Really Uncertain Business Cycles

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten and Stephen Terry^{*}

December 2011

Abstract

We propose uncertainty shocks as a new impulse driving business cycles. First, we demonstrate that measured uncertainty is strongly countercyclical. This is true both at the aggregate and at the industry level: slower aggregate and industry growth is associated with higher aggregate and industry uncertainty. Using trade and exchange rate instrumental variables we show that this slower industry growth is not causing the rise in uncertainty. Instead, uncertainty appears to be an exogenous process, suggesting recessions are driven by a combination of first and second moment shocks. We then build a DSGE model with heterogeneous firms, time-varying uncertainty and adjustment costs to quantify the impact of these second moment shocks. We find they typically lead to drops of about 2% in GDP, suggesting uncertainty could play an important role in driving business cycles. We also find that because uncertainty makes firms cautious it substantially reduces the response of the economy to stimulative policy, leading to pro-cyclical multipliers.

Keywords: uncertainty, adjustment costs and business cycles.

JEL Classification: D92, E22, D8, C23.

Disclaimer: This is a finished version of a prior draft by Bloom, Floetotto and Jaimovich which was suspended after Floetotto left academia. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. We are extremely grateful to the National Science Foundation for their financial support. We thanks seminar audiences in the AEA, Berkeley, Boston College, Chicago, Chicago Fed, Cleveland Fed, CMU, the Census Bureau, the ECB, Federal Reserve Board, HEC, the IMF, Minnesota, the NBER, Northwestern, NY Fed, Ohio State, Oxford, Penn, Philadelphia Fed, SF Fed, SITE, Stanford, UC Irvine, USC, UT Austin, Warwick and Yale.

^{*}Bloom, Saporta-Eksten and Terry at Stanford University, Floetotto at McKinsey, and Jaimovich at Duke University. Correspondence: Nick Bloom, Department of Economics, Stanford University, Stanford, CA 94305, nbloom@stanford.edu.

1 Introduction

Uncertainty has received substantial attention as a potential factor shaping the depth and duration of the Great Recession. This paper seeks to rigorously evaluate these claims in two parts. In the first part we develop new empirical measures of uncertainty using detailed Census micro data from 1972 to 2009, while in the second part we build a micro-to-macro stochastic dynamic general equilibrium model to quantify the impact of these changes in uncertainty on the economy.

On the empirical front we use Census data to detail four main results. First, the dispersion of plant-level shocks to TFP is strongly counter-cyclical, rising steeply in recessions (see Figure 1). The variance of these shocks in fact reached an all time high in the recession of 2008-2009. Surprisingly, higher moments like kurtosis and skewness of TFP shocks do not appear to be counter-cyclical. So recessions appear to be well characterized by a negative first moment and a positive second moment shock.¹

Second, uncertainty is also strongly counter-cyclical at the industry level. That is, within SIC 4-digit industries the yearly growth rates of output is negatively correlated with the dispersion of TFP shocks to establishments within the industry. Hence, both at the industry and at the aggregate level bad times - defined in terms of low growth rates of output - are also uncertain times in terms of increased cross-sectional dispersion of TFP shocks.

Third, this industry level increase in variance during periods of slow growth does not appear to be due to the slow-down itself. We use trade-reforms and exchange rate changes to instrument for industry growth rates and find this has no *causal* impact on industry uncertainty. Hence, the increase in uncertainty at the industry level that occurs during slow downs seems not to be directly driven by the slow down. This suggests that industry (and macro) recessions are combinations of negative first moment and positive second moment shocks. This seems quite intuitive - the types of events that cause recessions, like oil shocks, wars, policy shocks and financial crisis, are also likely to increase uncertainty.

Finally, we show that our plant-level TFP shock measures of uncertainty are highly correlated with firm and industry level measures of uncertainty. In particular, for plants with publicly listed parents the absolute size of their annual TFP shock is highly correlated to the volatility of their parent firms' daily stock-returns within the year. In the prior uncertainty literature daily stock returns volatility has been the most common measure of uncertainty, so it is reassuring that this is strongly correlated with our plant-level TFP shock measure. We also find similar correlations of the spread of plant-level TFP shocks to the volatility of their parent firms quarterly sales growth and the volatility of their industry's monthly output growth.

¹All data available at http://www.stanford.edu/~nbloom/RUBC_data.zip.

Given the robust evidence that uncertainty appears to rise sharply in recessions we build a dynamic stochastic general equilibrium model. Various features of the model are specified to conform as closely as possible to the standard frictionless real business cycle (RBC) model as this greatly simplifies comparison with existing work. We deviate from this benchmark in three ways. First, uncertainty is time-varying, so the model includes shocks to both the level of technology (the first moment) and its variance (the second moment) at both the micro and macro level. Second, there are heterogeneous firms that are subject to idiosyncratic productivity and demand shocks. Third, the model contains convex and non-convex adjustment costs in both capital and labor. The non-convexities together with time variation in uncertainty imply that firms become more cautious in investing and hiring when uncertainty increases.

Simulations allow us to study the response of our model economy to an uncertainty shock. We show that a rise in uncertainty makes it optimal for each individual firm to wait, leading to a significant fall in aggregate economic activity. In addition, we show that time-varying uncertainty reduces productivity growth during times of high uncertainty because it lowers the extent of reallocation in the economy. When uncertainty rises productive firms expand less and unproductive firms contract less.²

We then build on our theoretical model to investigate the effects of uncertainty on policy effectiveness. We use a simple illustrative example to show how time-varying uncertainty significantly dampens the effect of an expansionary policy. The key to this *policy ineffective-ness* is that a rise in uncertainty makes firms very cautious in responding to any stimulus and that includes the policy impulse. The impact of this stimulus is thus mitigated relative to its impact in low uncertainty times.

Our work is related to several strands in the literature. First, we add to the extensive literature building on the RBC framework that studies the role of productivity (TFP) shocks in causing business cycles. In this literature, recessions are generally caused by large negative technology shocks.³ The reliance on negative technology shocks has proven to be controversial, as it suggests that recessions are times of technological regress.⁴ As discussed above, our work provides a rationale for falls in measured productivity. Countercyclical increases in uncertainty lead to a freeze in economic activity, substantially lowering productivity growth during recessions. In our model, however, the drop in productivity is not

²In the actual U.S. economy, reallocation is a key factor driving aggregate productivity. See, for example, Foster, Haltiwanger and Krizan (2000, 2006), who report that reallocation, broadly defined to include entry and exit, accounts for around 50% of manufacturing and 80% of retail productivity growth in the US.

 $^{^{3}}$ See, for example, the review of this literature see King and Rebelo (1999) and Rebello (2005).

⁴This reasoning has lead many researchers to study models with other disturbances, which also mostly focus on first-moment (levels) shocks. A partial list of these alternative shocks includes oil shocks, investment specific shocks, monetary shocks, government expenditure shocks, news shocks, and terms-of-trade shocks. Yet, in most models, negative technology shocks continue to be an important driver of economic downturns.

causing the recession, but rather an artifact of a recession that is caused in turn by an increase in uncertainty.

Second, the paper relates to the literature on investment under uncertainty. A rapidly growing body of work has shown that uncertainty can directly influence firm-level investment and employment in the presence of adjustment costs.⁵ Many of the most recent papers have started to focus on stochastic volatility and its impacts on the economy, particularly focusing on the current recession.⁶ Third, the paper also builds upon a recent literature that studies the role of micro-rigidities in general equilibrium macro models.⁷

The remainder of this paper is organized as follows. Section 2 discusses the behavior of uncertainty over the business cycle. In Section 3 we formally present the model, define the recursive equilibrium, and present our non-linear solution algorithm which builds on the work of Krusell and Smith (1998), Kahn and Thomas (2008) and Bachman, Caballero and Engel (2008). The model is calibrated and simulated in Section 4, where we study the role of uncertainty shocks in driving the business cycle. Section 5 studies the impact of policy shocks in the presence of time-varying uncertainty. Section 6 concludes.

2 Measuring Uncertainty Over the Business Cycle

This section presents evidence on the cyclical behavior of uncertainty. Before presenting our empirical results, it is useful to briefly discuss what we mean by time-varying uncertainty in the context of our model.

A firm indexed by j produces output in period t according to the following production function

$$y_{j,t} = A_t z_{j,t} f(k_{j,t} n_{j,t}) \tag{1}$$

where $k_{t,j}$ and $n_{t,j}$ denote idiosyncratic capital and labor employed by the firm. Each firm's productivity is a product of two separate processes: an aggregate component, A_t , and an idiosyncratic component, $z_{j,t}$. More generally, we can think of this as a revenue function so that demand shocks will also be incorporate into the process for A_t and $z_{j,t}^8$, so these

⁵See, for example; Bernanke (1983), Romer (1990), Bertola and Caballero (1994), Dixit and Pindyck (1994), Abel and Eberly (1996), Hassler (1996), and Caballero and Engel (1999).

⁶See for example, Bloom (2009) partial equilibrium model with stochastic volatility, Fernandez-Villaverde et al.'s (2009) paper on uncertainty and exchange rates, Kehrig's (2011) paper on counter-cyclical productivity and mark-ups, Christiano et al. (2010), Arrelano et al. (2011) and Gilchrist et al.'s (2011) papers on uncertainty shocks in models with financial constraints, Basu and Bundick's (2011) paper on uncertainty shocks in a new-Keynsian model, Fernandez-Villaverde et al.'s (2011) paper on policy uncertainty, and Bachman and Bayer's (2011) paper on micro-uncertainty with capital adjustment costs.

⁷See for example, Thomas (2002), Veraciertio (2002), Kahn and Thomas (2008 and 2011), Bachman et al. (2008), and House (2008).

⁸See also Hoppenhayn and Rogerson (1993) where they discuss how productivity shocks at the micro level are isomorphic to consumer taste shocks shifting the demand curve.

processes can be labelled as macro and micro 'business conditions'.

To match the empirical evidence for macro and micro uncertainty, we assume that the aggregate and idiosyncratic components of business conditions follow autoregressive processes:

$$\log(A_t) = \rho^A \log(A_{t-1}) + \sigma^A_{t-1} \epsilon_t \tag{2}$$

$$\log(z_{j,t}) = \rho^Z \log(z_{j,t-1}) + \sigma^Z_{t-1} \epsilon_{j,t}$$

$$\tag{3}$$

We allow the variance of innovations, σ_t^A and σ_t^Z , to vary over time to generate periods of low and high macro and micro uncertainty.

There are two assumptions embedded in this formulation. First, the volatility in the idiosyncratic component, $z_{j,t}$, implies that productivity and demand dispersion *across* firms is time varying, while volatility in the aggregate component, A_t , implies that *all* firms are affected by more volatile shocks. Second, given the timing assumption in (2) – (3), firms learn in advance that the distribution of shocks from which they will draw in the next period is changing. This timing assumption captures the notion of uncertainty that firms face about future business conditions.

As we argue in our theoretical section, these two shocks have different implications in terms of the statistics that are driven by them. Volatility in $z_{t,j}$ implies that cross-sectional measures of firm performance (output, sales, stock market returns etc.) are time varying, while volatility in A_t induces higher variability in aggregate variables like GDP growth and the S&P500 index. Next we turn to our cross-sectional and macro uncertainty measures, details of the construction of which are contained in Appendix A.

2.1 Micro Uncertainty over the Business Cycle

In this section we present a set of results showing that shocks at the *establishment*, *firm* and *industry* level all increase in variance during recessions. In our model in section 3 we focus on units of production, ignoring multi-establishment firms or industry level shocks to reduce computational burden. Nevertheless, we present data at these three different levels to demonstrate the generality of the increase in idiosyncratic shocks during recessions.

Establishment level evidence: Our first set of measures come from the Census panel of manufacturing establishments. In summary (with extensive details in Appendix A) this contains detailed output and inputs data on over 50,000 establishments from 1972 to 2009. We focus on the subset of 15,673 establishments with 25+ years of data to ensure compositional changes do not bias our results, generating a sample of almost half a million establishment-year observations.⁹

⁹ The sampling issues arises both from the cyclicality of exit and from the sample stratification rules for the

To measure uncertainty we first calculate establishment level total-factor productivity (TFP) $(\hat{z}_{j,t})$ using the cost share approach to estimating input elasticities (see for example Foster, Haltiwanger and Krizan (2000)). We then define TFP shocks $(e_{j,t})$ as the residual from the following first-order autoregressive equation for establishment level log TFP

$$\log\left(\widehat{z}_{j,t}\right) = \rho \log\left(\widehat{z}_{j,t-1}\right) + \mu_i + \lambda_t + e_{j,t} \tag{4}$$

where μ_i is an establishment level fixed effect (to control for establishment level differences) and λ_t is a year fixed effect (to control for cyclical shocks). Since this residual will also contain plant-level demand shocks that are not controlled for by 4-digit price deflators (see Foster, Haltiwanger and Syverson (2008)) our measure will combine both TFP and demand shocks. Because our model is isomorphic in idiosyncratic productivity and demand shocks this is not a theoretical problem, but does highlight the difficulty in empirically distinguishing productivity shocks from demand shocks.

Finally, we then define micro-uncertainty, $\sigma_{t-1}^{\widehat{Z}}$, as the cross sectional standard deviation of $e_{j,t}$ calculated on a yearly basis. This is shown in Figure 1 as the interquartile range of this TFP shock within each year, displaying a clearly counter-cyclical behavior.

