates in the data, is approximately 33%.

Nikolov & Whited (2014) find that the managers' equity in the firm is between 2.5% and 7.8%, as a result I set the manager's profit share, λ, at an average value of 0.05. [25] Consistent with the household, I set $\bar{L^M} = 0.249$ to ensure the manager spends 25% of his time working. Finally, since τ is a free parameter I choose an intermediate value of $\tau = 5$ and then conduct sensitivity analysis.

I calibrate the mean-variance frontier by setting $\bar{a} = 0.0265$ and $b = 1.75$. These values correspond to the empirical mean-variance frontier in column (1) of Table 2. Given the heterogeneity in the empirical slope values in Table 2 and that firm-specific and/or industry-specific trends may bias the empirical mean-variance frontier estimations, in the robustness section I provide sensitivity analysis with respect to alternate slope values for the frontier. [26]

[24] If $\sigma = 1$ the utility function simplifies to log utility in consumption with Hansen(1985)-Rogerson(1988) preferences in labor, $U(C,L) = log(C) - \phi L$

[25] It should also be noted that the main results are relatively independent of exact value of λ.

[26] There are various issues with directly using the regression results to estimate the mean-variance frontier in the model. For example, firm-level data on TFP levels has firm and industry-specific time trends. Detrending such data to remove only the aggregate time-trend would not account for these firm and industry specific time trends. Alternatively, detrending to also remove firm and industry-specific time trends would necessarily remove the effects of firm/industry specific TFP, as the two are not separately identifiable. As a result, using data from either of these two

The remaining variables are calibrated internally. First, I calibrate $\sigma_\mu = 0.00435$, $\rho_\mu = 0.91$, $\sigma_a = 0.205$, and $\rho_a = 0.92$, such that the standard deviation and autocorrelation of the first and second moments of aggregate TFP in the data match those generated from simulations of the model.

Second, I choose the distribution of risk aversion across managers to be a truncated normal.[27] To internally calibrate a truncated normal I need to specify four statistics of the distribution, the minimum, the maximum, the mean, and the standard deviation. I choose these four statistics to generate a set of steady state TFP choices by the managers such that the minimum, maximum, mean, and standard deviation of the cross-sectional distribution of TFP's across the firms in the model match the 5^{th} percentile, 95^{th} percentile, mean, and standard deviation of the distribution of TFP growth rates across the 86 4-digit industries in the NBER-CES Manufacturing Industry Database. The 5^{th} percentile and 95^{th} percentile are used as target moments to account for any outliers. I should note that the mean level of risk aversion, $\gamma_{Mean} = 5$, calibrated in this way is consistent with other estimates of this parameter. For example, Guvenen (2009) estimates the risk aversion parameter in Epstein-Zin preferences (1989) to be approximately 6.

4 Results

This section begins by discussing how the model can explain both the 3 stylized facts and various business cycle moments. It then briefly explains how the model highlights issues with measuring first and second moment TFP shocks directly from the data. It closes with a number of robustness exercises.

For the results in this section, I solve the model by first simplifying it and rewriting it as a value function. I then use function iteration to find a numerical estimate of this value function. In turn, I use the resulting value function estimate to create policy rules which help me both construct impulse responses and calculate the moments for key variables. The exact details of the solution method and how the impulse responses and moments are calculated can be found in appendix A.

Also, this paper is mainly interested in the effects of uncertainty shocks, and as a result, unless stated otherwise, the moments in this section are calculated to only measure the effects of changes in aggregate uncertainty, a_t, and the impulse responses presented are for a two standard deviation shock to aggregate uncertainty, a_t, only. For completeness, however, Appendix B provides impulse responses to a two standard deviation shock to μ_t.

detrending procedures leads to biased estimate of the mean-variance frontier in detrended TFP levels. Regressions with TFP growth levels can also suffer from such issues

[27] This choice allows the distribution to be a bell-shaped curve with a well-defined minimum and maximum

4.1 Main Simulation

Figure 5 presents the impulse responses of key variables in the model to a 2 standard deviation (+2 std) shock to aggregate uncertainty, a_t. As can be seen in the figure an aggregate uncertainty shock causes the standard deviation of the aggregate TFP process to rise, while the aggregate output and mean of the aggregate TFP process fall.

In the model, higher aggregate uncertainty induces the risk-averse managers to endogenously move along the mean-variance frontier by altering their project choice to reduce the firm-specific component of uncertainty. Due to the positively sloped nature of the mean-variance frontier the reduction in the firm-specific component of uncertainty comes at the cost of a lower-mean TFP process at each firm. The lower-mean TFP process at the firm level aggregates to a lower aggregate mean TFP (panel (d) of Figure 5). Also, the lower-mean TFP at each firm reduces production at the firm level, which in turn aggregates to lower aggregate output in the economy (panel (b) of Figure 5). In this way endogenous movements along the mean-variance frontier cause uncertainty shocks to have a negative effect on both output and the mean level of TFP.

To highlight the importance of endogenous movements along the mean-variance, in Figure 5 I also plot the impulse responses for a model identical to the baseline model, but one where managers cannot endogenously choose projects and move along the mean-variance frontier. In this latter model, the managers' firm-specific component of uncertainty is fixed at the steady state level. Increases in aggregate uncertainty in this model cause the standard deviation of TFP to rise more than the baseline model because managers are unable to dampen the risk they face by moving along the mean-variance frontier. Also, if the managers cannot endogenously move along the frontier then the mean level of TFP remains constant, and with the mean level of TFP constant, output in this model does not fall. Instead, due to the Oi-Hartman-Abel effect output rises in this model. In the model, factor input decisions are made after the uncertainty is resolved (see Figure 4) and as a result for a fixed wage rate and rental rate the optimal labor and capital levels at the firm-level, equations (12) and (13), are convex in the level of TFP, $\epsilon_{i,t}$. This convexity in the TFP level causes capital and labor, and thus output, to rise in times of high aggregate uncertainty. A positive correlation between output and uncertainty is at odds with the data (see Figure 2).[28]

Table 4 presents the data moments and model generated moments for output and TFP for the baseline model. For the baseline model the table presents the moments for when the model is driven by just shocks to a_t (uncertainty shocks), just shocks to μ_t (traditional first moment shocks), and shocks to both a_t and μ_t. In the presence of uncertainty shocks both output and the mean TFP are negatively correlated with the second moment of the aggregate TFP process. These negative correlations coincide in the data with stylized facts 1 and 2. A model with no endogenous choice or a model that is just driven by traditional first moment shocks is unable to replicate these stylized

[28]Note that the Oi-Hartman-Abel effect is also present in the baseline model but is strictly dominated by the effect the changing TFP level has on the aggregate variables

facts, thus again showing the importance of the mean-variance frontier and endogenous choices in my framework.[29]

Next, due to the heterogeneity in risk aversions across the managers the effects of high aggregate uncertainty are not uniform across the firms in my model. Panel (a) of Figure 6 plots the mean TFP choice at the steady state level, and for a ± 2 standard deviation shock in a_t, for each level of managerial risk aversion, γ_i. Panel (b) gives the distribution of risk aversions. It is clear that the responsiveness to changes in aggregate uncertainty is increasing in the risk aversion of the manager. A manager with high risk aversion is more likely to dampen the effect of higher uncertainty as compared to a low risk-averse manager. This heterogeneity causes the distribution of firm level TFPs to spread out during times of high uncertainty and to contract in times of low uncertainty. As a result, in the model the cross-sectional dispersion of firm-level TFP is positively correlated with the second moment of aggregate TFP. As seen in panel (c), during recessionary times the cross-sectional dispersion of TFP increases. This is the third stylized fact. It should be noted that Kehrig (2011) documents that most of the countercyclical dispersion in TFP is driven by changes in the TFP for firms at the low end of the TFP distribution, causing cross-sectional skewness to be procyclical. My results are consistent with this additional stylized fact as well. Panel (d) of Figure 6 presents the impulse response of cross-sectional skewness to an uncertainty shock.

On the broader business-cycle facts: Figure 7 plots the impulse responses of consumption, labor for the household, and investment, to a $+2$ std shock to aggregate uncertainty, a_t. Consumption, labor for the household, and investment all fall in response to an increase in aggregate uncertainty.[30] Increases in aggregate uncertainty thus do generate an economy wide recession in my model. This is not true of all models in this literature; a number of models when driven by uncertainty shocks are unable to generate the empirically observed positive comovement between consumption and investment. Furthermore, unlike the model of this paper, models that do generate comovement in response to uncertainty shocks often need to rely on additional rigidities. The correlation of consumption and investment for my baseline model when driven by uncertainty shocks is 0.5, as compared to approximately 0.7 in the data.

In addition to generating the correct qualitative responses, the model is also able to generate relatively sizable business cycle fluctuations. Table 5 compares the data and model generated moments. When the model is driven by both uncertainty shocks (a_t) and traditional first moment shocks (μ_t) the model is able to explain a large portion of the variability in output, consumption, investment, and labor hours. A variance decomposition exercise shows that shocks to a_t are able to explain roughly a quarter of the variability in output, consumption, investment, and labor hours in

[29] In addition, because of the persistence of the TFP process, the model generates a negative correlation of 0.62 between the second moment of TFP *today* and first moment of *future* TFP (in 4 periods); this result is in line with VAR evidence from Bachmann & Bayer (2013) that high uncertainty today predicts low TFP in the future.

[30] I should note that as seen in panel (d) of Figure 7 the managers' hours in the model do indeed increase in response to higher aggregate uncertainty. However, this change is not important for the models aggregate implications because in the aggregate managers make up a very small portion of the total economy

the model with the remaining variability explained by shocks to μ_t. As a result, the model predicts that uncertainty shocks through endogenous movements along a mean-variance frontier can explain a significant portion, roughly a quarter, of the U.S. business cycle fluctuations. The model also does a good job of matching the relative standard deviations, autocorrelations, and correlations with GDP of the main aggregate variables.

Tables 4 and 5 also provide a comparison between uncertainty shocks (a_t) and traditional first moment shocks (μ_t). From Table 5, both shocks contribute to the observed variance in the *real* variables - output, consumption, investment, and labor hours. However, as seen in Table 4, this is not true for the variation in TFP moments. Whereas second moment shocks (a_t) drive part of the fluctuations in the observed first moment of TFP, traditional first moment shock (μ_t) have a negligible effect on fluctuations in the second moment of TFP. Furthermore, as seen in Appendix B, first moments shocks also have a very small effect on cross-sectional dispersion and skewness. This heterogeneity in responses allows me to disentangle and separately account for the two shocks. [31]

4.2 Identification of First and Second Moment TFP Shocks in the Data

The endogenous nature of firm-level project choices, and thus TFP choices, help sheds light on measurement issues related to (1) identifying first moment shocks in the data and (2) measuring the magnitude of second moment shocks.

The total variance of the mean level of TFP in the model is $(1.3\%)^2$ which decomposes into $(0.6\%)^2$ coming from shocks to a_t and $(1.1\%)^2$ coming from shocks to μ_t (see Table 4). As a result, roughly a quarter of the fluctuations in the mean level of TFP are purely endogenous and driven by uncertainty shocks that cause managers to reoptimize along the mean-variance frontier. This result highlights a problem with using raw moments of the TFP data to identify exogenous TFP shocks. As in the model, in the data it is possible that part of the observed movements in aggregate TFP are actually endogenous in nature, and not indicative of exogenous first moment shocks. Furthermore, this result also highlights the possibility that the observed negative correlation between the time series of the first and second moment of TFP is not coincidental, but structural in nature. Therefore, it is important that this sort of endogeneity - a structural negative correlation - be allowed for in models that attempt to jointly estimate the effects of first vs. second moment shocks.[32]

Next, endogenous TFP choices by managers along a mean-variance frontier act to dampen the effect of increased aggregate uncertainty and magnify the effect of low aggregate uncertainty. Figure 8 plots how different levels of aggregate uncertainty manifest themselves in the model. In

[31] It should, however, be noted that whereas traditional first moment shocks have negligible effects on the higher moments, in so much as a shock to a_t, the intercept of the mean-variance frontier, can be viewed of as a type of first moment shock, some types of first moment shocks can have large effects.

