

Really Uncertain Business Cycles

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten and Stephen J. Terry*

October 2014

Abstract

We propose uncertainty shocks as a new shock that drives business cycles. First, we demonstrate that microeconomic uncertainty is robustly countercyclical, rising sharply during recessions, particularly during the Great Recession of 2007-2009. Second, we quantify the impact of time-varying uncertainty on the economy in a dynamic stochastic general equilibrium model with heterogeneous firms. We find that reasonably calibrated uncertainty shocks can explain drops and rebounds in GDP of around 3%. Moreover, we show that increased uncertainty makes first-moment policies, like interest-rate and tax cuts, temporarily less effective because firms become more cautious in responding to price changes.

Keywords: uncertainty, adjustment costs, and business cycles.

JEL Classification: D92, E22, D8, C23.

Disclaimer: 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 thank our formal discussants Sami Alpanda, Eduardo Engel, Frank Smets, and Eric Swanson; Angela Andrus at the RDC; as well as numerous seminar audiences for comments.

*Bloom at Stanford, NBER and CEPR, Floetotto at McKinsey, Jaimovich at Duke University and NBER, Saporta-Eksten at Tel Aviv University and University College London, and Terry at Stanford 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. For example, the Federal Open Market Committee minutes repeatedly emphasize uncertainty as a key factor driving the 2001 and 2007-2009 recessions, while Stock and Watson (2012) state that “*the main contributions to the decline in output and employment during the [2007-2009] recession are estimated to come from financial and uncertainty shocks.*”

This paper seeks to evaluate these claims in two parts. In the first part, we develop new empirical measures of uncertainty using detailed Census microdata from 1972 to 2011, and highlight three main results. First, the dispersion of plant-level shocks to total factor productivity (TFP) is strongly countercyclical, rising steeply in recessions. For example, Figure 1 shows the dispersion of TFP shocks for a balanced panel of plants for the last two full years before the recent recession (2005 to 2006) and two years during the recession (2008 to 2009). This shows that plant-level TFP shocks increased in variance by 76% during the recession. Figure 2 shows that the dispersion of output growth increased even more, rising by a striking 152% during the recession. Surprisingly, higher microdata moments like the coefficient of skewness and kurtosis are not significantly countercyclical. So recessions appear to be characterized by a negative first-moment and a positive second-moment shock to the establishment-level driving processes.¹

Second, uncertainty is also strongly countercyclical at the industry-level. That is, within SIC 4-digit industries the yearly growth rate 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, we show that for plants owned by publicly traded Compustat parent firms, large plant-level TFP shocks are highly correlated with more volatile daily parent stock returns. Hence, daily stock returns volatility, a popular high-frequency financial measure of uncertainty, is tightly linked to yearly plant TFP shocks, our low-frequency real measure of uncertainty.

Given the robust evidence that uncertainty appears to rise sharply in recessions, in the second part of the paper we build a dynamic stochastic general equilibrium (DSGE) 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,

¹To be precise we find that the *coefficients* (rather than levels) of skewness and kurtosis are acyclical, where the *coefficient* of skewness= $M3/M2^{3/2}$ and kurtosis= $M4/M2^2$, and Mx is the level of the x^{th} centered moment. Data available at <http://www.stanford.edu/~nbloom/RUBC.zip>

so the model includes shocks to both the level of technology (the first moment) and its variance (the second moment) at both the microeconomic and macroeconomic levels. Second, there are heterogeneous firms that are subject to idiosyncratic shocks. Third, the model contains 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.

The model is numerically solved and estimated on macro and plant level data using simulated method of moments. Combining macro and micro data helps to overcome the identification problem associated with the limited sample size of macro data. Our parameter estimates suggest that micro and macro uncertainty increase between three to four fold during recessions.