Table 1 more systematically evaluates the relationship between the dispersion of TFP shocks and recessions. In column (1) we regress the cross-sectional standard-deviation (S.D.) of establishment TFP shocks on an indicator for the number of quarters in a recession during that year.¹⁰ We find a coefficient of 0.063 which is highly significant (a t-statistic of 6.3). In the bottom panel we report that this S.D. of establishment TFP shocks also has a highly significant correlation with GDP growth of -0.440. In columns (2) and (3) we look at the skewness and kurtosis of TFP shocks over the cycle and interestingly find no significant correlations.¹¹ This suggests that recessions can be well characterized at the micro level as a negative first-moment shock plus a positive second moment shock, with no shocks to any other higher moments. In column (4) we use our preferred outlier-robust measure of cross-sectional dispersion, which is the interquartile (IQR) range of TFP shocks, and again find this rises significantly in recessions.¹² In column (5) as another robustness test we use

Census, which rotates out smaller establishments at 5-yearly intervals. By restricting the sample to 25+ lived establishments we eliminate cyclical frequency sampling fluctuations. As we show in section 2.1 and in Bloom, Floetotto and Jaimovich (2009), our results also hold (and in some cases are stronger) using the full panel, and very similar when restricting to establishments with 35+ years of data.

¹⁰So, for example, this variable has a value of 0.25 in 2007 as the recession started in quarter IV, and values of 1 and 0.5 in 2008 and 2009 as the recession continued until quarter II in 2009.

¹¹This lack of significant correlation seemed pretty robust in a number of experiments we ran. For example, if we drop the time trend and census survey year controls the result is column (1) on the standard deviation remains highly significant at 0.061 (0.019), while the results in columns (2) and (3) on skewness and kurtosis remain insignificant at -0.241 (0.217) and -0.771 (2.797).

¹²Kehrig (2011) finds that the dispersion of TFP increases in recessions mostly for durables. We run column 4 separately for durables and nondurables. We find that in our sample the rise of IQR of TFP shocks for

output (rather than TFP) shocks and find a significant rise in recessions. We also run a range of other experiments on different indicators, measures of TFP and samples and always find that dispersion rises significantly in recession.¹³

In column (6) we use a different dataset which is the sample of all Compustat firms with 25+ years of data. This has the downside of being a much smaller selected sample containing only 2,465 publicly quoted firms, but spans all sectors of the economy, and provides quarterly sales observations going back to 1962. We find that the quarterly dispersion of sales growth in this Compustat sample is also significantly higher in recessions.

Or course, one important caveats when using the variance of productivity 'shocks' to measure uncertainty are that the residual $e_{j,t}$ is a productivity 'shock' only in the sense that it is unforecasted by the regression equation (4), rather than unforecasted by the establishment. Hence, it parallels the definition of a macro productivity 'shock' by Kydland and Prescott (1983) in being a forecast error from an AR(1) equation rather than necessarily a shock to economic agents. To address this concern in column (7) we examine cross-sectional spread of stock-returns - which reflects the volatility of news about firm performance - and find that this is again counter-cyclical, echoing the prior results in Campbell et al (2001). In fact we find (as discussed in detail below and shown in Table 4) that establishment-level shocks to TFP are significantly correlated to their parents stock-returns, so that (at least part of) these establishment TFP 'shocks' are news to the market.

Finally, Column (8) examines another measure of uncertainty, which is the cross-sectional spread of industry-level output growth rates, finding again this is strongly counter-cyclical.

Hence, in summary plant, firm and industry level measures of volatility and uncertainty all appear to be strongly countercyclical, suggesting that micro-uncertainty rises in recessions.

durables is larger, with a point estimate (standard error) of 0.077 (0.028), but that there is also a significant increase in dispersion for nondurables, with a point estimate (standard error) of 0.044 (0.019).

¹³For example, IQR of employment growth rates has a point-estimate (standard-error) of 0.053 (0.013), the IQR of TFP shocks measured using an industry by industry forecasting equation version of (4) has a point-estimate (standard-error) of 0.059 (0.019), using 2+ year samples for the S.D. of TFP shocks we find point estimates (standard-errors) of 0.042 (0.013), or even using the S.D. of shocks to labor productivity we find a point-estimate (standard-error) of 0.068 (0.009).

Two other recent papers have also reported similar findings of counter-cyclical increases in the variance of 'productivity shocks'. Bachman and Bayer (2009) use a panel of public and private German firms spanning manufacturing and retail, showing significant increases in the variance of innovations to productivity during recessions. Kehrig (2010) like our paper uses the US Census data, but takes a different approach to sampling and estimating productivity, and again finds a significant increase in the variance of productivity shocks in recessions. Finally, looking at price changes Berger and Vavara (2010) find that changes in product-level prices (e.g. the price of a 2 litre bottle of coke) are also more dispersed during recessions.

2.2 Industry business cycles and uncertainty

In Table 2 we report another set of results which highlights the generality of the relationship between uncertainty and recessions. To do this we exploit the size of our census dataset to examine the dispersion of productivity shocks *within* each SIC 4-digit industry year cell. The size of Census dataset means it has a mean (median) of 27.1 (17) establishments per SIC 4-digit industry-year cell, which enables us to examine the link between within-industry dispersion of establishment TFP shocks and industry growth.

Table 2 runs a series of industry panel regressions in which our dependent variable is the IQR of TFP shocks for all establishments in each industry-year cell. The explanatory variable in column (1) is the median growth rate of output in the industry-year cell, with a full set of industry and year fixed-effects also included.¹⁴ Column (1) of Table 2 shows that the within industry dispersion of TFP shocks is significantly higher when that industry is growing more slowly. Since the regression has a full set of year (and industry) dummies this is independent of the macroeconomic cycle. So at both the aggregate and industry level slow-downs in growth are associated with increases in the cross-sectional dispersion of shocks.

One immediate question is why is within industry dispersion of shocks is higher during industry slow-downs. Maybe this is because industry slow-downs impact some types of establishments differently? To investigate this columns (2) to (9) run a series of estimations to check if the increase in within industry dispersion is larger for some particular characteristics of the industry. In column (2) we interact industry growth with the median growth rate in that industry over the full period - perhaps faster growing industries are more counter-cyclical in their dispersion? We find no relationship, suggesting long-run industry growth rates are not linked to the increase in dispersion of establishment shocks they see in recessions. In column (3) we interact industry growth with the dispersion of industry growth rates - perhaps industries with a wide spread of growth rates across establishments are more counter-cyclical in their dispersion, but again find nothing. In columns (4) and (5) we look at the median and dispersion of plant size within each industry (measured by the number of employees) and again find nothing, in columns (6) and (7) we look at the median and dispersion of capital/labor ratios and again find nothing, and finally in columns (8) and (9) we look at TFP and geographical dispersion interactions and again find nothing.

So, in summary, it appears that firstly the within-industry dispersion of establishment TFP shocks rises sharply when the industry growth rates slow down. Secondly, somewhat surprisingly this relationship appears to be broadly robust across industries, in that no

¹⁴ We use the median rate of output growth in the industry-year to ensure the results are robust to establishment outliers. Results for column (1) using the mean of output growth across establishments are in fact slightly larger with a point-estimate (standard-error) of -0.151 (0.017).

industry characteristic is particularly linked to this.

2.3 Is uncertainty a cause or effect of slowdowns?

An obvious question regarding the relationship between uncertainty and the business cycle at the aggregate and industry level is what causes what. Does uncertainty drive the cycle, or do recessions drive increases in uncertainty? A recent literature has suggested a number of mechanisms for uncertainty to increase endogenously in recessions, so identifying the direction of causation is clearly important¹⁵.

To do this we need some kind of natural experiment or instrument that causes changes in the first moment, that we can use to investigate its causal impact on the second moment. Unfortunately no instrument obviously exists at the macro level because as Kocherlakota (2009) noted "the difficulty in macroeconomics is that virtually every variable is endogenous". But at the industry level we do have two sets of instruments for first moment shocks that we can use to evaluate what the causal impact of first moment shocks are on the second moment.

The first approach uses China joining the WTO as a quasi-experiment, which led to the abolition of import quotas on Chinese textiles and apparel in 2005. The origin of these quotas dates back to the 1950s when the US started putting quotas on Japanese imports. Over time, this quota system was expanded to take in most developing countries including China, and was formalized into the Multi-Fiber Agreement in 1974, which was itself integrated into GATT in 1994. After China joined the WTO (the successor to GATT) in December 2001 these quotas were scheduled for a multi-phase elimination, mostly in the final phase in 2005 (see Brambilla, Khandelwal and Schott, 2010). Importantly, the initial level of the quotas varied quasi-randomly across four-digit industries – for example, they covered 100% of women's dresses (2334) but only 5% of men's trousers (2325). This reflected the historic bargaining power of the various industries in the 1950s when these quotas were introduced, but by the 2000s was seemingly uncorrelated to any demand or technology trends (see Bloom, Draca and Van Reenen, 2011). The initial level of these quotas thus provides a good instrument for the growth of industry output between 2004 and 2006 because: (a) industries with initial quotas should (and did) see substantially lower growth rates than those without quotas, and (b) these quota effectively vary randomly across textiles and apparel industries.

Since these quotas operate at the HSIC 6-digit level we aggregate this up to the Census SIC 4-digit level by weighting each 6-digit sub-industry by its share of trade within the

¹⁵See, for example, the papers on information collection by Van Nieuwerburgh and Veldkamp (2006), on experimentation in Bachman and Moscarini (2011), on contracting by Narita (2011) and on search by Petrosky-Nadeau (2011).

overall 4-digit (details in Appendix A). So our instrument is a continuous variable from 0 to 1 reflecting the share of the 4-digit industry output covered by quotas before 2005.¹⁶

Table 3 starts in column (1) with the baseline results (copied from column (1) of Table 2) which shows that within-industry dispersion of TFP rises when industry growth rates fall. In column (2) we report the same results for the subset of textiles and apparel industries which is our China WTO quota instrument sample, and again see a significant negative relationship. Column (3) is our key result, which shows that once we instrument output growth at the industry level using the abolition of quotas the relationship with uncertainty becomes insignificant and in fact turns positive. This shows that - at least in this sample - changes in demand did not *cause* changes in the cross-sectional dispersion of plants. Finally, in column (4) we report the reduced form specification, again showing no significant causal impact of quota abolition on the dispersion of TFP shocks.

One obvious question is to what extent our IV estimates could be biased if the exclusion restriction fails, for example due to correlations between the instrument and other drivers of uncertainty? One potential concern is that the abolition of quotas should increase exit, trimming the left tail of plants in the US, causing a downward bias in our measure of uncertainty after the trade-shock. But we use the sample of plants with 25+ years of data so this result is not driven by compositional change.¹⁷ Another concern is that this trade reform itself should generate uncertainty. But this bias works against our results as it should generate a negative bias, making our positive coefficient even more striking. Finally, maybe our results are only valid for this sample in response to this trade-induced first-moment shock. So in columns (4) and (5) we use a different instrumental variable approach, which exploits movements in industry exchange rates to identify changes in industry growth.

To generate industry exchange rates we follow Bertrand (2004) in defining for each industry its trade weighted exchange rate for every major trade destinations. For example, if industry A trades 75% with Canada and 25% with Mexico this industry's exchange rate would be a 75% weight on the US-Canada rate and a 25% weight on the US-Mexico exchange rate (details in Appendix A). Using these exchange rates in an industry panel with industry and time fixed-effects allows us to identify variations in industry demand from the differential changes in exchange rates they face due to their differential mix of country trading partners.

Looking at Table 3, column (5), we see from the first stage (in the bottom panel) that industry exchange rates are effective at predicting industry output growth. But looking at the second stage (in the top panel) we again find no significant impact of demand growth on within-industry variance of TFP shocks. In column (6) we report the reduced form

¹⁶ All of the quotas were binding (fill rates above 95%). Hence, quota abolition had a large impact on exports from China, which rose by 240% on average between 2005 and 2006 in affected industries.

¹⁷In fact results for a completely balanced panel also show no significant relationship either.

regression again finding nothing significant.

So in summary, Table (3) presents evidence from two different instrumental variable estimations that first-moment demand shocks do not seem to drive second-moment increases in within-industry TFP spreads. These results of course are conditional on the data and validity of the experiment, but they are suggestive that uncertainty is not primarily driven by negative first-moment shocks, but instead moves independently. To us this seems very intuitive - the types of shocks that cause slower industry level growth are also likely to increase (at least in the short-run) industry-level uncertainty.

2.4 Are establishment level TFP shocks a good proxy for uncertainty?

The evidence we have provided for counter-cyclical aggregate and industry level uncertainty relies heavily on using the dispersion of establishment level TFP shocks as a measure of uncertainty. To cross check this Table 4 compares our establishment TFP shock measure of uncertainty with other measures of uncertainty - primarily the volatility of daily and monthly firm-stock returns - which have been used commonly in the prior uncertainty literature.¹⁸

In Column (1) we regress the mean absolute size of the TFP shock in the plants of public traded firms against their parent firms within year volatility of monthly stock-returns (plus a full set of firm and year fixed-effects). The positive and highly significant coefficient reveals that when plants of publicly quoted firms have large positive or negative TFP shocks in any given year their parent firms are likely to have significantly more volatile stock-returns over the course of that year. This is reassuring for both our TFP shock measure of uncertainty and stock-market volatility measures of uncertainty, as while neither measure is perfect, the fact that they are strongly correlated suggests that they are both proxying for some underlying measure of firm-level uncertainty. In column (2) we use daily returns rather than monthly returns and find similar results, while in column (3) following Leahy and Whited (1996) we leverage adjust the stock-returns and again find similar results.