[32] This result is in line with Bachmann and Bayer (2013) who find that allowing for correlations between first and second-moment shocks allows their model to better explain the data.

times of both relatively low and relatively high uncertainty the actual shock to the economy is quite different, than the shock measured from the data. This shows that using second moments of TFP data to measure second moment shocks can lead to a systematic mismeasurement of the true magnitude of these shocks. It is important to correctly measure the magnitude of the true underlying second moment shock because a mismeasured shock can cause models to erroneously conclude the lack of, or over-importance of, second moment shocks in the data. For example, model firms run by low risk-averse managers face the full magnitude of the uncertainty shock, and thus any mismeasurement of the true magnitude of the uncertainty shock will be unable to explain the behavior of these firms. In the model, the true magnitude of second moment shocks predicts important heterogeneities in uncertainty faced by the cross-section of firms which a mismeasured shock might miss.

In conclusion, the model highlights issues with naively measuring first and second moment shocks directly from TFP data.

4.3 Robustness

In this subsection, I perform a number of robustness exercises with respect to the calibration and modeling assumptions.[33]

4.3.1 Adjustment Costs

In calibrating the baseline model I treated τ as a free parameter. Recall that τ parameterizes the convex labor cost a manager must pay to move along the mean-variance frontier. Table 6 reports the real GDP and TFP moments for various values of τ. As can be seen for both when $\tau = 2.5$ (50% lower than baseline) and when $\tau = 7.5$ (50% higher than baseline) the second moment of TFP is negatively correlated with both output and the first moment of TFP - the first two stylized facts. Panel (a) of Figure 9 plots the changes in TFP dispersion in response to an uncertainty shock. For all three values of τ, the dispersion in cross-sectional TFP increases in response to a positive uncertainty shock. This illustrates the negative relationship between cross-sectional dispersion and GDP - the third stylized fact. Consequently, the model's ability to qualitatively replicate the three stylized facts is robust to small changes in τ. Further, as seen in Table 7 for these alternative τ values, uncertainty shocks in the model are still able to explain a large fraction of business cycle fluctuations.

The model, however, is unable to qualitatively replicate the stylized facts for extreme values of τ. When $\tau = \infty$ the model reduces to one where there are no endogenous choices, thereby eliminating the main channel through which increases in aggregate uncertainty have a negative effect on the

[33] The discussions in this section rely only on the model moments in Table 6 and 7 and the impulse responses of cross-sectional dispersion in Figure 9. For completeness, however, I do provide a full set of impulse responses in appendix B. Furthermore, for the sensitivity analysis in this section I provide partial derivative exercises in that I do not recalibrate the model after changing the parameter values or modeling assumptions.

mean level of TFP and output. On the other hand, when $\tau = 0$ the risk-averse manager is very responsive to changes in aggregate uncertainty. When $\tau = 0$, if aggregate uncertainty, a_t, increases the manager dampens the uncertainty so much so that he optimally chooses a TFP process that has both a lower-mean and lower-standard deviation than before. This effect can be illustrated as the medium risk-averse manager moving to H' instead of M' in Figure 1. As a result, with $\tau = 0$, even though an exogenous increase in aggregate uncertainty does cause a recession, due to the extreme endogenous movements the measured second moment of aggregate TFP *decreases* in response to an exogenous *increase* in aggregate uncertainty.

In summary, the model's main results are relatively robust as long as the value of τ is neither too close to 0, nor too large. It should additionally be noted that at lower values of τ uncertainty shocks in the model are able to explain a larger fraction of business cycle fluctuations in real variables (see Table 7). This is because as τ falls the marginal cost of moving along the mean-variance frontier falls, which causes the mean level of TFP, and thus real variables, to be more responsive to changes in aggregate uncertainty. The downside of the managers being more responsive is that at lower values of τ, the movements in the second moment of TFP are significantly endogenously dampened during recessions (and vice versa for expansions), thus decreasing the models ability to match the standard deviation of the second moment of aggregate TFP in data.

4.3.2 Risk Distribution

In the model, heterogeniety in managerial risk aversion plays a key role in generating countercylical dispersion and procyclical skewness of cross-sectional TFP. For the baseline results, I calibrate the managerial risk aversion (γ) distribution by matching the model simulated distribution of mean TFP growth rates to those in the 4-digit NBER-CES data. To understand how sensitive the results are to this calibration, I first look at how the main results change when I change the standard deviation of the γ distribution to 1.05 or 3.15, as compared to 2.1 in the baseline calibration. As can be seen in both Tables 6 and 7 the aggregate model moments are essentially unchanged by altering the standard deviation. The model is able to both qualitatively replicate the first two stylized facts and explain a large fraction of aggregate business cycle moments independent of the standard deviation of the managerial risk aversion distribution. Next, panel (b) of Figure 9 plots the response of the cross-sectional dispersion when the standard deviation of the γ distribution is 1.05 and 3.15. Once again, the results are essentially unchanged confirming robustness of the third stylized fact as well.[34] Additionally, I also look at, but do not report, how the main results change when the standard deviation is 0 (i.e. all the agents are identical in their risk aversion). In this extreme case the aggregate moments presented in Table 6 and 7, similar to when $\gamma_{Std} = 1.05$ and 3.15, are again unchanged from the baseline calibration. However, in this extreme case the model

[34]For completeness it should be noted that where as changing the standard deviation does not change the volatility of the cross-sectional moments, it does change the mean level

is unable to replicate changes in the cross-sectional moments because of the lack of heterogeneity in the managers.

Second, I look at how the results would change if the mean of the γ distribution was altered. I lower the mean of γ distribution to 50% of its original value, $\gamma_{Mean} = 2.6$. Tables 6 and 7 present the aggregate moments under this alternate calibration and panel (b) of Figure 9 plots the response of the cross-sectional dispersion. The main difference between the baseline calibration and when $\gamma_{Mean} = 2.6$ is that the managers are less risk-averse causing them to be less reactive to uncertainty shocks. Quantitatively this effect is very small. The standard deviation of all the variables are slightly lower, other than the standard deviation of the second moment of TFP which is marginally higher. The reason for why the standard deviation of the second moment of TFP is higher, is that because managers are less reactive the dampening effect on the second moment of TFP during recessions is weaker (and vice versa for expansions).

4.3.3 Insurance

For numerical simplicity, I assume that managers in my model do not have access to an intertemporal saving instrument. A drawback of this assumption is that in the absence of an intertemporal saving option managers are unable to smooth out their consumption over time, and are thus unable to reduce the effects of compensation risk on their consumption. To understand the importance of this assumption, I consider a case where managers have access to an ad-hoc insurance instrument that decouples their compensation risk from their project risk choice. In particular, I assume that if a manager chooses a project which results in their profits, π, having a variance of X, then the insurance instrument reduces the variance of their compensation (which is equal to $\lambda\pi$) to $0.75X$. This insurance instrument can be viewed as the managers intratemporally partially insuring each other's compensation stream.[35]

As can be seen in Table 6, under the partial insurance assumption the model continues to be able to generate a negative correlation between the second moment of TFP and both output and the first moment of TFP. The main difference between the partial insurance case and the baseline case is that managers are less reactive to uncertainty shocks. This, similar to the lower γ_{Mean} case above, results in lower model moments for all variables in Tables 6 and 7. The only exception to this is the standard deviation of the second moment of TFP which increases in the partial insurance case. This is because the lower responsiveness of the managers leads to less dampening of the standard deviation of the second moment of TFP during recessions (and vice versa for expansions). Also, panel (c) of Figure 9 plots the response of the cross-sectional dispersion in the partial insurance

[35]This intratemporal insurance scheme would work as follows: If a manager gets a good TFP draw, then he shares his compensation with another manager who gets a bad TFP draw, with the aim of reducing the overall risk each manager faces. In general such an insurance scheme is time-inconsistent as the manager who gets a good TFP draw has incentive to renege on his promise to share his compensation. However, because the managers in the model are infinitely lived a simple grim trigger strategy would make the scheme time-consistent.

case. There exists a negative relation between dispersion and GDP, however once again due to the lower responsiveness of the manager, the response of cross-sectional dispersion is significantly dampened. In summary, with partial insurance the model is still qualitatively able to explain all three stylized facts, even though the presence of partial insurance dampens the responses.

4.3.4 Mean-Variance Frontier

Table 2 presents the empirical mean-variance frontier for a number of different specifications. In particular, I estimate the mean-variance frontier for different levels of aggregation - from industry-level to firm-level. There is some amount of variation in the slope of the frontier across these specifications. As robustness I study how the results of the model change when I alter the slope to $b = 1$ and $b = 2.5$ (vs. the baseline value of $b = 1.75$). Tables 6 and 7 give the model moments and panel (d) of Figure 9 plots the impulse responses of cross-sectional dispersion. As can be seen the model is able to qualitatively replicate all three stylized facts for both alternative slope values.

Next, the mean-variance frontier need not be a linear function. As an alternative assumption I study a non-linear mean-variance frontier with the mean of $\epsilon_{i,t}$ still equal to $(\mu_t + \xi_{i,t})$, but the standard deviation increasing to $\left(a_t + b\xi_{i,t}^2\right)$. From the 4-Digit NBER-CES data I estimate $b = 8.12$ and $\bar{a} = 0.036$ for this alternative case. Tables 6 and 7 give the model moments for the non-linear mean-variance frontier. As in the baseline case the second moment of TFP is still negatively correlated with both GDP and the first moment of TFP - the first two stylized facts. Panel (d) of Figure 9 verifies the third stylized fact. The main difference between the linear frontier case and non-linear frontier case is that in the later the model moments, other than the standard deviation of the second moment of TFP, are dampened. I conjecture that the reason for this is similar to the partial insurance and lower γ_{Mean} cases discussed above. The non-linear slope increases the cost of adjustment making the managers less responsive to uncertainty shocks.

Tables 6 and 7 also compare the effects of different adjustment costs in the presence of a non-linear frontier. The main difference between the linear and non-linear frontier case is that in the latter case lower adjustment costs lead to less dampening of the standard deviation of the second moment of TFP. I conjecture that this is because the non-linear nature of the frontier itself acts as a convex cost on too much adjustment by the manager.

In summary, the critical assumption for this paper is the presence of a positively sloped mean-variance frontier. As long as the mean-variance frontier is positive sloped the model is able to replicate the three main stylized facts for both alternate linear slope values and a non-linear frontier.

5 Conclusion

This paper shows how changes in aggregate uncertainty - uncertainty shocks - can generate business cycle fluctuations. Uncertainty shocks in my model are able to account for a significant portion of the observed variation in aggregate GDP. The main innovation of the paper is that it documents the presence of a positively sloped mean-variance TFP frontier and and then constructs a model that allows agents to endogenously choose projects with different TFP processes along this frontier. In the model economy an increase in aggregate uncertainty induces the firms to adopt projects that exhibit lower firm-specific risk. However, low-risk projects along the positively sloped frontier are also low-return projects. As a result when firms adopts these project their production falls. In this way, increases in aggregate uncertainty in the model generate sizable recessions.

The model economy is also able to account for a number of stylized facts. At the aggregate level, the model generates a negative correlation between the second moment of TFP and GDP, and between the first and second moments of TFP. Cross-sectionally, the model generates increased TFP dispersion and decreased TFP skewness during recessions.