In column (4) we compare instead the within-year standard-deviation of firm quarterly sales growth against the absolute size of their establishment TFP shocks. We find again a strikingly significant positive coefficient, showing firms with a wider dispersion of TFP shocks across their plants tend to have more volatile sales growth within the year. Finally, in column (5) we generate an industry level measure of output volatility within the year by taking the standard-deviation of monthly production growth, and find that this measure is also correlated to the average absolute size of establishment level TFP shocks within the

¹⁸See, for example, Leahy and Whited (1996), Schwert (1989), Bloom, Bond and Van Reenen (2007) and Panousi and Papanikolaou (2011) who all use firm daily stock-returns volatility as an empirical measure of firm uncertainty.

industry in that year.

So in summary, establishment level TFP shocks are larger when the parent firms have more volatile stock returns and sales growth within the year, and the overall industry has more volatile monthly output growth within the year. This suggests these are all picking up some type of stochastic volatility process for uncertainty, which we will model in section (3).

2.5 Macroeconomic Measures of Uncertainty

The results discussed so far focus on establishing the countercyclically of idiosyncratic (establishment, firm and industry) uncertainty. Looking instead at macroeconomic uncertainty there is already a growing literature providing evidence that this is also counter-cyclical, for example Schwert (1989), Campbell et al. (2001) and Engel and Rangel (2008).

Rather than repeat this evidence here we simply include one additional model specific empirical measure of aggregate uncertainty, which is the conditional heteroskedasticity of productivity, A_t . This is estimated using a GARCH(1, 1) estimator on the Basu, Fernald and Kimball (2006) data on quarterly TFP updated from 1972Q1 to 2009Q4. We find that conditional heteroskedasticity of TFP is strongly countercyclical, rising by 98% during recessions which is highly significant (a t-statistic of 6.26).¹⁹

3 The General Equilibrium Model

We proceed by analyzing the quantitative impact of variation in uncertainty within a dynamic stochastic general equilibrium model where heterogeneous firms are subject to both first and second moment shocks. In the model, each firm uses capital and labor to produce a final good. Firms that adjust their capital stock and employment incur non-convex adjustment costs. As is standard in the RBC literature, firms are subject to an exogenous process for productivity. We assume that the productivity process has an aggregate and an idiosyncratic component. In addition to these first-moment shocks, we allow the second moment of the innovations to productivity to vary over time. That is, shocks to productivity can be fairly small in normal times, but become potentially large when uncertainty is high.

¹⁹We also estimated a GARCH(1, 1) for monthly industrial production, including as many as twelve lags and find very similar results. We also experimented with different specifications – such as ARCH(1) or using GDP growth rates – and results again were very similar.

3.1 Firms

3.1.1 Technology

The economy is populated by a large number of heterogeneous firms that employ capital and labor to produce a single final good. We assume that each firm operates a diminishing returns to scale production function with capital and labor as the variable inputs.²⁰ Specifically, a firm indexed by j produces output according to

$$y_{j,t} = A_t z_{j,t} k_{j,t}^{\alpha} n_{j,t}^{\nu} , \ \alpha + \nu < 1.$$
(5)

Each firm's productivity is a product of two separate processes: aggregate productivity, A_t , and an idiosyncratic component, $z_{j,t}$. Both the macro- and firm-level components of productivity follow autoregressive processes as noted in equations (2) and (3). We depart from the benchmark RBC model in that we allow the variance of innovations to the productivity processes, σ_t^A and σ_t^Z , to vary over time as noted in equations (??) and (??).

3.1.2 Capital and Labor Adjustment Costs

We allow for the presence of various types of convex and non-convex adjustment costs in capital and labor.²¹ With respect to capital, we assume that a firm's capital stock evolves according to the standard law of motion

$$\gamma k_{j,t+1} = (1 - \delta_k) k_{j,t} + i_{j,t} \tag{6}$$

where the $\gamma - 1$ is the trend growth rate of output and δ_k is the rate of capital depreciation. The first adjustment cost we allow for involves a non-convexity — conditional on undertaking an investment, a fixed cost F^K is incurred independently of the scale of investment. The second capital adjustment cost we consider is a partial irreversibility. Resale of capital occurs at a price that is only a share (1 - S) of its purchase price.

Similarly, we assume that the law of motion for hours worked is governed by

$$n_{t,t} = (1 - \delta_n)n_{j,t-1} + s_{j,t}.$$
(7)

At each period a constant fraction δ_n of hours worked is exogenously destroyed due to retirement, illness, maternity leave, exogenous quits, etc. Whenever the firm chooses to

²⁰An alternative model has a setup of monopolistically competitive firms in which each firm produces a differentiated good. Note that the assumption of decreasing returns to scale implies that there is a fixed factor of production that pins down firm size.

²¹See the literature focused on estimating labor and capital adjustment costs, including, Nickell (1986), Caballero and Engel (1999), Ramey and Shapiro (2002), Hall (2004), Cooper and Haltiwanger (2006), Merz and Yashiv (2007), and Bloom (2009).

adjust its stock of hours relative to $(1 - \delta_n)n_{j,t-1}$, it incurs a fixed cost F^L independently of the size of the change in hours. We also allow for hiring and firing costs which represent, for example, variable interviewing and training costs or severance packages. In our model, we assume that this cost is identical for hiring and firing and expressed as a share H of the annual wage bill per worker.

3.1.3 The Firm's Value Function

We denote by $V(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ the value function of a firm. The seven state variables are given by (1) a firm's capital stock k, (2) a firm's hours stock from the previous period n_{-1} , (3) the firm's idiosyncratic productivity $z_{j,t}$, (4) aggregate productivity A_t , (5) macro uncertainty σ_t^A , (6) micro uncertainty σ_t^Z and (7) the joint distribution of idiosyncratic productivity and firm-level capital stocks and hours worked in the last period μ_t , which is defined for the product space $S = Z \times R_+ \times R_+$.

The dynamic problem of the firm consists of choosing investment and hours to maximize the present discounted value of future profits

$$V(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu) =$$

$$\max_{i,n} \begin{cases} y - w(A, \sigma^{A}, \sigma^{Z}, \mu)n - i \\ -AC^{k}(k, k') - AC^{n}(n_{-1}, n) \\ +\mathbb{E}\left[m\left(A, \sigma^{A}, \sigma^{Z}, \mu; A', \sigma^{A'}, \sigma^{Z'}, \mu'\right)V(k', n, z'; A', \sigma^{A'}, \sigma^{Z'}, \mu')\right] \end{cases}$$
(8)

given a law of motion for the joint distribution of idiosyncratic productivity, capital and hours,

$$\mu' = \Gamma(A, \sigma^A, \sigma^Z, \mu), \tag{9}$$

and the stochastic discount factor, m. We denote by $AC^k(k, k')$ and $AC^n(n_{-1}, n)$ the capital and labor adjustment cost functions, respectively. $K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ and $N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ denote the policy rules associated with the firm's choice of capital for the next period and current demand for hours worked.

3.2 Households

The economy is populated by a large number of identical households that we normalize to a measure one. Households choose paths of consumption, labor supply, and investments in firm shares to maximize lifetime utility. We use the measure ϕ to denote the one-period shares in firms. The dynamic problem of the household is given by

$$W(\phi, A, \mu) = \max_{\{C, N, \phi'\}} \left\{ U(C, N) + \beta \mathbb{E} \left[W(\phi', A', \mu') \right] \right\}$$
(10)

subject to the law of motion for μ and a sequential budget constraint

$$C + \int q(k', n, z; A, \sigma^{A}, \sigma^{Z}, \mu) \phi'(dk dn dz)$$

$$\leq w(A, \sigma^{A}, \sigma^{Z}, \mu) N + \int \rho(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu) \phi(dk dn dz).$$

$$(11)$$

Households receive labor income as well as the sum of dividends and the resale value of their investments, $V(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$. With these resources the household consumes and buys new shares at a price $q(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ per share of the different firms in the economy. We denote by $C(\phi, A, \mu)$, $N^s(\phi, A, \mu)$, $\Psi(k', n, z; A, \sigma^A, \sigma^Z, \mu)$ the policy rules determining current consumption, time worked, and quantities of shares purchased in firms that begin the next period with a capital stock that equals k' and who currently employ n hours, respectively.

3.3 Recursive Competitive Equilibrium

A recursive competitive equilibrium in this economy is defined by a set of quantity functions $\{C, N^s, \Psi, K, N^d\}$, pricing functions $\{w, q, \rho, m\}$, and lifetime utility and value functions $\{W, V\}$. V and $\{K, N^d\}$ are the value function and policy functions solving (8) while W and $\{C, N^s, \Psi\}$ are the value function and policy functions solving (10). There is market clearing in the asset markets

$$\Psi(k', n, z; A, \sigma^A, \sigma^Z, \mu) = \mu(z, k', n)$$
 for every triplet $(z, k', n) \in S$,

the goods market

$$= \int_{S} \left[\begin{array}{c} Azk^{\alpha}N^{\nu}(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu)^{\nu} - \left(K(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu) - (1 - \delta_{k})k\right) \\ -AC^{k}(k, K(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu)) - AC^{n}(n_{-1}, N(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu)) \end{array} \right] \\ \mu \left(dkdndz\right),$$

and the labor market

$$N^{s}(\phi, A, \mu) = \int_{S} \left[N^{d}(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu) \right] \mu \left(dk dn dz \right).$$

Finally, the evolution of the joint distribution of z, k and n is consistent. That is, $\Gamma(A, \sigma^A, \sigma^Z, \mu)$ is generated by $K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$, $N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$, and the exogenous stochastic evolution of A, z, σ^Z and σ^A with the appropriate summation of firms' optimal choices of capital and hours worked given current state variables.

3.4 Sketch of the Numerical Solution

The model can be simplified substantially if we combine the firm and household problems into a single dynamic optimization problem as in Kahn and Thomas (2008). From the household problem we get

$$w = -\frac{U_N(C,N)}{U_C(C,N)} \tag{12}$$

$$m = \beta \frac{U_C(C', N')}{U_C(C, N)}$$
(13)

where equation (12) is the standard optimality condition for labor supply and equation (13) is the standard expression for the stochastic discount factor. To ease the burden of computation it is useful to assume that the momentary utility function for the household is separable across consumption and hours worked,

$$U(C_t, N_t) = \frac{C_t^{1-\eta}}{1-\eta} - \theta \frac{N_t^{\chi}}{\chi},$$
(14)

implying that the wage rate is a function of the marginal utility of consumption,

$$w_t = \phi N_t^{\chi - 1} \frac{\theta}{C_t^{-\eta}}.$$
(15)

Kahn and Thomas (2008) and Bachmann, Caballero and Engel (2008) define the intertemporal price of consumption goods as $p(A, \sigma^Z, \sigma^A, \mu) \equiv U_C(C, N)$. Using this approach, we can redefine the firm problem in terms of marginal utility, denoting the new value function as $\tilde{V} \equiv pV$. The firm problem can then be expressed as

$$\tilde{V}(k, n_{-1}, z; A, \sigma^{A}, \sigma^{Z}, \mu) = \max_{\{i,n\}} \left\{ \begin{array}{l} p(A, \sigma^{A}, \sigma^{Z}, \mu) \left(y - w(A, \sigma^{A}, \sigma^{Z}, \mu)n - i - AC^{k}(k, k') - AC^{n}(n_{-1}, n) \right) \\ + \beta \mathbb{E} \left[\tilde{V}(k', n, z'; A', \sigma^{A'}, \sigma^{Z'}, \mu') \right] \end{array} \right\} (16)$$

We employ non-linear techniques that build upon Krusell and Smith (1998). Specifically, we summarize μ with a small set of moments of the firm distribution which we denote by Ω . Specifically, in each iteration we preform four steps. We first forecast the intertemporal price \hat{p} and next period's moments $\hat{\Omega}'$ as functions of the current aggregate state:

$$\hat{p} = f_1^{(l)}(A, \sigma^A, \sigma^Z, \Omega)$$
$$\hat{\Omega}' = f_2^{(l)}(A, \sigma^A, \sigma^Z, \Omega)$$

Assuming that $\chi = 1$ (we discuss the choice of this parameter value below) we get for a given forecast of \hat{p} , the current period wage w from (15). We can then find the value function \tilde{V}^l associated with those forecasting functions by solving (16) substituting the approximated state Ω for the joint distribution μ and $f_2^{(l)}$ for the law of motion Γ . We then simulate the economy for many periods during which the forecasting rule for the intertemporal price is not used. Rather, in each period the market clearing price p_t is calculated as the price that combines firm optimization and goods market clearing. For a given price, the simplified firm optimization problem becomes

$$\max_{\{i,n\}} \left\{ p\left(y - wn - i - AC^k(k,k') - AC^n(n_{-1},n)\right) + \beta \mathbb{E}\left[\tilde{V}^l(k',n,z';A',\sigma^{A\prime},\sigma^{Z\prime},\hat{\Omega}')\right] \right\}$$

which uses the value function calculated in the second step and the moment forecasting function from the first step. Market clearing is achieved when aggregation of the optimal policies from this problem yield market clearing in the goods market

$$C = \int \left(y + i - AC^k - AC^n \right) \mu(dkdndz).$$

This simulation yields sequences of exogenous states $\{A_t, \sigma_t^A, \sigma_t^Z\}$ states, prices $\{p_t\}$ and moments $\{\Omega_t\}$. As the final step we then update the forecasting functions $f_1^{(l+1)}$ and $f_2^{(l+1)}$ from the observed moments and equilibrium prices and restart the algorithm at the first step. We iterate until the forecasting functions converge.²²

4 Simulation

This section motivates the choice of parameter values used in the simulations (see Table 2) and also presents simulation results for our preferred specification.