Finally, the paper highlights some measurement issues. These include: (1) the second moment of TFP in the data may be a systematically downward biased estimates of uncertainty shocks during recessions and an upward biased estimate during expansions; (2) Uncertainty shocks can generate TFP dynamics that mimic traditional first moment shocks, and thus changes in the level of TFP in the data may be driven by second moment shocks. This latter result gives evidence that first and second moments shocks in the data may be endogenously correlated, and thus any joint estimation of these shocks in the data must allow for this correlation.

With uncertainty shocks gaining a lot more interest as drivers of the business cycle, it is important to be able to correctly measure them, develop models that illustrate the channels through which they work, and to understand how the various aggregate and cross-sectional notions of TFP moments may be related. This paper sheds light on all these issues.

Bibliography

[1] Abel, Andrew, “Optimal Investment under Uncertainty,” *American Economic Review*, 73:228-233, 1983.

[2] Arellano, Cristina, Yan Bai, and Patrick Kehoe, “Financial markets and fluctuations in uncertainty,” *Federal Reserve Bank of Minneapolis Working Paper*, 2012.

[3] Bachmann, Rudiger, and Christian Bayer, “’Wait-and-See’business cycles?” *Journal of Monetary Economics*, 60:704-719, 2013.

[4] Bachmann, Rudiger, and Giuseppe Moscarini, “Business Cycles and Endogenous Uncertainty,” *Mimeo*, Yale University, 2012.

[5] Bachmann, Rudiger, Steffen Elstner, and Eric R. Sims, “Uncertainty and Economic Activity: Evidence from Business Survey Data,” *American Economic Journal: Macroeconomics*, 5:217-49, 2013.

[6] Basu, Susanto, and Brent Bundick, “Uncertainty shocks in a model of effective demand,” *National Bureau of Economic Research Working Paper 18420*, 2012.

[7] Berger, David, and Joseph Vavra, “Dynamics of the US price distribution,” *Mimeo*, Yale University, 2011.

[8] Bernanke, Ben S., “Irreversibility, Uncertainty, and Cyclical Investment,” *Quarterly Journal of Economics*, 98:85−106, 1983.

[9] Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin, “Finance and Development: A Tale of Two Sectors,” *American Economic Review*, 101:1964-2002, 2011.

[10] Bloom, Nicholas, “The impact of uncertainty shocks,” *Econometrica*, 77:623-685, 2009.

[11] Bloom, Nicholas, “Fluctuations in Uncertainty”, *Journal of Economic Perspectives*, 28:153−176, 2014.

[12] Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry, “Really uncertain business cycles,” *National Bureau of Economic Research Working Paper 18245*, 2012.

[13] Bond, Stephen, and Jason G. Cummins, “Uncertainty and investment: An empirical investigation using data on analysts’ profits forecasts”, *Mimeo*, Federal Reserve Board, 2004.

[14] Christiano, Lawrence, Roberto Motto, and Massimo Rostagno, “Risk shocks,” *National Bureau of Economic Research Working Paper 18682*, 2013.

[15] Chugh, Sanjay, “Firm Risk and Leverage-Based Business Cycles,” *Mimeo*, Boston College, 2014.

[16] Curdia, Vasco, and Ricardo Reis, “Correlated disturbances and US business cycles,” *National Bureau of Economic Research Working Paper 15774*, 2010.

[17] Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda, “Volatility and dispersion in business growth rates: Publicly traded versus privately held firms,” *NBER Macroeconomics Annual 2006*, 21:107-180, 2007.

[18] Decker, Ryan, Pablo N. D'Erasmo, and Hernan M. Boedo, "Market Exposure and Endogenous Firm Volatility over the Business Cycle," *Mimeo*, University of Maryland, 2014.

[19] Dopke, Jorg, Michael Funke, Sean Holly, and Sebastian Weber, "The cross-sectional dynamics of German business cycles: a bird's eye view," *Mimeo*, Deutsche Bundesbank, 2005.

[20] Dopke, Jorg, and Sebastian Weber, "The within-distribution business cycle dynamics of German firms," *Mimeo*, Deutsche Bundesbank, 2006.

[21] Dorofeenko, Victor, Gabriel S. Lee, and Kevin D. Salyer, "Time-varying uncertainty and the credit channel," *Bulletin of Economic Research*, 60:375-403, 2008.

[22] Eisfeldt, Andrea L., and Adriano A. Rampini, "Capital reallocation and liquidity," *Journal of monetary Economics*, 53: 369-399, 2006.

[23] Epstein, Larry, and Stanley Zin, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica*, 57:937-969, 1989.

[24] Fernandez-Villaverde, Jesus, Pable Guerro, Juan F. Rubio-Ramirez, and Martin Uribe, "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101:2530-2561, 2011.

[25] Fostel, Ana, and John Geanakoplos, "Why Does Bad News Increase Volatility and Decrease Leverage?" *Journal of Economic Theory*, 147:501−525, 2012.

[26] Gilchrist, Simon, Jae Sim, and Egon Zakrajsek, "Uncertainty, financial frictions, and investment dynamics," *Mimeo*, Boston University, 2010.

[27] Gourio, Francois, "Estimating firm-level risk," *Mimeo*, Boston University, 2008.

[28] Graham, John, Campbell Harvey, and Manju Puri, "Managerial attitudes and corporate actions," *Journal of Financial Economics*, 109:103-121, 2013.

[29] Guay, Wayne, "The Sensitivity of CEP wealth to Equity Risk: An Analysis of the Magnitude and Determinants," *Journal of Financial Economics*, 53:43-71, 1999.

[30] Guiso, Luigi, and Giuseppe Parigi, "Investment and demand uncertainty," *The Quarterly Journal of Economics*, 114:185-227, 1999.

[31] Guvenen, Fatih, "A parsimonious macroeconomic model for asset pricing," *Econometrica*, 77:1711-1750, 2009.

[32] Hansen, Gary, "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16:309327, 1985.

[33] Hartman, Richard, "The effect of price and cost uncertainty on Investment," *Journal of Economic Theory*, 5:258-266, 1972.

[34] Higson, Chris, Sean Holly, and Paul Kattuman, "The cross-sectional dynamics of the US business cycle: 1950−1999," *Journal of Economic Dynamics and Control*, 26:1539−1555, 2002

[35] Higson, Chris, Sean Holly, Paul Kattuman, and Stelios Platis, "The Business Cycle, Macroeconomic Shocks and the Cross-Section: The Growth of UK Quoted Companies," *Economica*, 71:299-318, 2004.

[36] Imrohorglu, Ayse and Selale Tuzel, "Firm Level Productivity, Risk, and Return," *Management Science*, 60:2073-2090, 2014.

[37] Kehrig, Matthias, "The cyclicality of productivity dispersion," *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*, 2011.

[38] King, Robert G., Charles I. Plosser, and Sergio T. Rebelo, "Production, growth and business cycles: I. The basic neoclassical model," *Journal of monetary Economics*, 21:195-232, 1988.

[39] Leduc, Sylvain, and Zheng Liu, "Uncertainty shocks are aggregate demand shocks," *Federal Reserve Bank of San Francisco Working Paper*, 2012.

[40] Lee, Junghoon, "Does an Aggregate Increase in Idiosyncratic Volatility Cause a Recession?", *Mimeo*, Emory University, 2013.

[41] Low, Angie, "Managerial Risk-Taking Behavior and Equity-Based Compensation", *Journal of Financial Economics*, 92:470-490, 2009.

[42] Mericle, David, "Uncertainty, the Liquidity Trap, and Social Insurance," *Mimeo*, Harvard University, 2011.

[43] Moll, Benjamin, "Productivity Losses from Financial Frictions: Can Self-financing Undo Capital Misallocation?" *American Economic Review*, 104:3186-3221, 2014.

[44] Narita, Futoshi, "Hidden Actions, Risk-Taking, and Uncertainty Shocks," *Mimeo*, University of Minnesota, 2011.

[45] Nikolov, Boris, and Toni M. Whited, "Agency Conflicts and Cash: Estimates from a Dynamic Model," *Journal of Finance*, 69:1883-1921, 2014.

[46] Oi, Walter Y., "The Desirability of Price Instability Under Perfect Competition," *Econometrica* 29:58-64, 1961

[47] Panousi, Vasia, and Dimitris Papanikolaou, "Investment, idiosyncratic risk, and ownership," *The Journal of Finance*, 67:1113-1148, 2012.

[48] Rogerson, Richard, "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 21:316, 1988.

[49] Ross, Stephen, "Compensation, Incentives, and the Duality of Risk Aversion and Riskiness," *Journal of Finance*, 59:207-225, 2004.

[50] Salgado, Sergio, Fatih Guvenen, and Bloom, Nicholas, "Skewed Business Cycles," *Mimeo* Stanford University, 2015.

[51] Schaal, Edouard, "Uncertainty, productivity and unemployment in the great recession," *Mimeo*, Federal Reserve Bank of Minneapolis, 2012.

[52] Tauchen, George, "Finite state Markov-chain approximations to univariate and vector autoregressions," *Economics letters*, 20:177-181, 1986.

[53] Tian, Can, "Technology choice and endogenous productivity dispersion over the business cycles," *Mimeo*, University of Pennsylvania, 2011.

[54] Van Nieuwerburgh, Stijn, and Laura Veldkamp, "Learning asymmetries in real business cycles," *Journal of monetary Economics*, 53:753-772, 2006.

Table 1: Countercyclical TFP Dispersion

Paper	Dispersion Calculated using	Correlation with Output	Data
Bachmann & Bayer (2013)	Firm-Specific TFP Growth	-0.47	USTAN Database (Germany)
Bloom et. al (2012)	Establishment-Specific TFP Growth	-0.46	Census/ASM (U.S.)
Kehrig (2011)	Establishment-Specific TFP Level (Durables)	-0.50	Census/ASM (U.S.)
	Establishment-Specific TFP Level (Non-durables)	-0.29	
Tian (2011)	Industry-Specific TFP Growth	-0.46	NBER-CES (U.S.)
Eisfeldt & Rampini (2006)	Industry-Specific (2-Digit SIC) TFP Growth	-0.47	Multifactor Productivity (BLS)
	Industry-Specific (4-Digit SIC) TFP Growth	-0.39	

Table 2: Positively Sloped Mean-Variance Froniter

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable:	Mean TFP							
Standard Deviation of TFP	0.572** (2.409)	0.241*** (2.954)	0.276*** (3.775)	0.173** (2.438)	0.443*** (21.84)	0.270*** (58.72)	1.368*** (9.30)	1.679*** 6.17
NBER (N) or Compustat (C)	N	N	N	N	C	C	C	C
NAICS Digit Level	4	4	4	6	-	-	-	-
No Intercept Term		✓						
Weighted Regression			✓					
No Outliers			✓					
Minimum Firm Age > 10	-	-	-	-		✓		✓
Growth (G) or Level (L)	G	G	G	G	G	G	L	L
Observations	86	86	80	462	11,689	3,851	11,689	3,851

t-statistics are in parentheses. t-statistics are computed using White heteroskedasticity robust standard errors.

*** denotes 1% , ** 5% and * 10% significance

Data Source: NBER-CES MID (Economic Variables: dtfp4 and vadd, Dates: 1958-2005) and Compustat TFP Data from Imrohoroglu and Tuzel (2014) (Dates: 1963-2009).

Note: In the regressions above a NAICS digit level of 4 implies 4-digit level data was used conditional on the first two digits being 31-33 (Manufacturing). TFP growth for higher levels of aggregation was calculated yearly as a weighted average of TFP growth at the 6-digit NAICS level. The weighting was done by the value-added of the 6-digit industry for that year. For regression (3) outliers are dropped dropped. An outlier here is defined as any point ±1.5 standard deviations from the mean of either the dependent or independent variable Industries for which data was not available prior to 1997 were dropped for the NAICS 6-digit level regression above.