4.1 Calibration

4.1.1 Frequency and Preferences

We set the time period to equal a quarter and the household's discount rate, β , is calibrated to 0.985. η is set equal to one which implies that the momentary utility function features an elasticity of intertemporal substitution of one. Following Kahn and Thomas (2008) and Bachmann, Caballero and Engel (2008) we make the simplifying assumption that the Frisch

²²For the forecasting functions, we use the aggregate productivity state, the first moments of the distribution over capital and labor as well as the aggregate uncertainty state. Interestingly, this provides a very good fit and an R² of above 0.985. Additional moments of the distributions over capital, labor and idiosyncratic productivity could be added to the forecasting functions. We are currently working on extending the model in this way, but the computational burden quickly becomes very large.

labor supply elasticity is infinite, corresponding to $\chi = 1$. This assumption implies that we do not need to forecast the wage rate in addition to the forecast of p because when $\chi = 1$ we get

$$w_t = \frac{\theta}{C_t^{-1}} = \frac{\theta}{p}$$

Hence, once we construct a forecast for p, we immediately obtain a forecast for w, eliminating the need to forecast it separately and simplifying the computational problem. We set the parameter θ such that households spend a third of their time working in the non-stochastic steady state. The trend growth rate of per capita output is set to equal 1.6% annually.

4.1.2 Production Function, Depreciation, and Adjustment Costs

We set δ_k to match a 10% annual capital depreciation rate. The annual exogenous quit rate of labor is a key parameter set to 15%. This estimate is based on the quit rate reported in the Bureau of Labor Statistic JOLTS data.²³ We set the exponents on capital and labor in the firm's production to be $\alpha = 0.25$ and $\nu = 0.5$, consistent with a capital cost share of 1/3 and a 33% markup when the firm faces an iso-elastic demand curve.

The existing literature provides a wide range of estimates for capital and labor adjustment costs.²⁴ We set our adjustment cost parameters to match Bloom (2009), which to our knowledge is the only paper that jointly estimates capital and labor convex and non-convex adjustment costs. Fixed costs of capital adjustment are set to 1.5% of annual sales, and the resale loss of capital amounts to 40%. The fixed cost of adjusting hours, is set to 2.1% of annual wages, and the hiring and firing costs equal 1.8% of annual wages.

4.1.3 Aggregate and Idiosyncratic TFP Processes

Productivity both at the aggregate and the idiosyncratic level is determined by AR1 processes as specified in equations (2) and (3). The serial autocorrelation is taken directly from Khan and Thomas (2008) and adjusted to the quarterly frequency. Hence, ρ^A and ρ^Z are set to yield an annual persistence parameter of 0.859. In our model, the variance of the innovations to these processes is time-varying. The exact calibration is presented in some detail in the subsequent paragraph, but on average, σ_t^A and σ_t^Z are set to 1.59% and 8.50%. We ap-

²³ JOLTS stands for Job Openings and Labor Turnover Data, which the BLS has been collecting since January 2001. Hence, this data spans two NBER defined recessions. It distinguishes between quits, layoffs, and other separations. Our figures are seasonally adjusted for total private employment. In JOLTS, the monthly quit figure varies between 1.6% and 2.4%, with the lowest value occurring in November 2008 during the depths of the recent recession. Annualizing the November 2008 quit rate we get a value of 19.2%. Our calibration using a lower value of 15% is thus a conservative calibration.

²⁴See, for example, Hayashi (1982), Nickel (1986), Shapiro (1986), Caballero and Engel (1999), Ramey and Shapiro (2001), Hall (2004), Cooper, Haltiwanger and Willis (2004), Cooper and Haltiwanger (2006) as well as Mertz and Yashiv (2007).

proximate the autoregressive processes with Markov chains. The support for the processes are set to include three standard deviations on either side of the mean.

4.1.4 The Calibrated Process for Uncertainty

In the benchmark calibration we assume that the uncertainty process is independent of the first-moment shocks. This implies that we are not artificially creating the drop in economic activity following a second-moment shock by correlating it with the first-moment shock.

We assume for simplicity that the stochastic volatility processes, σ_t^A and σ_t^Z , each follow a two-point Markov chain

$$\sigma_t^A \in \left\{ \sigma_L^A, \sigma_H^A \right\} \quad \text{where } Pr(\sigma_{t+1}^A = \sigma_j^A | \sigma_t^A = \sigma_k^A) = \pi_{k,j}^{\sigma A}$$
(17)

$$\sigma_t^Z \in \{\sigma_L^Z, \sigma_H^Z\} \quad \text{where } Pr(\sigma_{t+1}^Z = \sigma_j^Z | \sigma_t^Z = \sigma_k^Z) = \pi_{k,j}^{\sigma_Z}$$
(18)

Since we cannot directly observe the stochastic process of uncertainty in the data the calibration has to be guided the impact of uncertainty on observable cross sectional and aggregate time series moments. There are eight parameters that need to be calibrated: σ_L^A , σ_H^A , σ_L^Z , σ_H^Z , $\pi_{L,H}^{\sigma A}$, $\pi_{H,L}^{\sigma A}$, $\pi_{L,H}^{\sigma Z}$ and $\pi_{H,L}^{\sigma Z}$. The empirical section (2) suggests that uncertainty at the micro and at the macro level is highly correlated. So as a simplification to ease computational constraints we assume in the benchmark calibration that a single process determines the economy's uncertainty regime. This reduces the number of parameters to six: : σ_L^A , σ_H^A , σ_L^Z , σ_L^Z , $\pi_{L,H}^\sigma$ and $\pi_{H,L}^\sigma$ since σ^A and σ^Z follow the same Markov process (with different levels, of course).

With these six parameters we try to match eight moments. The first four are based on the cross-sectional IQR of sales growth rates in the establishment data, which is our largest and most representative micro dataset. We calculate the mean, standard deviation, skewness and serial correlation of this IQR time series, which characterizes how much microuncertainty changes over time. These four moments are reported in the first column in the top panel of Table 3. The second set of moments is based on the GARCH(1,1) estimated conditional heteroskedasticity of GDP growth shown in Figure 10. We again calculate the mean, standard deviation, skewness and serial correlation of this series, which characterizes how much macro-uncertainty changes over time. The first column in the lower panel of Table 3 reports these four moments.

We calibrate our uncertainty parameters by aligning these eight moments from the actual data with the simulated data. In particular, in order to get the first four micro moments we run the simulation and then aggregate the firm-level values to an annual value. Using this we construct the cross sectional IQR of these growth rates, and then generate the mean, standard-deviation, skewness and serial correlation of this series. Similarly, in order

to get the four macro moments we simulate the model, estimate a GARCH(1,1) process on log(GDP) and four lags, and generate the mean, standard-deviation, skewness and serial correlation of this. We thus try to match the simulated counterpart of each empirical moment to calibrate the underlying uncertainty process. In order to understand how robust the estimates are, we recalculate the exact same eight moments using slightly different data samples. Specifically, the second and third column in Table 3 show the minimum and maximum value that was recorded for each moment when either the first or last five years are dropped from the sample. This immediately highlights that the skewness is obviously difficult to measure in short samples and varies significantly in between variations of the sample.

Based on our preferred calibration we find that periods of high uncertainty occur with a quarterly probability of slightly below 5%. The period of heightened uncertainty is quite persistent with a quarterly probability of 88.5% of staying in the high uncertainty state. Idiosyncratic volatility is set to a lower value of 6.7%, but almost doubles in the heightened uncertainty state. Aggregate volatility is at a low of 0.81%, but more than quadruples when an uncertainty shock hits. The baseline calibration yields moments that are shown in column four of Table 3. We are obviously constrained by having to match eight moments with only six parameters. As a result, some of our moments are difficult to be matched both at the micro and macro level. Specifically, the skewness is too low at the micro and too high at the macro level. Conversely, the serial autocorrelation is too high for the simulated micro process at the micro too low at the macro level. Even though the fit is not perfect we find the results to be surprisingly in accordance with the data.

The last two columns of Table 3 report the simulated moments for two alternative specifications of the model. In column five, σ^A does not vary and is instead set to its long run average value. Interestingly, the resulting micro moments are very similar to those obtained in the benchmark calibration. Turning off macro uncertainty, for instance, only slightly reduces the estimated standard deviation of the micro measure from 3.53% to 3.34%. Analogously, the calibration in column six sets σ^Z to its long run average and again there seems to be only a small effect on the resulting macro moments. While the two channels are not completely independent of each other, the cross effects are rather small. This suggests that our approach of linking the process of micro uncertainty to dispersion measures from establishment data and macro uncertainty to the conditional heteroskedasticity of aggregate output is reasonable.

4.2 The Effects of an Uncertainty Shock

We first study the effects of an isolated increase in uncertainty. We simulate the model economy 500 times with 4000 firms, and let each simulation run for 250 periods to initialize

the distribution over z, k and n. We then force uncertainty to be low for 10 periods. Finally we induce an uncertainty shock in time period zero. This is to mimic a typical business cycle shock to uncertainty that occurs after a period of low uncertainty. We average over the 500 simulated economies to yield an average effect of a business cycle sized increase in uncertainty.

Before showing the simulation results it is important to note that the unweighted average of firm level productivity $\sum_{j}^{N} (A_t z_{j,t})/N$ does not vary in this experiment as shown in Figure 12. Thus, the results shown on the figures below are driven entirely by changes in uncertainty. This uncertainty shock leads to three phases of activity:

The drop: When uncertainty rises in period zero, the investment and hiring thresholds move out as the real options effect leads firms to defer spending and hiring projects. That is higher uncertainty makes firms cautious as they don't want to make a costly hiring or investment mistake. So because most firms pause hiring there is an immediate fall in aggregate hours worked. The fall in aggregate hours worked manifests itself into a drop in aggregate output, as can be seen in Figure 13. Output falls by almost 2% upon impact and continues to fall for another quarter to a low of about 2.1% from the effects of further falls in hours and the drop in capital from the pause in investment in period 0. Both hours and capital falls are driven by exogenous quits and depreciation - that is workers leaving (for retirement, sickness, maternity etc.) and capital depreciating that are not replaced.

The rebound: By the second quarter uncertainty has fallen sufficiently, and firms built up enough pent-up demand for hiring and investment that output begins to rebound. During this transition the thresholds slowly begin to move back in as more economies leave the high uncertainty state. At the same time the distribution of firm specific productivity fans out, so that more and more firms begin hitting the new, wider thresholds, accelerating the rebound.

The overshoot: By the fourth quarter the economy has rises above its long-run trend for a few quarters before returning to it's long-run average. The reason for this overshooting is that many firms are bunched near their Ss investment and hiring thresholds due to depreciation, labor attrition and trend growth. So small increases in productivity causes firms to hit those hiring and investment thresholds, while small decreases in productivity move them towards the interior of their (S, s) bands. As a result, the increased variance of idiosyncratic productivity shocks induced by higher uncertainty increases aggregate medium-run hiring and investment. That is, most firms that receive a positive productivity shocks hit their (S, s) bands and invest and hire, while most firms that receive a negative shock move to the interior of the (S, s) bands and do nothing.

In the top left panel of Figure 14 we show hours over the cycle which fall immediately after the shock arrives because we assume hiring happens instantaneously (capital has one quarter time to build). So hours fall by about 2.4% on impact, begin to recover after two periods and overshoot from 4 onwards. The uncertainty shock also induces a similar drop and subsequent rebound in investment, as shown in the top right panel of Figure 14.

The lower right hand panel plots the time profile of consumption. When the uncertainty shock occurs in period zero, consumption jumps up immediately and then falls below trend for about three quarters. The reason for the initial spike in consumption is that the freeze in investment and hiring reduces the resources spent on capital and labor adjustment. Since the interest rate drops upon impact of the shock, consumers are signaled that consumption is cheap, which leads to an increase in consumption in period zero. In other words, even though consumers know they face higher uncertainty in the future and they would like to save more, they do not increase savings in the first period because the returns to saving have become (temporarily) low and very risky.²⁵

Finally, the lower left hand panel plots the value of aggregate productivity, defined as $\sum_{j}^{N} (A_t z_{j,t} n_{j,t}) / \sum_{j}^{N} n_{j,t}$. Productivity also has a clear drop and subsequent rebound following the uncertainty shock, despite the fact that the average micro and macro productivity shocks are unchanged, as shown in Figure 12. The reason is that uncertainty freezes the reallocation of capital and labor from low- to high-productivity firms. In normal times, unproductive firms contract and productive firms expand, helping to maintain high productivity levels. When uncertainty is high, firms reduce expansion and contraction, shutting off much of this productivity-enhancing reallocation, leading to a fall in productivity growth rates. When uncertainty reverts back to normal, firms rapidly address their pent-up demand for reallocation so that productivity returns to its long-run trend.