Table 3: Parameter Values (Quarterly Model)

Parameter	Value	Description	Motivation
β	0.985	Subjective Discount Factor	Annual interest rate of $\approx$ 6%.
$\frac{1}{\sigma}$	1	Intertemporal Elasticity of Substitution	Standard log utility assumption.
ϕ	2.4	Scale Parameter for Labor	Time spent working is $\approx$ 0.25 in steady state.
δ	0.025	Depreciation Rate	Annual capital depreciation rate is 0.1.
η	0.5	Labor Share in Production	Labor $\div$ capital share of output $\approx$ 2 & markup = 0.33.
α	0.25	Capital Share in Production.	
λ	0.05	Manager's share of profit	Set at 5% of profit from Nikolov and Whited (2014).
τ	5	Adjustment Cost	Sensitivity analysis is provided.
$\bar{L^M}$	0.249	Steady State Labor Hours for Manager	Time spent working is $\approx$ 0.25 in steady state.
$\bar{a}$	0.0265	Steady State Intercept of Mean-Variance Frontier	From Mean-Variance frontier regression.
b	1.75	Steady State Slope of Mean-Variance Frontier	From Mean-Variance frontier regression.
σ_μ	0.0044	Standard Deviation of AR(1) Shocks to μ	Standard deviation of mean TFP is $\approx$ 0.013
σ_a	0.21	Standard Deviation of AR(1) Shocks to a	Standard deviation of the St. Dev. of TFP is $\approx$ 0.007
ρ_μ	0.91	Persistence of AR(1) Shocks to μ	Autocorrelation of mean TFP is $\approx$ 0.9
ρ_a	0.92	Persistence of AR(1) Shocks to a	Autocorrelation of the St. Dev. of TFP is $\approx$ 0.9
γ_{Min}	2.8	Minimum of Risk Aversion Distribution	Chosen to match the distribution of TFP mean choices in the data ($Mean = 0.005$, $St.Dev. = 0.008$, 5^{th} Perc. = -0.0044, and 95^{th} Per.c = 0.021).
γ_{Max}	10.5	Maximum of Risk Aversion Distribution	
γ_{Mean}	5.2	Mean of Risk Aversion Distribution	
γ_{Std}	2.1	Standard Deviation of Risk Aversion Distribution	

Table 4: Real GDP and TFP Statistics

Statistic	Variable	Data	Baseline Model			No Endogenous Choice (Only a)
			Only a	Only μ	Both a & μ	
Standard Deviation	Output	1.6	1.0	1.8	2.2	0.2
	Mean TFP	1.3	0.6	1.1	1.3	0.0
	St. Dev. Of TFP	0.7	0.7	0.03	0.7	1.9
Corr w/ Output	Output	1.0	1.0	1.0	1.0	1.0
	Mean TFP	0.4	1.0	1.0	1.0	–
	St. Dev. Of TFP	-0.4	-1.0	1.0	-0.4	0.9
Corr w/ Mean TFP	Output	0.4	1.0	1.0	1.0	–
	Mean TFP	1.0	1.0	1.0	1.0	–
	St. Dev. Of TFP	-0.5	-1.0	1.0	-0.4	–

Data Source: GDP series is from FRED and TFP Series from John Fernald's website (Economic Variables: dtfp and GDPC1, Dates: 1974-2010).
Note: Mean TFP is the 8 quarter moving average and the St. Dev. of TFP is the estimate of Conditional Heteroskedasticity from a GARCH(1,1) model. GDP data is calculated as percentage deviations from HP trend with a smoothing factor of 1600.

Table 5: Business Cycle Statistics

Statistic	Variable	Data	Baseline Model			No Endogenous Choice (Only a)
			Only a	Only μ	Both a & μ	
Standard Deviation	Output	1.6	1.0	1.8	2.2	0.2
	Consumption	1.3	0.6	1.2	1.5	0.1
	Investment	7.1	3.8	7.1	8.3	0.7
	Hours	2.1	0.5	1.0	1.2	0.1
Standard Deviation Relative to Output	Output	1.0	1.0	1.0	1.0	1.0
	Consumption	0.8	0.7	0.7	0.7	0.6
	Investment	4.4	3.9	3.9	3.8	4.7
	Hours	1.3	0.6	0.6	0.5	0.5
Corr w/ Output	Output	1.00	1.00	1.00	1.00	1.00
	Consumption	0.87	0.87	0.86	0.87	0.74
	Investment	0.93	0.88	0.88	0.87	0.90
	Hours	0.89	0.79	0.80	0.78	0.71
Autocorrelation	Output	0.85	0.91	0.91	0.92	0.83
	Consumption	0.89	0.98	0.98	0.99	0.99
	Investment	0.85	0.83	0.84	0.84	0.72
	Hours	0.93	0.82	0.83	0.83	0.64

Data Source: Real GDP, Real Personal Consumption Expenditure, Real Gross Private Domestic Investment, and Nonfarm Business Sector: Hours of All Persons from FRED (Economic Variables: GDPC1, GPDIC96, and PCECC96, HOANBS, Dates: 1974-2010).
Note: All data is calculated as a percentage deviations from HP trend with a smoothing factor of 1600.

Table 6: Robustness: Real GDP and TFP Statistics (Only a shocks)

Adjustments Costs

Statistic	Variable	Baseline Model	Alternate τ Values $\tau=\infty$	$\tau=7.5$	$\tau=2.5$	$\tau=0$
Standard Deviation	Output	1.0	0.2	0.6	1.3	1.6
	Mean TFP	0.6	0.0	0.4	0.8	1.0
	St. Dev. Of TFP	0.7	1.9	1.1	0.2	0.0
Corr w/ Output	Output	1.0	1.0	1.0	1.0	1.0
	Mean TFP	1.0	–	1.0	1.0	1.0
	St. Dev. Of TFP	-1.0	0.9	-1.0	-1.0	1.0
Corr w/ Mean TFP	Output	1.0	–	1.0	1.0	1.0
	Mean TFP	1.0	–	1.0	1.0	1.0
	St. Dev. Of TFP	-1.0	–	-1.0	-1.0	1.0

Risk Distributions + Partial Insurance

Statistic	Variable	Baseline Model	Alternate γ Distributions $\gamma_{Std}=1.05$	$\gamma_{Std}=3.15$	$\gamma_{Mean}=2.6$	*Partial Insurance*
Standard Deviation	Output	1.0	1.0	1.0	0.8	0.7
	Mean TFP	0.6	0.6	0.6	0.5	0.5
	St. Dev. Of TFP	0.7	0.7	0.7	0.8	1.0
Corr w/ Output	Output	1.0	1.0	1.0	1.0	1.0
	Mean TFP	1.0	1.0	1.0	1.0	1.0
	St. Dev. Of TFP	-1.0	-1.0	-1.0	-1.0	-1.0
Corr w/ Mean TFP	Output	1.0	1.0	1.0	1.0	1.0
	Mean TFP	1.0	1.0	1.0	1.0	1.0
	St. Dev. Of TFP	-1.0	-1.0	-1.0	-1.0	-1.0

Mean-Variance Frontier

Statistic	Variable	Baseline Model	Alternate b Values $b=1$	$b=2.5$	Non-Linear $\tau=5$	$\tau=2.5$
Standard Deviation	Output	1.0	1.1	0.9	0.4	0.7
	Mean TFP	0.6	0.5	0.5	0.4	0.6
	St. Dev. Of TFP	0.7	1.4	0.4	2.4	2.6
Corr w/ Output	Output	1.0	1.0	1.0	1.0	1.0
	Mean TFP	1.0	1.0	1.0	1.0	1.0
	St. Dev. Of TFP	-1.0	-1.0	-1.0	-1.0	-1.0
Corr w/ Mean TFP	Output	1.0	1.0	1.0	1.0	1.0
	Mean TFP	1.0	1.0	1.0	1.0	1.0
	St. Dev. Of TFP	-1.0	-1.0	-1.0	-1.0	-1.0

Note: The non-linear mean-variance frontier case is the standard deviation of $\epsilon_{i,t} = a_t + 8.12\xi_{i,t}^2$ with $\bar{a} = 0.036$. Also, the model moments are calculated to only capture the effects of changes in a_t.

Table 7: Robustness: Business Cycle Statistics (Only a shocks)

Adjustments Costs

Statistic	Variable	Baseline Model	Alternate τ Values			
			$\tau = \infty$	$\tau = 7.5$	$\tau = 2.5$	$\tau = 0$
Standard Deviation	Output	1.0	0.2	0.6	1.3	1.6
	Consumption	0.6	0.1	0.4	0.9	1.1
	Investment	3.8	0.7	2.2	5.4	6.4
	Hours	0.5	0.1	0.3	0.8	0.9
Standard Deviation Relative to Output	Output	1.0	1.0	1.0	1.0	1.0
	Consumption	0.7	0.6	0.7	0.7	0.7
	Investment	3.9	4.7	3.7	4.1	4.0
	Hours	0.6	0.5	0.6	0.6	0.6
Corr w/ Output	Output	1.00	1.00	1.00	1.00	1.00
	Consumption	0.87	0.74	0.88	0.86	0.86
	Investment	0.88	0.90	0.87	0.89	0.88
	Hours	0.79	0.71	0.78	0.80	0.79
Autocorrelation	Output	0.91	0.83	0.92	0.90	0.91
	Consumption	0.98	0.99	0.99	0.98	0.98
	Investment	0.83	0.72	0.84	0.83	0.82
	Hours	0.82	0.64	0.82	0.82	0.81

Risk Distributions + Partial Insurance

Statistic	Variable	Baseline Model	Alternate γ Distributions			*Partial Insurance*
			$\gamma_{Std} = 1.05$	$\gamma_{Std} = 3.15$	$\gamma_{Mean} = 2.6$	
Standard Deviation	Output	1.0	0.9	1.0	0.8	0.7
	Consumption	0.6	0.6	0.6	0.5	0.5
	Investment	3.8	3.5	3.8	3.0	2.5
	Hours	0.5	0.5	0.5	0.4	0.4
Standard Deviation Relative to Output	Output	1.0	1.0	1.0	1.0	1.0
	Consumption	0.7	0.7	0.7	0.7	0.7
	Investment	3.9	3.9	3.9	3.9	3.8
	Hours	0.6	0.6	0.6	0.6	0.6
Corr w/ Output	Output	1.00	1.00	1.00	1.00	1.00
	Consumption	0.87	0.87	0.86	0.87	0.88
	Investment	0.88	0.88	0.88	0.88	0.87
	Hours	0.79	0.79	0.80	0.80	0.79
Autocorrelation	Output	0.91	0.90	0.91	0.90	0.91
	Consumption	0.98	0.98	0.98	0.98	0.98
	Investment	0.83	0.82	0.83	0.82	0.82
	Hours	0.82	0.81	0.82	0.81	0.81

Mean-Variance Frontier

Statistic	Variable	Baseline Model	Alternate b Values		Non-Linear	
			$b = 1$	$b = 2.5$	$\tau = 5$	$\tau = 2.5$
Standard Deviation	Output	1.0	0.6	0.9	0.4	0.7
	Consumption	0.6	0.5	0.6	0.3	0.5
	Investment	3.8	2.0	3.4	1.4	2.6
	Hours	0.5	0.4	0.5	0.2	0.4
Standard Deviation Relative to Output	Output	1.0	1.0	1.0	1.0	1.0
	Consumption	0.7	0.7	0.7	0.7	0.7
	Investment	3.9	3.6	3.9	3.4	3.6
	Hours	0.6	0.6	0.6	0.6	0.6
Corr w/ Output	Output	1.00	1.00	1.00	1.00	1.00
	Consumption	0.87	0.89	0.86	0.90	0.89
	Investment	0.88	0.86	0.88	0.86	0.87
	Hours	0.79	0.79	0.79	0.80	0.80
Autocorrelation	Output	0.91	0.90	0.90	0.91	0.92
	Consumption	0.98	0.98	0.98	0.98	0.98
	Investment	0.83	0.81	0.82	0.81	0.84
	Hours	0.82	0.79	0.81	0.81	0.83

Note: The non-linear mean-variance frontier case is the standard deviation of $\epsilon_{i,t} = a_t + 8.12\xi_{i,t}^2$ with $\bar{a} = 0.036$. Also, the model calculated to only capture the effects of changes in a_t.