A second-moment shock induces a fall in investment, hours, and output. Intriguingly, aggregate productivity falls even though this is an effect rather than a cause of the drop in economic activity. Consumption also exhibits a fall from quarter two onwards, although there is an initial one-period jump in the basic model.

4.3 Model Simulations

4.3.1 Business cycle statistics

We have shown that our model can generate expansions and contractions in response to an increase in uncertainty. One natural question is whether this success comes at a cost of the model's ability to generate empirically recognizable business fluctuations. That is, can

²⁵This logic suggests that if we extended the model to allow for some alternative savings technology – for example, inventories or savings abroad – this initial spike in consumption would disappear as the representative consumer would just increase savings through this channel when the uncertainty shock hit to reduce the subsequent drop in consumption. Due to computational constraints we cannot currently increase the state space of the model in this way.

the model, when calibrated with the same parameters as used in the experiments discussed so far, generate levels of comovement and volatility of macroeconomic aggregates that are empirically plausible? To answer this question, we simulate our model for 4000 periods and compute the standard set of business cycle statistics.

Table 4 illustrates that the calibration generates second-moment statistics (panel 2) that resemble their empirical counterparts in U.S. data (panel 1). Investment is more volatile than output, while consumption is less volatile. Investment and hours comove with output. Similar to many variants of the RBC model, hours are not as volatile relative to output in the model as they are in the data. Panel 3 of the same table reports the same statistics for an alternative calibration of the model without aggregate productivity shocks and an otherwise identical calibration. While this is only a crude experiment, it is interesting to note that time-variation in micro uncertainty can produce about 30% of the volatility in output that we find using the benchmark model

4.3.2 Establishment Level Moments

Time averaged moments of establishment level investment rates provide an alternative way to evaluate the empirical realism of our model. Table 5 illustrates this idea. The first row contains a set of moments that are taken directly from previous work by Cooper and Haltiwanger (2006). They construct annual investment rates using data on capital retirements and investment for a balanced panel of plants from the Longitudinal Research Database (LRD). This table groups plants as investing (investment rate greater or equal than 0.01), showing an investment spike (i/k greater than 0.2) or inactive (absolute i/k < 0.01).

The rest of the table shows those same moments from simulations of our baseline model aggregated to the annual frequency. The four rows within each panel refer to different levels of unit aggregation. For example, the second row in panel 2 shows results for plants that contain 10 independent production lines. The reason for displaying various levels of aggregation is that the establishments in the Cooper and Haltiwanger (2006) are a balanced panel of large continuing plants with an average of almost 600 employees, compared to about 10 employees for the average of all Census establishments. Hence, in some senses these moments are for more aggregated production units.

Our model can match those moments relatively well. Depending on the exact degree of unit aggregation, our model can easily match the share of plants that are inactive, investing and showing an investment spike. As alternatives to further improve the fit the model could be extended to allow small adjustments to the capital and labor stock without adjustment costs as in Khan and Thomas (2008) or by introducing maintenance investment as in Bachmann, Caballero and Engel (2008).

4.4 Inspecting the Mechanism

We report in this subsection five experiments that are meant to highlight the different forces at work in our model. First, we consider alternative specifications of adjustment costs. Here, we simulate models with only capital adjustment costs, only labor adjustment costs and no adjustment costs at all. Second, we are interested in the differential effect of uncertainty at the macro and micro level. We thus simulate economies with time variation in only one of the two.

Adjustment Costs: To highlight the differential effects of adjustment costs in capital and labor, we use alternative calibrations that turn off adjustment costs in either capital or labor. Figure 15 illustrates the effects of an uncertainty shock for those alternative cases. In the model with adjustment costs in capital only, there is no impact of an uncertainty shock in period 0 because of the time-to-built assumption in capital. Investment, however, falls immediately and output thus falls in the subsequent period. The drop in output is smaller at a maximum of almost 1% as opposed to 2% in the baseline model. A model with adjustment costs in labor only, produces a slightly bigger fall in aggregate output by about 1.3%. In this case, the drop occurs in the period the uncertainty hits as the hiring freeze affects the labor input in production in the same period.

Figure 15 also shows that when there are no adjustment costs of any type in the economy, economic activity actually increases following an uncertainty shock. The reason for this result is related to the Hartman (1976) and Abel (1983) effect whereby a higher variance of productivity increases investment, hiring and output because the marginal revenue product of capital and labor is convex in productivity.²⁶

Only Micro or Only Macro Uncertainty: We perform two additional experiments where we consider an economy that exhibits either only micro or only macro uncertainty. In those experiments, we turn off time variation in one of the two uncertainty processes.²⁷ Those experiments are reported in Figure 16. As the figure suggests the presence of both macro and micro uncertainty magnifies the effect of the uncertainty shock. Note however that each shock separately induces a similar effect on output.

The main difference between micro and macro uncertainty shocks is in the recovery and overshoot. The recovery is faster and the overshoot larger in the case of only micro uncertainty relative to the only macro uncertainty case. The reason for this is that is it the

²⁶To be precise if $Y = AK^a L^b$ and the per period rental cost of capital is r and labor is w, then the without adjustment costs the optimal choice of K and L are $K^* = \phi_1 A^{\frac{1}{1-a-b}}$ and $L^* = \phi_2 A^{\frac{1}{1-a-b}}$ where ϕ_1 and ϕ_2 are functions of a, b, r and w. Hence, it is clear that K^* and L^* are convex in A so that higher variance in A will increase the average levels of K and L, which is commonly known as the Hartman-Abel effect after it was pointed out in Hartman (1972) and refined by Abel (1983).

²⁷Those experiments are indeed identical to the ones mentioned in column five and six of Table 3 in the calibration section.

cross-sectional spread of shocks that generates the overshoot for the baseline uncertainty simulation arises primarily from the micro-uncertainty. So with macro uncertainty there is very little overshoot.

5 Policy in the Presence of Uncertainty

In this section, we analyze the effects of stimulative policies in the presence of uncertainty shocks. It is important to emphasize that any such policy is *not* optimal within the context of our model as the competitive equilibrium is Pareto optimal. Rather, we see our policy experiment as a means of documenting that in this framework a given policy can be less effective in times of heightened uncertainty. The reason is that during times of increased uncertainty, firms are far away from their hiring and investment thresholds, making them less responsive to the policy stimulus. Our quantitative model thus allows us to shed some light on the *effectiveness* of the policy as opposed to its *desirability*.

We are interested in a policy that attempts to temporarily stimulate hiring by reducing the effective wage paid by firms. More specifically, the policy consists of an unanticipated 1% wage bill subsidy paid for a period of two quarters. We simulate this policy impulse once during an uncertainty shock and also in an economy that is not hit by an uncertainty shock. By comparing the effect in those two cases, we can attempt to identify the dampening effect of uncertainty on policy effectiveness.²⁸

Figure 17 illustrates our experiment and depicts the response of output to a temporary wage subsidy that takes place at period zero. The green line (marked with a cross) refers to the case with no accompanying uncertainty shock. Not surprisingly, the artificially reduced wage stimulates hiring and increases output for two quarters which then gradually returns to its long run trend. The red line (squares) shows the response of output to such a policy when the policy is introduced at the exact time that the uncertainty shock hits the economy. The uncertainty shock will still result in an immediate drop in output and a subsequent rebound an overshot. To ease comparison, the figure also shows the effect of an uncertainty shock on output in the absence of any policy (blue, circles). The difference between the second and the third line can be interpreted as the policy's impact policy during an uncertainty shock.

Figure 18, compares the policy's impact on the same scale. The blue line (marked with a circle) shows the marginal impact of the wage subsidy with no accompanying shock and the green line (crosses) shows the marginal impact with uncertainty. As it is clear from the figure, the presence of uncertainty substantially reduces the effects of such a policy relative to an economy that is in the normal, or low uncertainty state. For example, in this

²⁸ In this version of the experiment, we abstract from balanced budget considerations. Those can obviously have important general equilibrium effects, but we are focusing here on the relative effectiveness of a stimulative policy in normal and highly uncertain times and not on the absolute size of the policy multiplier.

particular case the period 2 increase in hours is 1.2% with low uncertainty and 0.5% with high uncertainty, so that uncertainty reduces the impact of the policy stimulus by almost 60%.

To better understand the mechanism that is causing the policy ineffectiveness result, it is useful to look at Figure 19. This plots the hiring thresholds and the evolution of the cross-sectional distribution of firm level TFP for a given combination of k and n following an uncertainty shock. It is clear that on impact the hiring threshold jumps out.²⁹ As a result no firms are near the hiring threshold in period 1, so that any policies to reduce the cost of hiring - for example from a hiring subsidy - will have very little impact since no firms are close to their hiring threshold. That is uncertainty makes firms cautious so that they do not respond to policy stimulus. But, slowly over time however, these thresholds slowly fall back as uncertainty drops. At the same time the distribution of firm-level TFP fans out towards the thresholds. So that over time as uncertainty falls the responsiveness to policy rises.

Two messages arise from this experiment. First, in order for such a policy to have any effect on hiring (or investment) in the presence of uncertainty it has to be larger than the equivalent policy that would be implemented during normal times - in this particular simulation, for example, it would need to be 140% larger. Second, to avoid overshooting once uncertainty falls, the policy stimulus has to be abandoned as uncertainty reduces the short-run impact of policy much more than the medium and long run impact.

6 Conclusions

We propose uncertainty shocks as a new impulse helping to drive business cycles. First, we demonstrate that measured uncertainty is strongly countercyclical. This is true both at the aggregate and the industry level: slower industry growth is associated with higher industry uncertainty. Using trade and exchange rate instrumental variables we show that this slower industry growth is not causing the rise in uncertainty. Instead, uncertainty appears to be an exogenous process, suggesting recessions are driven by a combination of first and second moment shocks. We then build a DSGE model with heterogeneous firms, time-varying uncertainty and adjustment costs to quantify the impact of these second moment shocks. We find they typically lead to drops of about 2% in GDP, suggesting uncertainty could play an important role in driving business cycles. We also find that because uncertainty makes

²⁹ We find that the firing thresholds are not as reactive. This is due to the presence of exogenous labor attrition in the model that assures the firms that it can get "for free" firing in the model without having to pay the adjustment cost. Hence, absent very big shocks we find the firing thresholds to be more stable than those of hiring. Interestingly this result is consistent with the claims in Hall (2005) and Shimer (2005) that firing is acylical and that the majority of the cyclical adjustment in the labor force is done through a reduction in hiring.

firms cautious it substantially reduces the response of the economy to stimulative policy, leading to pro-cyclical multipliers.

References

Abel, A. (1983), "Optimal Investment under Uncertainty", American Economic Review, 73,228-233.

Abel, A.B. and Eberly, J.C. (1996), "Optimal Investment with Costly Reversibility", Review of Economic Studies, 63(4), 581-593.

Abraham, K.G. and Katz, L.F. (1986), "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?", Journal of Political Economy, 9(3), 507-522.

Alexopoloulos, M. and Cohen, J. (2009), "Uncertain times, uncertain measures", Toronto mimeo.

Arellano, C., Y. Bai, P. Kehoe, 2010, "Financial Markets and Fluctuations in Uncertainty", Minneapolis FED

Bachman, R. and Bayer, R. (2009), "Firm-specific productivity risk over the business cycle: facts and aggregate implications", Yale mimeo.

Bachman, R. and Bayer, R. (2011), "Uncertainty Business Cycles - Really?", NBER WP 16862.

Bachman, R., Caballero, R.J. and Engel, E.M.R.A. (2008), "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model", Yale mimeo.

Bachman, R. and Moscarini, G. (2010), "Business Cycles and Endogenous Uncertainty", Yale mimeo.

Basu, S. and Bundick, B. (2011), "Uncertainty Shocks in a Model of Effective Demand", Boston College WP 774.

Bayer, C and Bachman, R. (2009), "Firm specific productivity risk over the business cycle: facts and aggregate implications", Michigan mimeo.

Barlevy, G. (2004), "The Cost of Business Cycles Under Endogenous Growth", American Economic Review, 94(4), 964-990.

Basu, S., Fernald, J. and Kimball, M.S. (2006), "Are Technology Improvements Contractionary?", American Economic Review, 96(5), 1418-1448.

Bernanke, B. (1983), "Irreversibility, Uncertainty and Cyclical Investment", Quarterly Journal of Economics, 98, 85-106.

Bentolila, S. and Bertola, G. (1990), "Firing Costs and Labor Demand: How Bad is Eurosclerosis", Review of Economic Studies, 57(3), 381-402.

Berger, D. and J. Vavra, 2010, "Dynamics of the U.S. Price Distribution", Yale mimeo.

Bertrand, M. (2004), "From the Invisible Handshake to the Invisible Hand? How Import Competition Changes the Employment Relationship.", Journal of Labor Economics, 22(4), 723-765.

Bertola, G. and Caballero, R.J. (1994), "Irreversibility and Aggregate Investment", Review of Economic Studies, 61, 223-246.

Bloom, N. (2007), "Uncertainty and the Dynamics of R&D", American Economic Review Papers and Proceedings, 97(2), 250-255.

Bloom, N. (2009), "The Impact of Uncertainty Shocks", Econometrica, May.

Bloom, N., Bond, S. and Van Reenen, J. (2007), "Uncertainty and Investment Dynamics", Review of Economic Studies, April.

Bloom, N., Draca, M. and Van Reenen, J. (2011), "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and productivity", NBER WP 16717.

Brambilla, I., Khandewal, A. and Schott, P. (2010) "China's Experience under the Multi-Fiber Agreement (MFA) and the Agreement on Textile and Clothing (ATC)", in Robert Feenstra and Shang-Jin Wei (eds) China's Growing Role in World Trade, Chicago: Chicago University Press.

Caballero, R.J. and Engel, E.M.R.A (1999), "Explaining Investment Dynamics in U.S. Manufacturing: a Generalized (S,s) Approach", Econometrica, 67(4), 783-826.

Campbell, J., Lettau, M., Malkiel B. and Xu, Y. (2001), "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk", Journal of Finance, 56(1), 1-43.

Christiano, L.J., Motto, R. and Rostagno, M. (2010), "Financial factors in economic fluctations", Northwestern mimeo.

Christiano, L.J., Eichenbaum, M. and Evans, C.L. (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy", Journal of Political Economy, 113(1), 1-45.

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

Davis, S.J., Haltiwanger, J., Jarmin, R. and Miranda, J. (2006), "Volatility and Dispersion in Business Growth Rates: Publicly Traded Versus Privately Held Firms", NBER WP 12354.

Diebold, F. and Yilmaz, K. (2008), "Macroeconomic volatility and stock market volatility worldwide", in T. Bollerslev, J. Russell and M. Watson (eds.), Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle. Oxford: Oxford University Press.

Dixit, A.K. and Pindyck, R.S. (1994), "Investment Under Uncertainty", Princeton University Press, Princeton.

Eisfeldt, A.L. and Rampini, A. (2008), "Managerial Incentives, Capital Reallocation, and the Business Cycle", Journal of Financial Economics, 87(1), 177-199.

Engle, R.F. and Rangel, J.G. (2008), "The Spline-GARCH Model for Low-Frequency Volatility and Its Global Macroeconomic Causes", Review of Financial Studies, 21(3), 1187-1222.

Fernandez-Villaverde, J., Guerro, P., Rubio-Ramirez, J. and Uribe, M. (2009), "Risk matters: the real effects of volatility shocks", Penn mimeo.

Fernandez-Villaverde, J., Guerrón-Quintana, P. A., Kuester, K. and Rubio-Ramirez, J. (2011), "Fiscal Volatility Shocks and Economic Activity", NBER Working Paper No. 17317.

Foster, L., Haltiwanger, J. and Krizan, C.J. (2000), "Aggregate Productivity Growth: Lessons from Microeconomic Evidence", New Developments in Productivity Analysis, NBER, University of Chicago Press.

Foster, L., Haltiwanger, J. and Krizan, C.J. (2006), "Market Selection, Reallocation and Restructuring in the U.S. Retail: Trade Sector in the 1990s", Review of Economics and Statistics, 88, 748-758. Foster, L., Haltiwanger, J. and Syversion, C. (2008), "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" American Economic Review, 98(1), 394-425.

Gilchrist, S., Sim, J. and Zakrajsek, E (2009), "Uncertainty, credit spreads and investment dynamics", Boston University mimeo.

Hamilton, J.D. and Lin, G. (1996), "Stock Market Volatility and the Business Cycle", Journal of Applied Econometrics, 11(5), 573-593.

Hartman, R. (1972), "The effect of price and cost uncertainty on Investment", Journal of Economic Theory, 5, 258-266.

Hammermesh, D. (1989), "Labor demand and the structure of adjustment costs", American Economic Review, 79, 674-89.

Hassler, J. (1996), "Variations in Risk and Fluctuations in Demand - a Theoretical Model", Journal of Economic Dynamics and Control, 20, 1115-1143.

House, C.L. (2008), "Fixed Costs and Long-Lived Investments", NBER WP 14402.

Hopenhayn, H and Rogerson, R. (1993), "Job turnover and policy evaluation: a general equilibrium analysis", Journal of Political Economy, vol 101, pp 915-938.

Kehrig, Matthias. (2011), "The Cyclicality of Productivity Dispersion", CES WP 11-15.

Jaimovich, N. and Siu, H. (2008), "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility", forthcoming American Economic Review.

Justiniano, A and Primiceri, G. (2008) "Time varying volatility of macroeconomics fluctuations", American Economic Review, 98 (3), 604-641.

Leahy, J. and Whited, T (1996) "The Effects of Uncertainty on Investment: Some Stylized Facts", Journal of Money Credit and Banking, 28, 64-83.

Kehrig, Matthias, (2010), "The cyclicality of productivity dispersion", Northwestern mimeo.

Keynes, J.M. (1936), "The General Theory of Empployment, Interest and Money", Macmillan, London.

Khan, A. and Thomas, J.K. (2008), "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," Econometrica, 76(2), 395-436.

Khan, A. and Thomas, J.K. (2011), "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity," NBER Working Paper 17311.

King, R.G. and Rebelo, S.T. (1999), "Resuscitating Real Business Cycles", in Handbook of Macroeconomics, John B. Taylor and Michael Woodford (eds.), Elsevier.

Kocherlakota, N., (2010), "Modern Macroeconomic Models as Tools for Economic Policy", The Region, May 2010 issue.

Krusell, P and Smith, A.A. (1998), "Income and Wealth Heterogeneity in the Macroeconomy", Journal of Political Economy, 106(5), 867-896. Kydland, F. E. and Prescott E. C. (1982), "Time to Build and Aggregate Fluctuations", Econometrica 50 (6), 1345-1370

Lillien, D.M. (1982), "Sectoral Shifts and Cyclical Unemployment", Journal of Political Economy, 90(4), 777-793.

Mele, A. (2007), "Asymmetric Stock Market Volatility and the Cyclical Behavior of Expected Returns", Journal of Financial Economics, 86, 446-478.

Merz, M. and Yashiv, E. (2007), "Labor and the Market Value of the Firm", American Economic Review, 97(4), 1419-1431.

Narita, F. (2011), "Hidden Actions, Risk-Taking, and Uncertainty Shocks", University of Minnesota mimeo.

Nickell, S.J. (1986), "Dynamic Models of Labor Demand", in O. Ashenfelter and R. Layard (eds.), Handbook of Labor Economics, North-Holland, Amsterdam.

Panousi, V. and Papanikolaou, D. (2011), "Investment, Idiosyncratic Risk, and Ownership", forthcoming, Journal of Finance.

Pindyck, R.S. (1988), "Irreversible Investment, Capacity Choice, and the Value of the Firm", American Economic Review, 78(5), 969-985.

Ramey, V.A. and Shapiro, M.D. (2001) "Displaced Capital: A Study of Aerospace Plant Closings", Journal of Political Economy, 109, 605-622.

Rebelo, S.T. (2005), "Real Business Cycle Models: Past, Present, and Future", Scandinavian Journal of Economics, 107(2), 217-238.

Romer, C. (1990) "The Great Crash and the Onset of the Great Depression.", Quarterly Journal of Economics, 105, 597-624

Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?", Journal of Finance, 44, 1115–1153.

Sim, J., (2006) "Uncertainty, Irreversible Investment and General Equilibrium", Boston University mimeo.

Storesletten, K., Telmer, C.I. and Yaron, A. (2004), "Cyclical Dynamics in Idiosyncratic Labor-Market Risk", Journal of Political Economy, 112(3), 695-717.

Syverson, Chad, 2004a. 'Product Substitutability and Productivity Dispersion', Review of Economics and Statistics 86(2), 534-550.

Thomas, J.K. (2002), "Is Lumpy Investment Relevant for the Business Cycle?", Journal of Political Economy, 110(3), 508-534.

Van Nieuwerburgha, S. and Veldkamp, L. (2006), "Learning asymmetries in real business cycles", Journal of Monetary Economics, 53(4), 753-772.

Veldkamp, L. (2005), "Slow boom, sudden crash", Journal of Economic Theory, 124(2), 230-257.

Veracierto, M.L. (2002), "Plant Level Irreversible Investment and Equilibrium Business Cycles", American Economic Review, 92(1), 181-197.

A Appendix: Census Uncertainty Data

We use data from the Census of Manufacturing (CM) and the Annual Survey of Manufacturing (ASM) from the US Census Bureau to construct an establishment level panel. Using the Compustat-SSEL bridge (CPST-SSEL) we merge the establishment level data with Compustat and CRSP high frequency firm level financial and sales data which allows us to correlate firm and industry level cross-sectional dispersion from census data with stock-returns volatility measures. For industry level deflators, and to calculate production function elasticities, we use industry level data from the NBER-CES productivity database, the Federal Reserve Board (Prices and depreciation), the BLS (Multifactor productivity) and the BEA (Fixed assets tables). We use exchange rates and product level quotas and trade data to construct industry-year demand instruments. In this appendix we describe each of our data sources, the way we construct our samples and the way each variable is constructed. In constructing the TFP variables we closely follow Syverson (2004).

A.1 Data Sources

A.1.1 Establishment level

The establishment level analysis uses the CM and the ASM data. The CM is conducted every 5 years (for years ending 2 and 7) since 1967 (another CM was conducted at 1963). It covers all establishments with one or more paid employees in the manufacturing sector (SIC 20-39 or NAICS 31-33) which amounts to 300,000 to 400,000 establishments per survey. Since the CM is conducted at the establishment level, a firm which operates more than one establishment files a separate report for each establishment. As a unique establishment level id we use the LBDNUM variable which allows us to match establishments over time within the CM and between the CM and the ASM. We use the FIRMID variable to match establishments to the Compustat-SSEL bridge which allows us to match to Compustat and CRSP firm's data using the Compustat CUSIP identifier.

For annual frequency we add the ASM files to the CM files constructing a panel of establishments from 1972 to 2009 (using the LBDNUM identifier). Starting 1973, the ASM is conducted every year in which a CM is not conducted. The ASM covers all establishments which were recorded in the CM above certain size and a sample of the smaller establishments. The ASM includes 50,000 to 75,000 observations per year. Both the CM and the ASM provide detailed data on sales, value added, labor inputs, labor cost, materials' cost, capital expenditures, inventories and more. We give more details on the variables we use in the variables construction subsection below.

A.1.2 Firm level

We use Compustat and CRSP to calculate volatility of sales and returns at the firm level.³⁰ The Compustat-SSEL bridge is used to match census establishment data to Compustat and CRSP firm's data using the Compustat CUSIP identifier. The bridge includes a mapping (m:m) between FIRMID (which can be found in the CM and ASM) and CUSIP8 (which can be found in Compustat and CRSP). The bridge covers the years 1976 to 2005. To extend the bridge to the entire sample of our analysis (1972-2009), we assigned each FIRMID after 2001 to the last observed CUSIP8 and before 1976 to the first observed CUSIP8³¹.

From the CRSP data set we obtain daily and monthly returns at the firm level (RET). From Compustat we obtain firm level quarterly sales (SALEQ) as well as data on equity (SEQQ) and debt (DLTTQ and DLCQ) which is used to construct the leverage ratio (in book values).

A.1.3 Industry level

We use multiple sources of industry level data for variables which do not exist at the establishment or firm level including price indices, cost shares, depreciation rates, market to book ratio of capital, import-export data and industrial production.

The NBER-CES Manufacturing Industry Database is the main source for industry level price indices for total value of shipments (PISHIP), capital expenditures (PIINV), cost of materials (PIMAT) and cost of energy (PIEN).³² It is also the main source for industry level total cost of inputs for labor (PAY), cost of materials (MATCOST) and cost of energy (ENERGY). These total cost variables are used in the construction of industry cost shares. We match the NBER data to the establishment data using 4 digit SIC87 codes for the years 1972-1996 and 6 digit NAICS codes starting 1997. We complete our set of price indices using FRB capital investment deflators (separate deflators for equipment and structures) kindly provided to us by Randy Becker.

The BLS Multifactor productivity data base is used for constructing data on industry level cost of capital and capital depreciation.³³ In particular data from the tables "Capital Detail Data by Measure for Major Sectors" is used. From these tables we obtain data on depreciation rates (table 9a: EQDE, STDE), capital income (table 3a: EQKY STKY), productive capital (table 4a: EQPK, STPK) and index of the ratio of capital input to productive stock (table 6b: EQKC, STKC). All measures are obtained separately for equipment and for structures (there are the EQ and ST prefix respectively). We use these measures to recover the cost of capital in production at the industry level. We match the BLS data to the establishment data using 2 digit SIC87 codes for the years 1972-1996 and 3 digit NAICS

³⁰The access to CRSP and Compustat data sets is through WRDS: https://wrds-web.wharton.upenn.edu/wrds/.

³¹We do this assignment for 2002-2005 since the bridge have many missing matches for these years.

³²See: http://www.nber.org/data/nbprod2005.html for the public version. We thank Wayne Gray for sharing his version of the dataset that is updated to 2009.

³³See: http://www.bls.gov/mfp/mprdload.htm

codes starting 1997.

We use the BEA Fixed assets tables to transform establishment level capital book value to market value. For historical cost we use tables 3.3E and 3.3S for equipment and for structures respectively.³⁴ For current cost we use tables 3.1E and 3.1S.