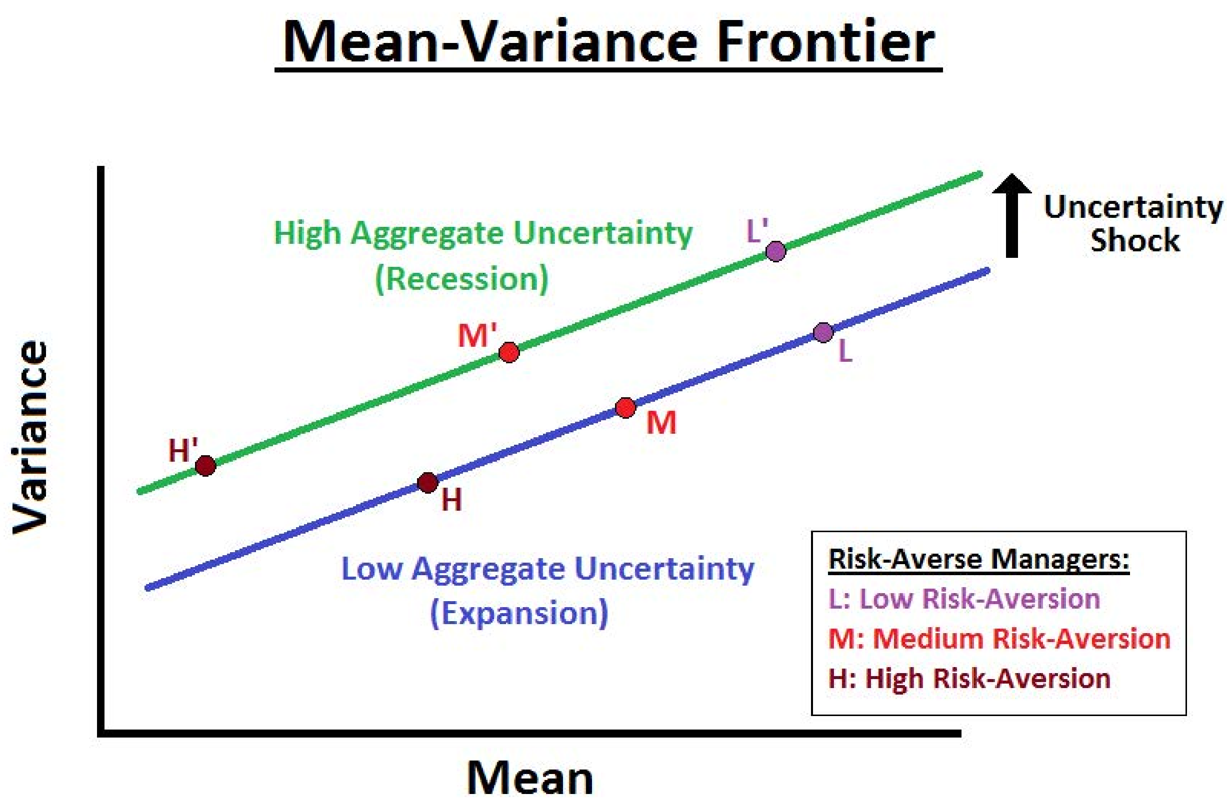

Figure 1: The Mean-Variance Frontier

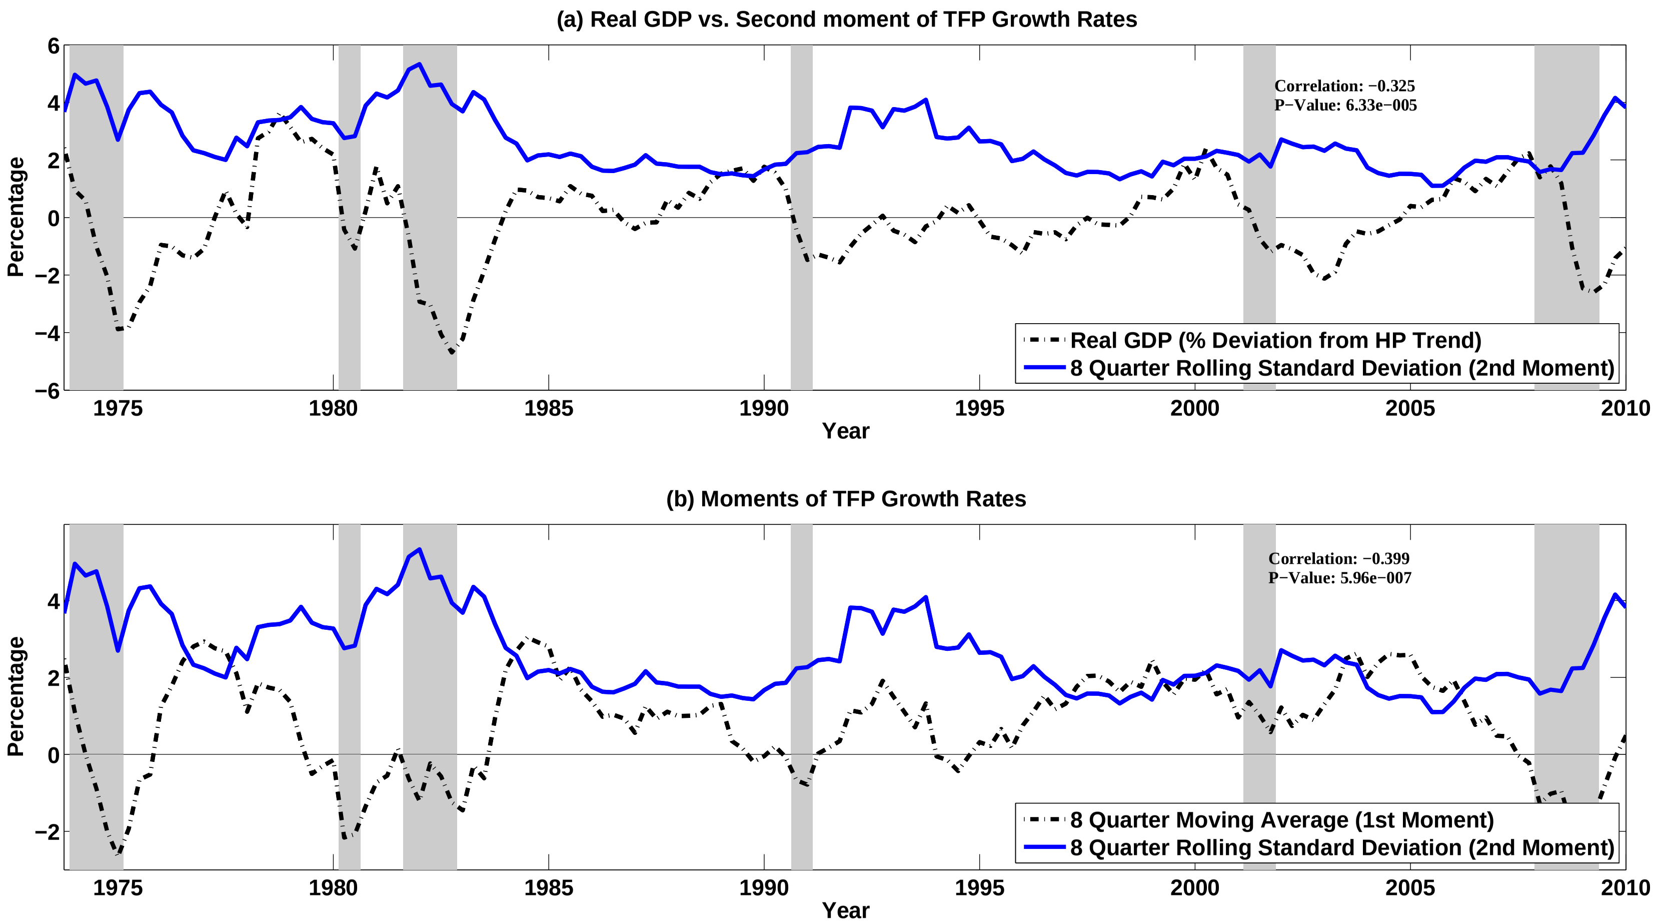

Figure 2: Real GDP vs. 2nd Moment of TFP Growth

Data Sources: TFP Series from John Fernald's website (Economic Variable: dtfp), Real GDP from FRED, and NBER Business Cycle Dates
Note: I also calculate the second moment of TFP as the conditional heteroskedasticity derived from a GARCH(1,1) model. The correlations of real GDP and the first moment of GDP with this alternative second moment measure are -0.41(significant at 1%) and -0.465 (significant at 1%) respectively. In addition, I also calculate the correlations above using the purified TFP series from John Fernalds website (Economic Variable:dtfp_util). The correlation for the purified measure of TFP with Real GDP is -0.246 (significant at 1%) and with the first moment of TFP is -0.137 (significant at 10%) . Finally, the grey bars above indicate U.S. recessions as defined by the NBER Business Cycle Dating Committee.

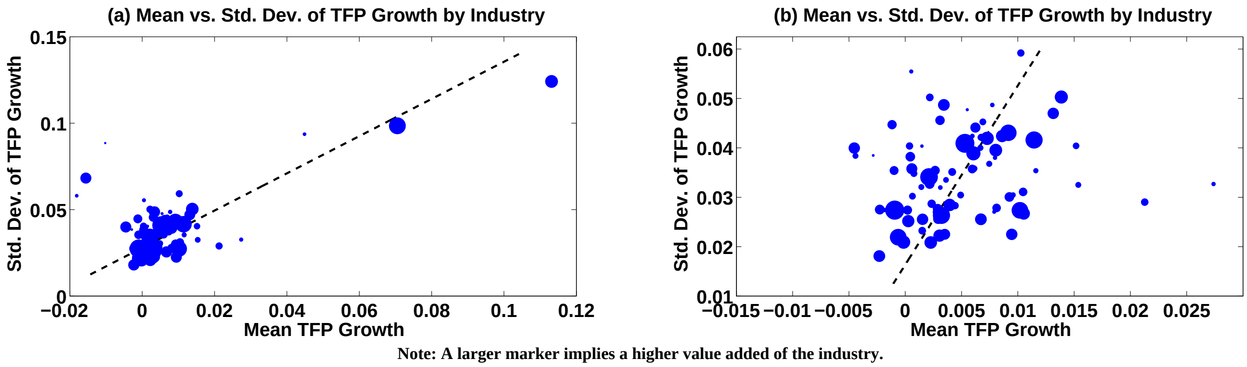

Figure 3: Mean-Variance TFP Frontier

Data Source: NBER-CES Manufacturing Industry Database (Economic Variables: dtfp4 and vadd, Dates: 1958-2005)).
Note: For the figures above manufacturing sector data at the 4-digit NAICS level was used (i.e. the first two digits were 31-33). TFP growth at the 4-digit level was calculated yearly as a weighted average of TFP growth at the 6-digit NAICS level. The weighting was done by the value-added of the 6-digit industry for that year. For panel (b) outliers are dropped. An outlier here is defined as any point ± 1.5 standard deviations from the mean along either the x or y axis. Trend lines in panel (a) and (b) are from a weighted regression of the data in each panel.