The industrial production index constructed by the Board of Governors of the Federal Reserve System (FRB) is used to construct annual industry level volatility measures.³⁵ The data is at a monthly frequency and is provided at NAICS three to six digits level. We match the data to establishment level data using the most detailed NAICS value available in the FRB data. Since ASM and CM records do not contain NAICS codes prior to 1997, we obtain for each establishment in our sample a modal NAICS code which will be non-missing in the case that the establishment appears for at least one year after 1996. For establishments who do not appear in our sample after 1996 we use an empirical SIC87-NAICS concordance. This concordance matches to each SIC87 code its modal NAICS code using establishments which appear in years prior to 1997 and after 1997.

We use data from Peter K. Schott's website for exports and imports originally purchased from the U.S. Census Bureau and given in a 4 digit SIC87 codes.³⁶ We use the Cost, Insurance and Freight (CIF) of import by industry from this data set when we construct weighted industry level exchange rate indices which are used as instruments for demand shocks. We match the data to establishment level data using SIC87 codes. Since SIC87 codes are not available for all years, we follow the procedure described above for NAICS to assign SIC87 codes for all establishments.

A.1.4 Additional Data Sets

We use three additional data sets in the construction of the instruments for demand shocks. The IMF IFS website is used for downloading exchange rates between local currencies of 15 countries and the US dollar.³⁷ We focus on G-20 countries as these have large diversified economies so should have exchange rate movements which are exogeneous to shocks to any particular industry. Within the G-20 we exclude Argentina, Brazil, and Russia since these have hyperinflations over this period making the construction of real exchange rates very hard. We obtain price deflators for the 15 countries from the OECD website.³⁸

For the construction of the trade instruments, we use data on the change in quotas on imports from China constructed by Bloom, Draca and Van Reenen (2011), and provided to us by the authors (see the data appendix in their paper for more details on construction of the quotas).

³⁴See: http://www.bea.gov/national/FA2004/SelectTable.asp

³⁵See: http://www.federalreserve.gov/releases/G17/Current/default.htm

³⁶See: http://faculty.som.yale.edu/peterschott/sub_international.htm. These data are an update of Schott (2008) and use the concordances from Pierce and Schott (2009) and Bartelsman et al (2000).

³⁷See: http://www.imfstatistics.org/IMF/imfbrowser.aspx?branch=ROOT

³⁸See: http://stats.oecd.org/index.aspx?querytype=view&queryname=221

A.2 Sample Selection

We have three main establishments samples which are used in our analysis of the manufacturing sector. The largest sample includes all establishments which appear in the CM or ASM for at least two consecutive years (implicitly implying that we must have at least one year from the ASM, therefore ASM sampling applies). In addition we exclude establishment which are not used in for census tabulation (TABBED=N), establishment with missing or non-positive data on total value of shipments (TVS) and establishments with missing values for LBDNUM, value added (VA), labor inputs or investment. This sample consists of 211,939 establishments and 1,340,793 establishment-year observations.

Our baseline sample for the establishment level only keeps establishments which appears at least 25 years between 1972 and 2009. This sample consists of 15,673 establishments and 446,051 establishment-year observations.

The third sample we use is based on the baseline sample limited to establishments for which their firms have CRSP and Compustat records, with non-missing values for stock returns, sales, equity and debt. The sample includes 10,498 establishments with 172,074 establishment-year observations.

When calculating annual dispersion measures using CRSP and Compustat (see table 1), we use the same sampling criteria as in the baseline ASM-CM sample, keeping only firms which appear at least 25 years.

A.3 Variables Construction

A.3.1 Value Added

We use the census value added measure which is defined for establishment j at year t as

$$v_{j,t} = Q_{j,t} - M_{j,t} - E_{j,t}$$

where $Q_{j,t}$ is nominal output, $M_{j,t}$ is cost of materials and $E_{j,t}$ is cost of energy and fuels. Nominal output is calculated as the sum of total value of shipments and the change in inventory from previous year (both finished inventory and work in progress inventory).

In most of the analysis we use real value added. In this case, we deflate value added by the 4 digit industry price of shipment (PISHIP) given in the NBER-CES data set.

A.3.2 Labor Input

The CM and ASM report for each establishment the total employment (TE), the number of hours worked by production workers (PH), the total salaries for the establishment (SW) and the total salaries for production workers (WW). The surveys do not report the total hours for non-production workers. In addition, one might suspect that the effective unit of labor input is not the same for production and non-production workers. We calculate the following measure of labor inputs

$$n_{j,t} = \frac{SW_{j,t}}{WW_{j,t}}PH_{j,t}$$

A.3.3 Capital Input

There are two issues to consider when constructing the capital measure. First, capital expenditures rather than capital stock are reported in most survey years, and when capital stock is reported it is sensitive to differences in accounting methods over the years. Second, the reported capital in the surveys is book value. We deal with the latter by first converting book to market value of capital stocks using BEA fixed asset tables which include both current and historical cost of equipment and structures stocks by industry-year. We address the first issue using the perpetual inventory method, calculating plant level series of capital stocks using the plant initial level of capital stock, the plant level investment data and industry level depreciation rates and. To apply the perpetual inventory method we first deflate the initial capital stock (in market value) as well as the investment series using FRB capital investment deflators. We than apply the formula³⁹

$$K_t = (1 - \delta_t) K_{t-1} + I_t$$

This procedure is done separately for structures and for equipment. However, starting 1997, the CM does not report separately capital stocks for equipment and structures. For plants which existed in 1992, we can use the investment data to back out capital stocks for equipment and structures separately after 1992. For plants born after 1992, we assign the share of capital stock to equipment and structures to match the share of investment in equipment and structures.

A.3.4 TFP and TFP Shocks

For establishment j in industry i at year t we define value added based total factor productivity (TFP) $\hat{z}_{j,i,t}$ as

$$\log(\hat{z}_{j,i,t}) = \log(v_{j,i,t}) - \alpha_{i,t}^{S} \log(k_{j,i,t}^{S}) - \alpha_{i,t}^{E} \log(k_{j,i,t}^{E}) - \alpha_{i,t}^{N} \log(n_{j,i,t})$$

where $v_{j,i,t}$ denotes value added (output less materials and energy inputs), $k_{j,i,t}^S$ represents the capital stock of structures, $k_{j,i,t}^E$ represents the capital stock of equipment and $n_{j,i,t}$ the total hours worked as described above.

 $\alpha_{i,t}^S$, $\alpha_{i,t}^E$ and $\alpha_{i,t}^N$ are the cost shares for structures, equipment and labor. These cost shares are recovered at 4-digit industry level by year, as is standard in the establishment productivity estimation literature (see, for example, the survey in Foster, Haltiwanger and Krizan, 2000). We generate the cost shares such that they sum to one. Define $c_{i,t}^x$ as total

³⁹The stock is reported for the end of period, therefore we use last period's stock with this period depreciation and investment.

cost of input x for industry i at year t. Then for input x

$$\alpha_{i,t}^{x} = \frac{c_{i,t}^{x}}{\sum_{x \in X} c_{i,t}^{x}}, X = \{S, E, N\}.$$

We use industry level data to back out $c_{i,t}^x$. The total cost of labor inputs c_i^N is taken from the NBER-CES Manufacturing Industry Database (PAY). The cost of capital (for equipment and structures) is set to be capital income at the industry level. The BLS productivity dataset includes data on capital income at the two digit industry level. To obtain capital income at four digit industry level we apply the ratio of capital income to capital input calculated using BLS data to the four digit NBER-CES capital data.

Given the cost shares, we can recover $\log(\hat{z}_{j,i,t})$. We then define TFP shocks $(e_{j,t})$ as the residual from the following first-order autoregressive equation for establishment level log TFP

$$\log\left(\widehat{z}_{j,i,t}\right) = \rho \log\left(\widehat{z}_{j,i,t-1}\right) + \mu_i + \lambda_t + e_{j,i,t} \tag{19}$$

where μ_i are plant fixed effects and λ_t are year dummies.

A.3.5 Micro Uncertainty Dispersion-Based Measures

Our main micro uncertainty measures are based on establishment level TFP shocks $(e_{j,t})$ and on establishment level growth in employment and in sales. For variable x we define establishment's *i* growth rates for year t as $\Delta x_{i,t} = (x_{i,t+1} - x_{i,t})/(0.5 \times x_{i,t+1} + 0.5 \times x_{i,t})$.

Aggregate level (Table 1): To measure uncertainty at the aggregate level, we use the interquartile range (IQR) and the standard deviation of both TFP shocks and sales and employment growth by year. We consider additional measures for TFP shocks that allow for more flexibility in the AR regression (19) used to back out the shocks. In particular we report the dispersion of TFP shocks which were calculated by running (19) at the 3 digit industry level (industry by industry), effectively allowing for ρ and for λ_t to vary by industry.

We use three additional aggregate uncertainty measures which are not based on census data. We use CRSP to calculate the firms' cross-sectional dispersion of monthly stock returns at a monthly frequency, and Compustat to calculate the cross-sectional dispersion of sales' growth at a quarterly frequency, where sales growth is defined as $(x_{i,t+4} - x_{i,4})/(0.5 \times x_{i,t+4} + 0.5 \times x_{i,4})$. We use the industrial production index constructed by the FRB to calculate the cross-sectional dispersion of industry production growth $(x_{i,t+12} - x_{i,12})/(0.5 \times x_{i,t+12} + 0.5 \times x_{i,12})$ at the monthly level.

Firm level (Table 4): To measure uncertainty at the firm level, we use the weighted mean of the absolute value of TFP shocks and sales growth, where we use establishments' total value of shipments as weights. As an example, the uncertainty measure for firm f at

year t using TFP shocks is calculated as

$$\frac{\sum_{j \in f} TVS_{j,t} * |e_{j,t}|}{\sum_{j \in f} TVS_{j,t}}$$

(and similarly for growth measures).

Industry level: At the industry level we use both IQR (Table 2 and Table 3) and weighted mean of absolute values (Table 4) as uncertainty measures.

A.3.6 Micro Volatility Measures

Using CRSP, Compustat and FRB data, we construct firm and industry level annual volatility measures which are used in Table 4.

Firm level: At the firm level we construct volatility measures using firms' stock returns. We use standard deviation of daily and monthly returns over a year to generate the stock volatility of a firm. For the monthly returns we limit our samples to firms with data on at least 6 months of returns in a given calendar year. For monthly returns we winsorize records with daily returns which are higher or lower than 25%. As an alternative measure we follow Leahy and Whited (1996) and Bloom, Bond and Van Reenen (2007) in implementing a leverage adjusted volatility measure which eliminates the effect of gearing on the variability of stock returns. To generate this measure we multiply the standard deviation of returns for firm f at year t by the ratio of equity to (equity + debt), with equity measured using the book value of shares (SSEQ) and debt measured using the book value of debt (DLTTQ + DLCQ). To match the timing of the TFP shock in the regressions (calculated between t and t+1, see eq. (19)), we average over the standard deviation of returns at year t and the standard deviation at year t+1.

For volatility of sales we use the standard deviation over a year of firm's annual growth calculated at a quarterly rate $(x_{i,t+4} - x_{i,4})/(0.5 \times x_{i,t+4} + 0.5 \times x_{i,4})$.

Industry level: For industry level measures of volatility we use the standard deviation over a year of industry's annual growth calculated at a monthly rate $(x_{i,t+12} - x_{i,12})/(0.5 \times x_{i,t+12} + 0.5 \times x_{i,12})$ using the industrial production index constructed by the FRB.

A.3.7 Industry Characteristics

In Table 2 we use measures for industry business conditions and for industry characteristics. To proxy for industry business conditions we use either mean or median plants' real sales growth rates within industry year. Industry characteristic are constant over time and are either level or dispersion measures. For levels we use medians, implying that a typical measure would look like

$$Median_{j\in i}\left(\frac{1}{T}\sum_{t=1}^{T}x_{jt}\right)$$

where x_{jt} is some characteristic of plant j at year t (e.g. plant total employment). The industry level measure is calculated as the median over all plants in industry i of the within plant mean over time of x_{jt} . The dispersion measures are similar but use IQR instead of medians:

$$IQR_{j\in i}\left(\frac{1}{T}\sum_{t=1}^{T}x_{jt}\right)$$

One exception is the measure of industry geographic dispersion, which is calculated as the Ellison-Glaeser dispersion index at the county level.

A.3.8 Demand Instruments

In Table 3 we use two instruments for first moment shocks, both are at the industry level. The first instrument is based on the abolishment of the China textile quotas in 2005, which translated to a negative demand shock to the local textile industry. The second instrument is in the spirit of the instruments in Bertrand (2004). It is constructed as a weighted industry level exchange rate index, where the weight of a particular country's exchange rate is given by the exposure of the industry to the particular currency. An increase in the industry exchange rate is a negative demand shock to this industry in the US since it reduces the demand for exports from this industry and increases import of goods for the particular industry.