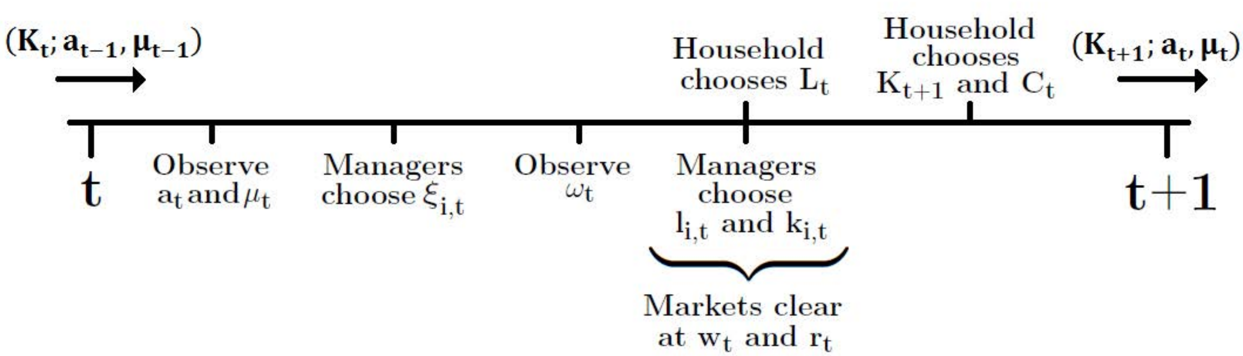

Figure 4: Model Timeline

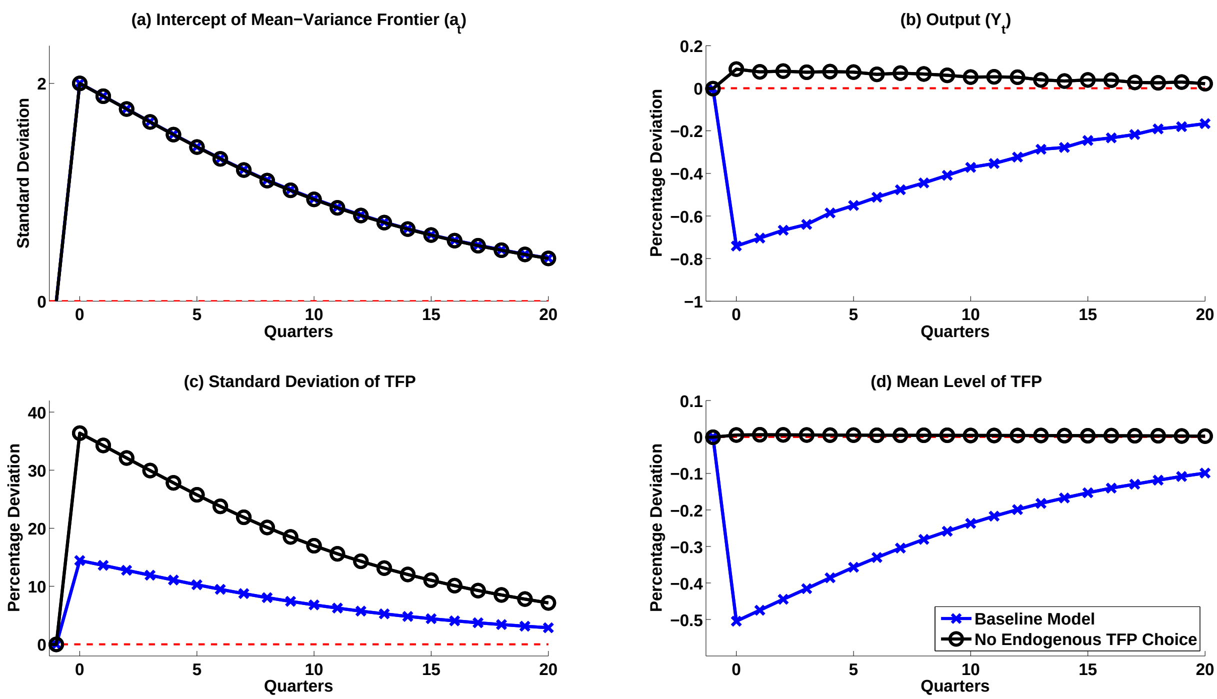

Figure 5: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

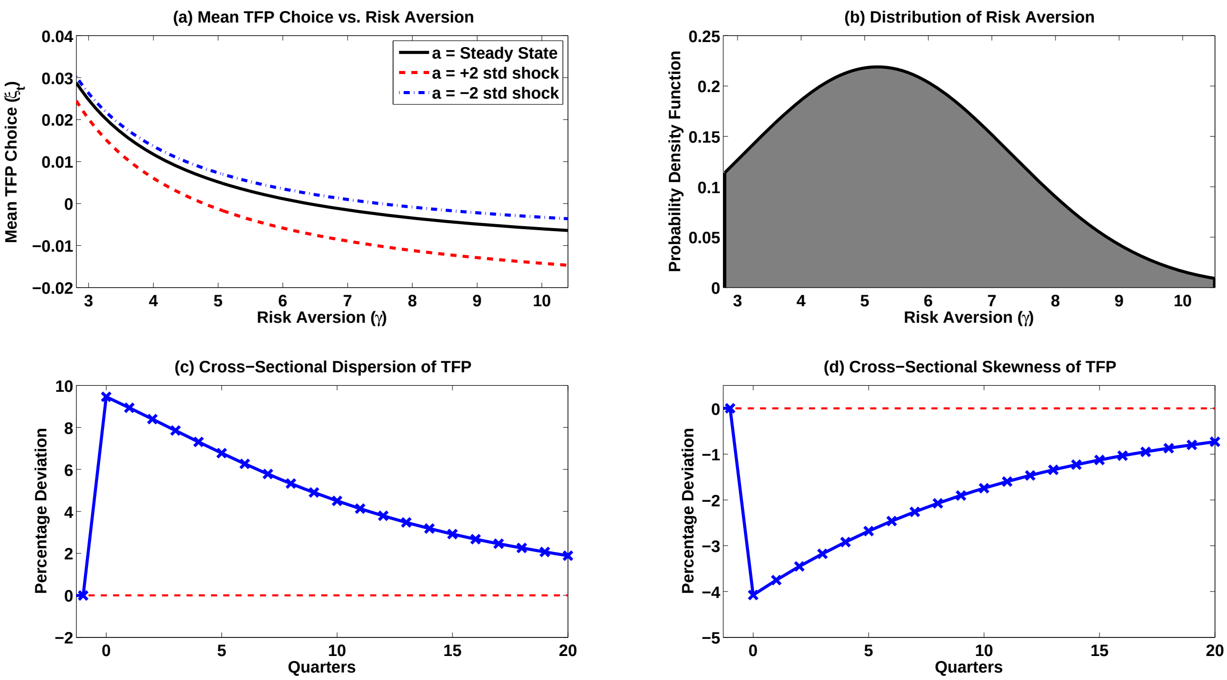

Figure 6: Choices along the Mean-Variance Frontier

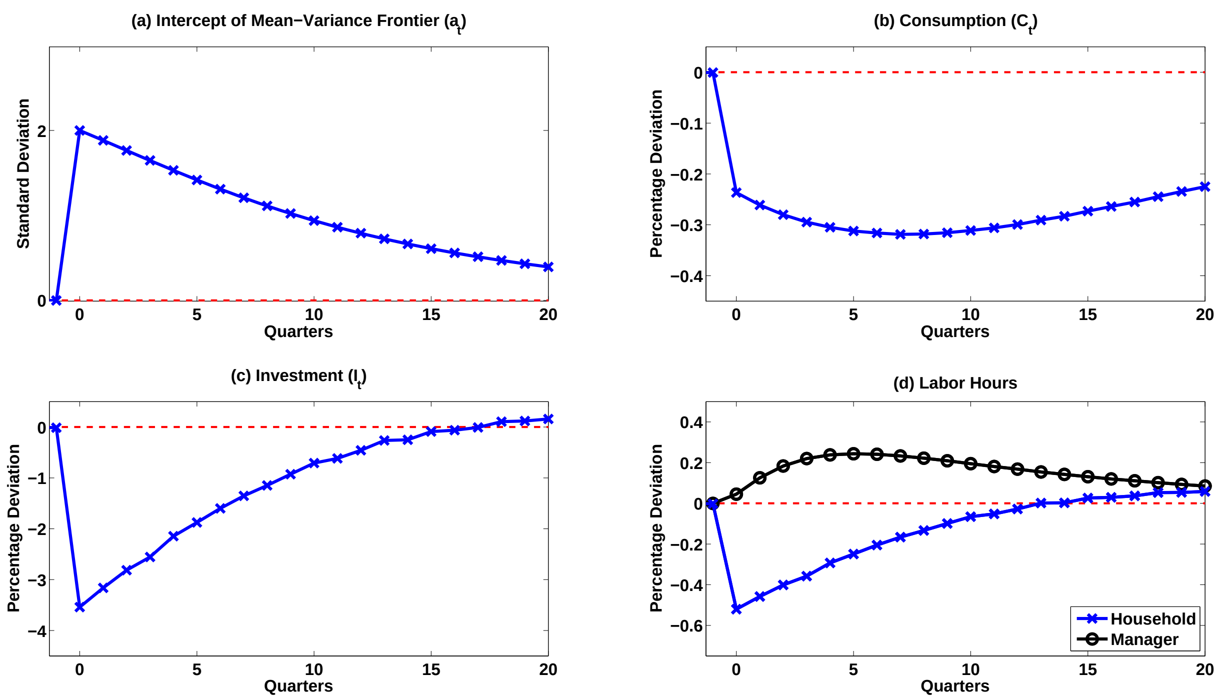

Figure 7: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

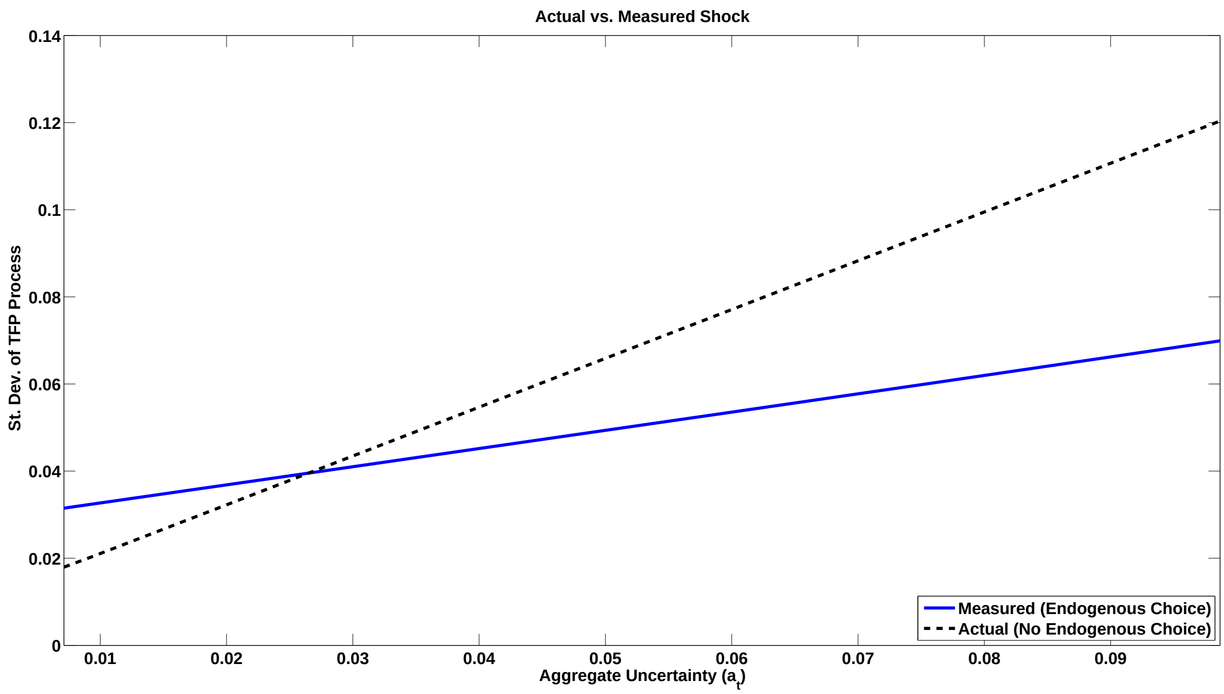

Figure 8: Measurement of Shocks

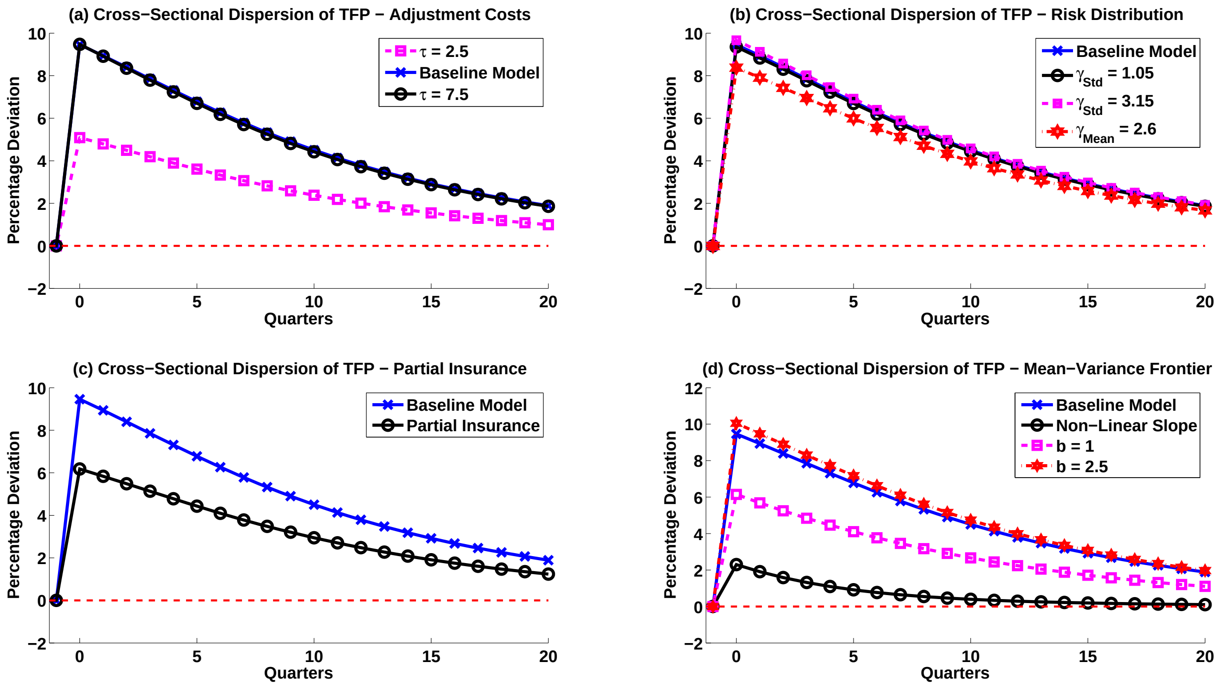

Figure 9: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

A Appendix: Solution Method (For online publication)

By assumption the manager-firm pair is single period lived, resulting in only one endogenous state variable, K_t, in the model. The evolution of K_t, can be fully described by the household's decision making. Therefore, I can solve for the evolution of this aggregate state variable by solving the household's problem and using the market clearing conditions to determine the equilibrium wage and rental rate faced by the household each period.

Mathematically, this reduces to solving the following bellman equation:

$$V\left(K_t; a_t, \mu_t, \omega_t\right) = \max_{K_{t+1}} \left\{ \tilde{U}\left(K_t, K_{t+1}; a_t, \mu_t, \omega_t\right) + \beta E\left[V\left(K_{t+1}; a_{t+1}, \mu_{t+1}, \omega_{t+1}\right) | a_t\right] \right\}$$

where the expectation is over the stochastic variables a_t, mu_t, and ω_t.

$$\tilde{U}\left(K_t, K_{t+1}; a_t, \mu_t, \omega_t\right) = U\left(C_t^{H*}, L_t^{H*}\right)$$

here $C_t^{H*}, L_t^{H*}, w_t^*, r_t^*$, and Π_t^* are given by the solution to the following 5 equation system:

$$C_t^{H*} = w_t^* L_t^{H*} + r_t^* K_t + (1-\delta) K_t + \Pi_t^* - K_{t+1} \tag{A.1}$$

$$U_{C_t^{H*}} w_t = -U_{L_t^{H*}} \tag{A.2}$$

$$L_t^{H*} = \int_{i \in [0,1]} \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t^*}{\eta}\right)^{\frac{\alpha-1}{1-\alpha-\eta}} \left(\frac{r_t^*}{\alpha}\right)^{\frac{-\alpha}{1-\alpha-\eta}} di \tag{A.3}$$

$$K_t^* = \int_{i \in [0,1]} \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t^*}{\eta}\right)^{\frac{-\eta}{1-\alpha-\eta}} \left(\frac{r_t^*}{\alpha}\right)^{\frac{\eta-1}{1-\alpha-\eta}} di \tag{A.4}$$

$$\Pi_t^* = (1-\lambda) \int_{i \in [0,1]} (1-\alpha-\eta)\, \epsilon_{i,t}^{\frac{1}{1-\alpha-\eta}} \left(\frac{w_t^*}{\eta}\right)^{\frac{-\eta}{1-\alpha-\eta}} \left(\frac{r_t^*}{\alpha}\right)^{\frac{-\alpha}{1-\alpha-\eta}} di \tag{A.5}$$

where $\epsilon_{i,t} = \mu_t + \xi(\gamma_i, a_t, \mu_t,) + (a_t + b\xi(\gamma_i, a_t, \mu_t))\, \omega_t$. Equation (A.1) defines consumption, equation (A.2) gives the households optimal labor decision (the labor supply curve), equations (A.3) and (A.4) give the market clearing conditions with the right hand sides being equal to the factor input demand curves from the manager-firm pairs, and equation (A.5) defines the households share of the profits.

I solve the bellman equation using value function iteration. The value function is defined on a 301x11x9x9 grid in the (K, a, μ, ω) dimensions. To evaluate the expectation operator I discretize the a, mu, and ωdistributions and then use numerical integration. The expectations in a and mu

are conditional and thus their discretization requires the construction of a markov chain. I use the Tauchen (1986) method to construct this Markov chain.

Next, to solve the bellman equation I need to solve for the function $\tilde{U}(K_t, K_{t+1}; a_t, \mu_t, \omega_t)$. To find this function, I first find the function $\mu(\gamma_i, a_t, \mu_t)$ and then use this to evaluate the integrals in (A.3) - (A.5), and then use the resulting equations to solve the 5 equation system given by (A.1) - (A.5):

1. To find $\mu(\gamma_i, a_t, \mu_t)$ I first discretize the ω distribution over 501 points and then use numerical integration to evaluate the expectations operator in (14). I then use golden section search to find the maximizer for each of 501 different γ_i values and the 11 a_t and 9 μ_t values. The 501 different γ_i values are chosen to match the calibrated γ_i distribution.

2. To evaluate the integrals in (A.3) - (A.5) I use numerical integration over the 501 different γ_i values.

3. To solve the 5 equation system I use substitution to simplify the problem to a one equation system in L_t^{H*}. I then use bisection search to find the root, L_t^{H*}, of this one equation.

Once I have solved for the value function, I use it to construct the policy function, $K_{t+1}(K_t; a_t, \mu_t, \omega_t)$, on the same 301x11x9x9 grid in the (K, a, μ, ω) dimensions. I then use this policy function, along with the discretized a_t, μ_t, and ω_t processes to simulate the time series for K_t, a_t, μ_t, and ω_t. The variable K_t, a_t, μ_t, and ω_t together fully describe the state of the economy and thus I can use their simulated values to calculate the time series for other aggregate variables.

In the results section, I present both impulse response functions and moments from my simulations. To calculate the moments, and specifically to isolate the effects of a_t on the moments, I simulate the economy for 10500 periods with both a_t, μ_t, and ω_t shocks. I then repeat this exercise 60,000 times keeping the a_t shocks the same but allowing the ω_t and μ_t shocks to be different across simulations. I then average across these 60,000 simulation exercises to produce a 10500 period simulation that is relatively purified of the effects of ω_t and μ_t shocks. Next, I drop the first 500 periods of this 10500 period simulation and use the remaining 10000 periods to calculate the moments.[36] Dropping the first 500 periods allows the economy to settle into its stochastic steady state.

To calculate the impulse responses I simulate the economy for 260 periods with the a_t, μ_t, and ω_t shocks. I force a_t to be +2 std from the mean in period 210. I repeat this exercise 2.4 million times over different realizations of a_t, μ_t, and ω_t shocks each time and calculate the impulse response as the average of these 2.4 million simulations. The figures in the paper plot periods 209 through 230 of these impulse responses and label period 210, the period of the shock, as 0.

[36] I perform similar exercises to isolate the effects of μ_t, and both a_t and μ_t for Table 4 and 5.

B Appendix: Impulse Responses for Robustness Exercises (For online publication)

B.1 First Moment Shocks

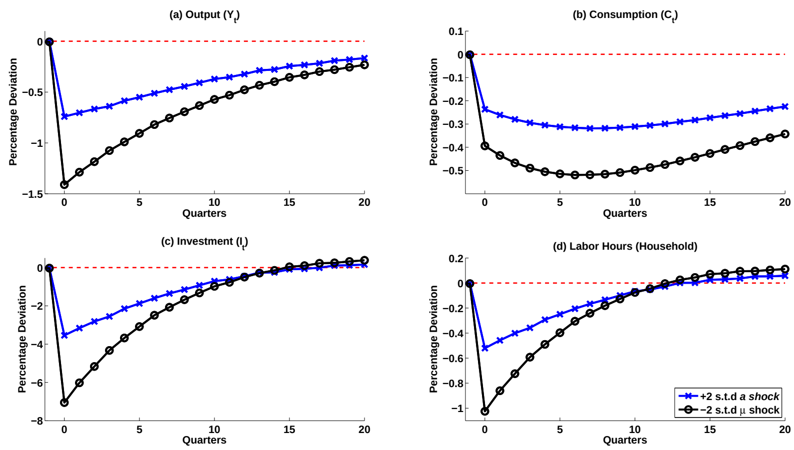

Figure 10: Real Variables: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock) & a −2 St. Dev. Shock to μ (Traditional First Moment Shock)

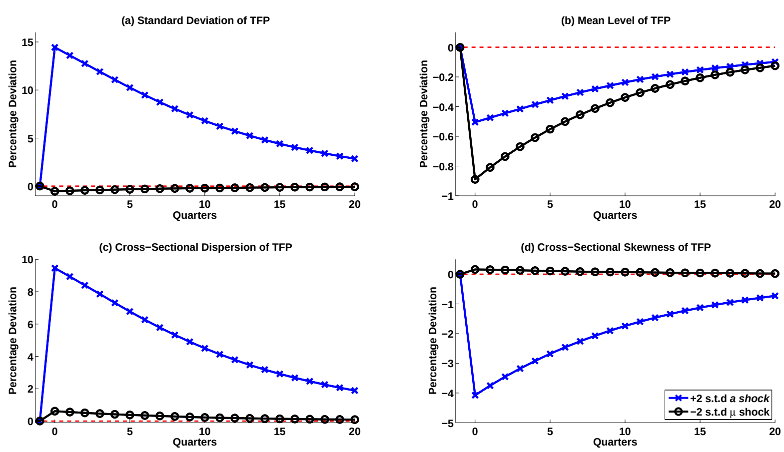

Figure 11: TFP Moments: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock) & a −2 St. Dev. Shock to μ (Traditional First Moment Shock)