Textile Quotas Instrument: The relaxation of quotas for China started when it joined the WTO in 2001, and peaked in 2005 when the quotas were completely removed. The removal of the quotas generated an increase in the imports of Chinese goods in the categories for which the quotas were removed. We use the 2005 quota variable constructed by Bloom, Draca and Van Reenen (2011). For each four-digit industry this variable stores the trade weighted proportion of product categories that were covered by a quota in 2005 by 4 digit industry categories.⁴⁰ The instrument is then constructed as the interaction between the quota level and a dummy which takes the value of 1 for all years starting 2005. We limit the analysis to industries which are similar to the treated group, thus focusing on the textile and related industry (SIC codes 22, 23, 28 and 29, which were the 2-digit industries including sub-industries impacted by the quotas). We restrict the analysis to a 7 years window around the change (2002-2008).

Exchange rate instrument: We use the OECD and IMF data to construct 15 series of real exchange rates for Australia, Canada, China, France, Germany, India, Indonesia, Italy, Japan, Mexico, Russia, Saudi Arabia, South Africa, Republic of Korea, Turkey and the UK. Three of these countries (Germany, France and Italy) are part of the Euro zone, therefore changed their currency to Euro on January 1st 1999. To keep the exchange rate series smooth for these countries we convert the Euro series to the currency used in the country

⁴⁰For detailed description of the quotas system and its effect on China's export see Bloom, Draca and Van Reenen (2011) and Brambilla et al. (2010).

prior to 1999 using the December 31st 1998 exchange rate.⁴¹ We then use the data from Schott's website on imports and exports to generate a weighted exchange rate by industry. The weights are based on the Cost, Insurance and Freight (CIF) of import variable. These are collapsed to the sum of imports at the country-industry level over 1972 to 2005 (so weights are fixed over time) and zeros are assigned for missing values. For industry *i* at time *t*, the instrument is constructed as

$$exc_inst_{i,t} = \sum_{c} w_{i,c} * \log(exc_{c,t})$$

where c is index for country and the weights $w_{i,c}$ are defined as

$$w_{i,c} = \frac{CIF_{i,c}}{\sum_{c} CIF_{i,c}}$$

⁴¹This is available in Wikipedia: http://en.wikipedia.org/wiki/Euro

Table 1: Uncertainty is Higher During Recessions

Dependent Variable:	(1) S.D. of log(TFP) shock	(2) Skewness of log(TFP) shock	(3) Kurtosis of log(TFP) shock	(4) IQR of log(TFP) shock	(5) IQR of output growth	(6) IQR of sales growth	(7) IQR of stock returns	(8) IQR of industrial prod. growth
Sample:	Establishments (manufacturing)	Establishments (manufacturing)	Establishments (manufacturing)	Establishments (manufacturing)	Establishments (manufacturing)	Public firms (all sectors)	Public firms (all sectors)	Industries (manufacturing)
Recession	0.063*** (0.010)	-0.244 (0.179)	-1.432 (2.088)	0.060*** (0.021)	0.076*** (0.021)	0.032*** (0.009)	0.025*** (0.004)	0.044*** (0.006)
Mean of Dep. Variable:	0.499	-1.527	20.514	0.39	0.195	0.186	0.104	0.101
Corr. with GDP growth	-0.440***	0.131	0.038	-0.446***	-0.560***	-0.185**	-0.283***	-0.340***
Frequency	Annual	Annual	Annual	Annual	Annual	Quarterly	Monthly	Monthly
Years	1972-2009	1972-2009	1972-2009	1972-2009	1972-2009	1962:1-2010:3	1960:1-2010:9	1972:1-2010:11
Observations	37	37	37	37	37	191	609	455
Underlying sample	446,051	446,051	446,051	446,051	446,051	320,306	931,143	70,487

Notes: Each column reports a time-series OLS regression point estimate (and standard error below in brackets) of a measure of uncertainty on a recession indicator. The recession indicator is the share of quarters in that year in a recession in columns (1) to (5), whether that quarter was in a recession in column (6), and whether the month was in recession in columns (7) and (8). Recessions are defined using the NBER data. In the bottom panel we report the mean of the dependent variable and its correlation with real GDP growth. In columns (1) to (5) the sample is the population of manufacturing establishments with 25 years or more of observations in the ASM or CM survey between 1972 and 2009, which contains data on 15,673 establishments across 38 years of data (one more year than the 37 years of regression data since we need lagged TFP to generate a TFP shock measure). We include plants with 25+ years to reduce concerns over changing samples. In column (1) the dependent variable is the cross-sectional standard deviation (S.D.) of the establishment level 'shock' to Total Factor Productivity (TFP). This 'shock' is calculated as the residual from the regression of log(TFP) at year t+1 on its lagged value (year t), a full set of year dummies and establishment fixed effects. In column (2) we use the cross-sectional skewness of the TFP 'shock', in column (3) the cross-sectional kurtosis and in column (6) the dependent variable is the inter-quartile range of plants' sales growth. In column (6) the dependent variable is the inter-quartile range of firms' sales growth by quarter for all public firms with 25 years or more in Compusta theween 1962 and 2010 (with 25+ years or more in CRSP between 1960 and 2010. Finally, in column (8) the dependent variable is the inter-quartile range industrial production growth by month for manufacturing industries from the Federal Reserve Board's monthly industrial production database. All regressions include a time trend and census year dummies (for census year and for 3 lags). Newey-Wes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variable:	IQR of establishment TFP shocks within each industry-year cell								
Specification:	Baseline	Median industry output growth	IQR of industry output growth	Median industry establishment size	IQR of industry establishment size	Median industry capital/labor ratio	IQR of industry capital/labor ratio	IQR of industry TFP spread	Industry geographic spread
Industry Output Growth	-0.132*** (0.021)	-0.142*** (0.021)	-0.176*** (0.047)	-0.119*** (0.024)	-0.116*** (0.022)	-0.111*** (0.034)	-0.111*** (0.030)	-0.191*** (0.041)	-0.133*** (0.028)
Interaction of industry output growth with the variable in specification row		0.822 (0.630)	0.882 (0.996)	-0.032 (0.038)	-0.033 (0.026)	-0.197 (0.292)	-0.265 (0.330)	0.123 (0.084)	0.007 (0.122)
Years	1972-2009	1972-2009	1972-2009	1972-2009	1972-2009	1972-2009	1972-2009	1972-2009	1972-2009
Observations	16,451	16,451	16,451	16,451	16,451	16,451	16,451	16,451	16,451
Underlying sample	446,051	446,051	446,051	446,051	446,051	446,051	446,051	446,051	446,051

Notes: Each column reports the results from an industry by year OLS panel regression, including a full set of industry and year fixed-effects. The dependent variable in every column is the interquartile range (IQR) of establishment level TFP 'shocks' within each SIC 4-digit industry-year cell. The regression sample is the 16,451 industry-year cells of the population of manufacturing establishments with 25 years or more of observations in the ASM or CM survey between 1972 and 2009 (which contains 446,051 underlying establishment years of data). These industry-year cells are weighted in the regression by the number of establishment observations within that cell, with the mean and median number of establishments per industry-year cell 27.1 and 17 respectively. The TFP 'shock' is calculated as the residual from the regression of log(TFP) at year t+1 on its lagged value (year t), a full set of year dummies and establishment fixed effects. In column (1) the explanatory variable is the median of the establishment-level output growth in that industry year. In columns (2) to (9) a second variable is also included which is an interaction of that explanatory variable with an industry level characteristic. In columns (2) and (3) this is the median and IQR of industry level output growth, in columns (4) and (5) this is the median and IQR of industry level stablishment size denotes in employees, in columns (6) and (7) this is the median and IQR of industry level capital/labor ratios, in column (8) this is the IQR of industry level TFP levels (note the mean is zero by construction), while finally in column in (9) this interaction is the dispersion of industry level concentration measured using the Ellison-Glaeser dispersion index. Standard errors clustered by industry are reported in brackets below every point estimate. *** denotes 1% significance, ** 5% significance and * 10% significance.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	IQR of establishment TFP shocks within each industry-year cell					
Sample	Manufacturing, 25+ year establishments	Textile and apparel establishments	Textile and apparel establishments	Textile and apparel establishments	Manufacturing, 25+ year establishments.	Manufacturing, 25+ year establishments
Specification:	OLS	OLS	IV	Reduced form	IV	Reduced form
Median real output growth rates	-0.132***	-0.387**	0.428		0.230	
	(0.021)	(0.188)	(0.776)		(0.411)	
2005 Quotas				-0.036 (0.070)		
Industry ex. rate						-0.0209
						(0.035)
Industry ex. rate at t-1						-0.0300
						(0.035)
First Stage:						
2005 Quotas			-0.084***			
			(0.027)			
Industry exchange rate					-0.0840***	
					(0.021)	
Industry exchange rate at t-1					0.064***	
					(0.019)	
F-test			9.91		7.79	
Years	1972-2009	2002-2008	2002-2008	2002-2008	1973-2008	1973-2009
Observations	16,451	474	474	474	16,009	16,009
Mean obs per industry year	27.1	22.0	22.0	22.0	27.3	27.3
Underlying sample size	446,051	10,703	10,703	10,703	436,261	436,261

Table 3: First Moment Shocks do not seem to be Driving Industry-Level Counter-Cyclical Uncertainty

Notes: Each column reports the results from an industry by year panel regression, including a full set of industry and year fixed-effects. The dependent variable in every column is the interquartile range (IQR) of establishment level TFP 'shocks' within each SIC 4-digit industry-year cell. The regression sample in column (1) is the 16,451 industry-year cells of the population of manufacturing establishments with 25 years or more of observations in the ASM or CM survey between 1972 and 2009 (which contains 446,051 underlying establishment years of data), in columns (2) to (4) is the apparel and textiles subset of this dataset, and in columns (5) and (6) is this dataset less the data for 1972 (since Breton-Woods created fixed exchange rates until 1971). These industry-year cells are weighted in the regression by the number of establishment observations within that cell. The TFP 'shock' is calculated as the residual from the regression of log(TFP) at year t+1 on its lagged value (year t), a full set of year dummies and establishment fixed effects. Columns (1), (2), (4) and (6) are estimated by OLS while columns (3) and (5) are estimated by instrumental variables (IV) with the first-stage results shown in the bottom panel below. The instrument in column (3) is the share of the industry trade covered by quotas on Chinese exports before China joined with WTO in 2005 (details in the Appendix), while the instruments in column 5) are the industry-level exchange rates (details in the Appendix). Standard errors clustered by industry are reported in brackets below every point estimate. *** denotes 1% significance, ** 5% significance and * 10% significance.

Dependent variable	(1) Mean of e	(2) stablishment absolu	(3) te (TFP shocks) wi	(4) ithin firm year	(5) Mean of establishment absolute (TFP shocks) within industry year
Sample	Establishme	nts (in manufacturir	ng) with a parent fin	rm in Compustat	Manufacturing industries
Regression panel dimension	Firm by Year			Industry by Year	
S.D. of parent monthly stock returns within year	0.275*** (0.083)				
S.D. of parent daily stock returns within year		0.317*** (0.091)			
S.D. of parent monthly stock returns within year, leverage adjusted			0.330*** (0.109)		
S.D. of parent quarterly sales growth within year				0.134***	
S.D of monthly industrial production within year				(0.027)	0.330*** (0.060)
Fixed effects and clustering	firm	firm	firm	firm	industry
Firms/Industries	1,838	1,838	1,838	1,838	466
Observations	25,302	25,302	25,302	25,302	16,406
Underlying observations	172,074	172,074	172,074	172,074	446,051

Table 4: Cross-sectional establishment uncertainty measures are correlated with firm and industry time-series uncertainty measures

Notes: The dependent variable is the mean of the absolute TFP shock at the firm-year level (columns (1) to (4)) and industry-year level (column (5)). This TFP 'shock is calculated as the residual from the regression of log(TFP) at year t+1 on its lagged value (year t), a full set of year dummies and establishment fixed effects. The regression sample in columns (1) to (4) are the 25,302 firm-year cells of the population of manufacturing establishments with 25 years or more of observations in the ASM or CM survey between 1972 and 2009 which are owned by Compustat (publicly listed) firms. This covers 172,074 underlying establishment years of data. The regression sample in column (5) is the 16,406 industry-year cells of the population of manufacturing establishments with 25 years or more of observations in the ASM or CM survey between 1972 and 2009. The explanatory variables in columns (1) to (3) are the annual standard deviation of the parent firm's stock returns, which is calculated using the 12 monthly values in columns (1) and (3) and the 260 daily values in column (3). The monthly stock returns in column (2) are normalized by the (equity/debt+equity) ratio to control for leverage effects. In column (4) the explanatory variable is the standard deviation of the parent firm's quarterly sales growth. Finally, in column (5) the explanatory variable is the standard deviation of the industries monthly industrial production data from the Federal Reserve Board. All columns have a full set of year fixed effects with columns (1) to (4) also having firm fixed-effects while column (5) has industry fixed effects. Standard errors clustered by firm/industry are reported in brackets below every point estimate. *** denotes 1% significance, ** 5% significance and * 10% significance.

Figure 1: TFP 'shocks' are more dispersed in recessions



Notes: Constructed from the Census of Manufacturers and the Annual Survey of Manufacturing establishments establishments with 25+ years to address sample selection. Grey shaded columns are share of quarters in recession within a year.

Figure 12: Unweighted TFP



Figure 13: Effect of rise in uncertainty on output



Figure 14: Labor, investment, weighted TFP, consumption



Figure 15: Output – alternative AC specifications



Figure 16: Output – only micro / macro uncertainty



Figure 17: Output – policy effectiveness



Figure 18: Output – differential policy effect



Figure 19: Policy - thresholds