B.2 Adjustment Costs

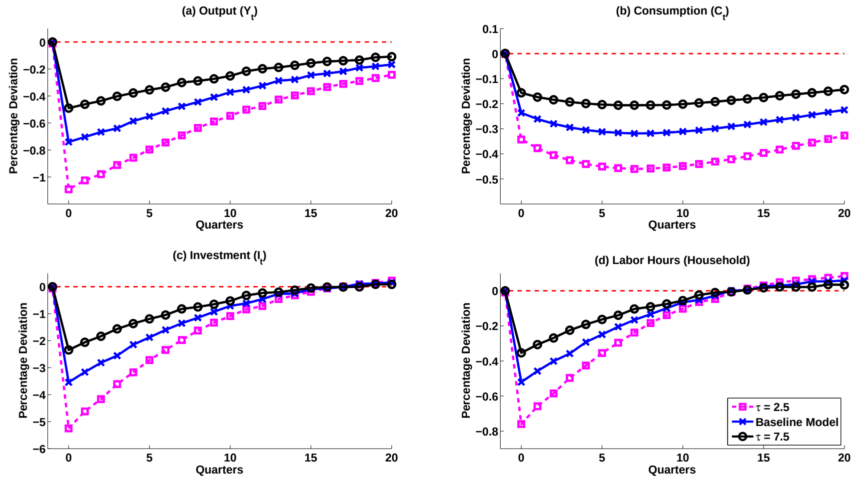

Figure 12: Real Variables: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

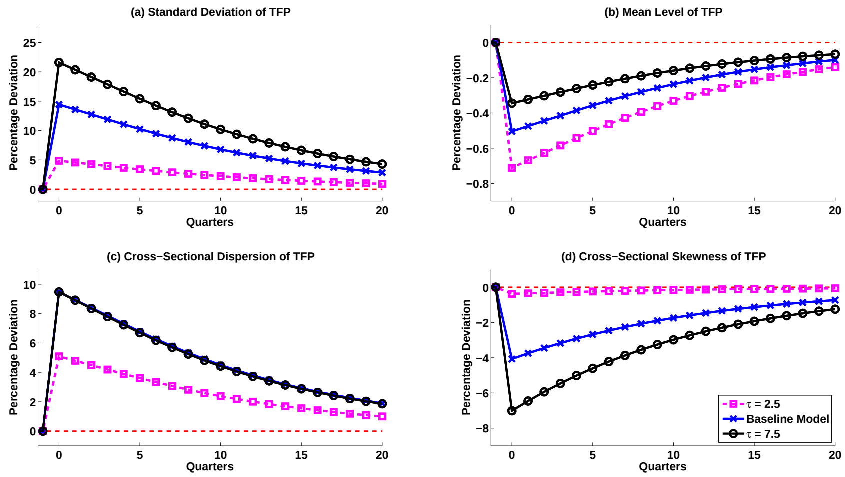

Figure 13: TFP Moments: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

B.3 Risk Distributions

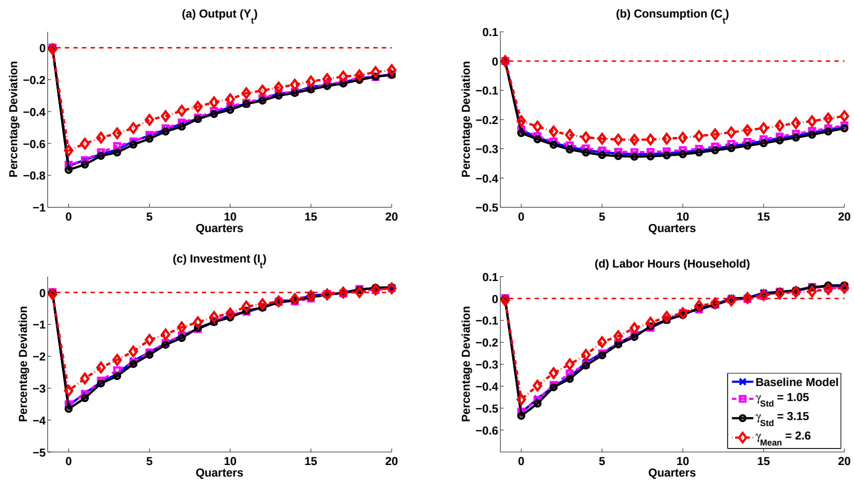

Figure 14: Real Variables: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

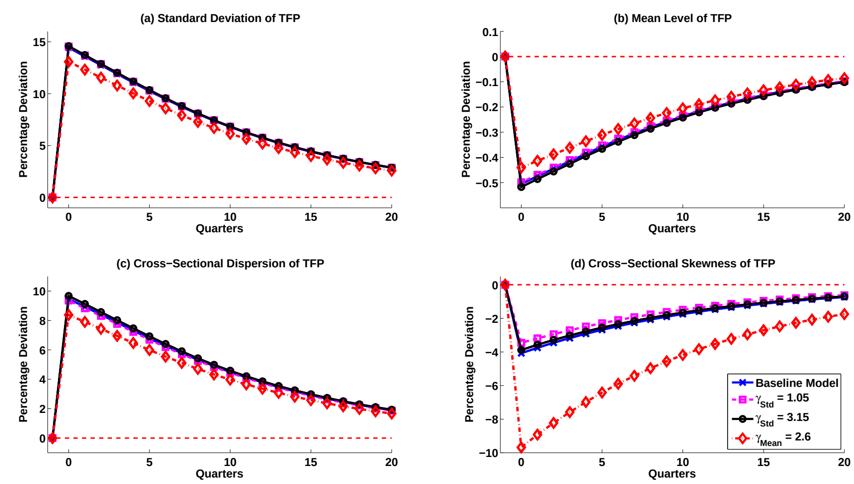

Figure 15: TFP Moments: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

B.4 Partial Insurance

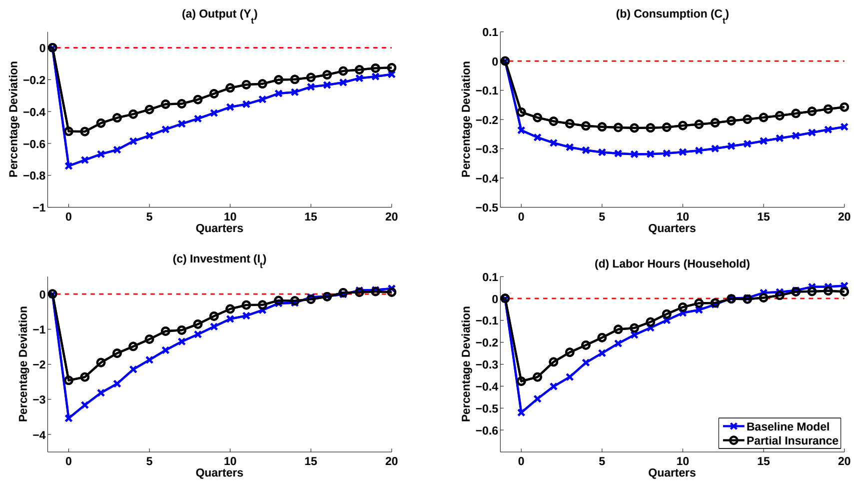

Figure 16: Real Variables: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

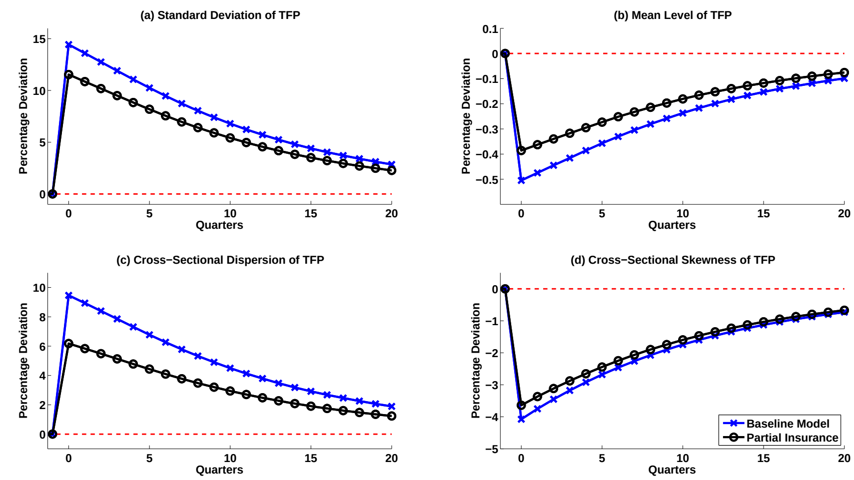

Figure 17: TFP Moments: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

B.5 Mean-Variance Frontier

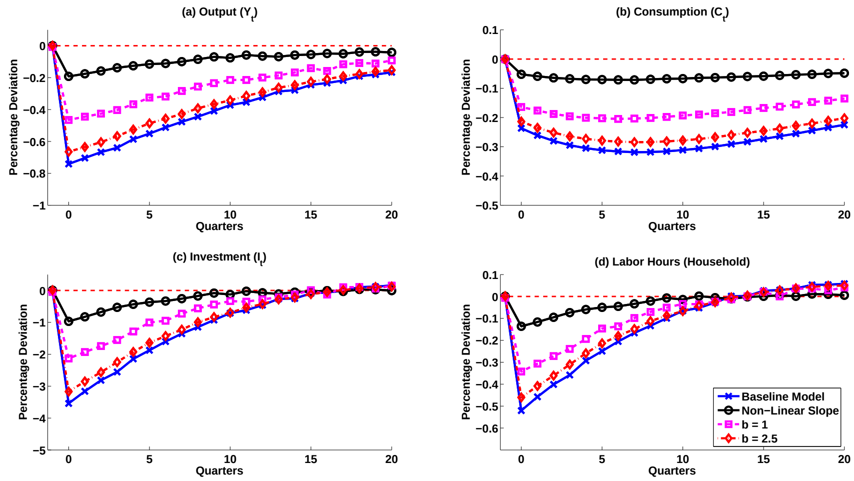

Figure 18: Real Variables: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)

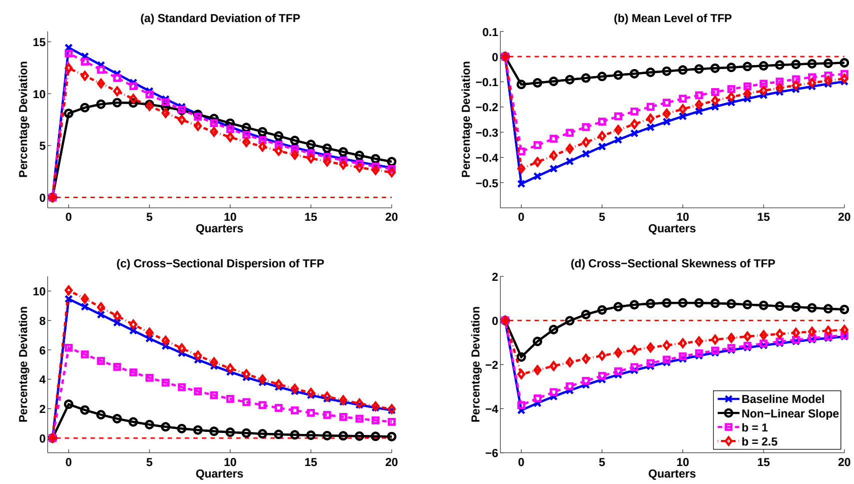

Figure 19: TFP Moments: Impulse Responses to a +2 St. Dev. Shock to the Intercept of the Mean-Variance Frontier (Uncertainty Shock)