Simulations of the model then allow us to study the response of our model economy to an uncertainty shock. Increased uncertainty makes it optimal for firms to wait, leading to significant falls in hiring, investment and output. In addition, we show that increased uncertainty also reduces productivity growth because it reduces the degree of reallocation in the economy. Higher uncertainty leads productive plants to pause expanding and unproductive plants to pause contracting, which in the model as in the U.S. economy drives much of aggregate productivity growth.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 initially dampens the effect of an expansionary policy. The key to this initial *policy ineffectiveness* is that a rise in uncertainty makes firms very cautious in responding to any stimulus. Once the uncertainty shock has begun to pass, the increased dispersion in actual TFP returns the responsiveness back to its long-run average level.

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 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

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.

³See, for example, the review of this literature in King and Rebelo (1999) and Rebelo (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.

during recessions. In our model, however, the drop in productivity is not 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 Finally, the paper also builds upon a recent literature that studies the role of microeconomic 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. The model is calibrated and simulated in Section 4, and in Section 5 we study the role of uncertainty shocks in driving the business cycle. Section 6 decomposes the effect of an uncertainty shock on the economy. Section 7 studies the impact of policy shocks in the presence of time-varying uncertainty. Section 8 concludes.

2 Measuring Uncertainty over the Business Cycle

Before presenting our empirical results, it is useful to briefly discuss what we mean by time-varying uncertainty in the context of our model.

We assume that a firm, indexed by j , produces output in period t according to the following production function

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

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

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)'s partial equilibrium model with stochastic volatility, Fernandez-Villaverde, Guerron-Quintana, Rubio-Ramirez, and Uribe (2011)'s paper on uncertainty and real exchange rates, Kehrig (2011)'s paper on countercyclical productivity dispersion, Christiano, Motto, and Rostagno (2014)'s, Arelano, Bai, and Kehoe (2012)'s and Gilchrist, Sim, and Zakrajsek (2011)'s papers on uncertainty shocks in models with financial constraints, Basu and Bundick (2011)'s paper on uncertainty shocks in a new-Keynesian model, Fernandez-Villaverde, Guerron-Quintana, Kuester, and Rubio-Ramirez (2014)'s paper on fiscal policy uncertainty, Knotek and Khan (2011)'s paper on durables consumption and uncertainty, and Bachmann and Bayer (2012)'s paper on microeconomic level uncertainty with capital adjustment costs.

⁷See for example, Thomas (2002), Veracierto (2002), Kahn and Thomas (2008, 2013), Bachmann, Caballero, and Engel (2013), and House (2014).

that demand shocks will also be incorporated into the process for A_t and $z_{j,t}$.⁸

We assume that the aggregate and idiosyncratic components of business conditions follow autoregressive processes:

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

$$\log(z_{j,t}) = \rho^Z \log(z_{j,t-1}) + \sigma_{t-1}^Z \epsilon_{j,t}. \quad (3)$$

We allow the variance of innovations, σ_t^A and σ_t^Z , to move over time according to two-state Markov chains, generating 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.

These two shocks are driven by different statistics. Volatility in $z_{j,t}$ implies that cross-sectional dispersion-based 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 macroeconomic uncertainty measures, details of the construction of which are contained in Appendix A.

2.1 Microeconomic Uncertainty over the Business Cycle

In this section we present a set of results showing that shocks at the *establishment*-level, *firm*-level 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.

Our first set of measures come from the Census panel of manufacturing establishments. In summary, with extensive details in Appendix A, this dataset contains detailed output and inputs data on over 50,000 establishments from 1972 to 2011. We focus on the subset of 15,673 establishments with 25+ years of data to ensure that compositional changes do not bias our results, generating a sample of almost half a million establishment-year

⁸See also Hopenhayn and Rogerson (1993), in which they discuss how productivity shocks at the microeconomic level are isomorphic to consumer taste shocks shifting the demand curve.

observations.⁹

To measure uncertainty we first calculate establishment-level TFP ($\hat{z}_{j,t}$) using the standard approach from, 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:

$$\log(\hat{z}_{j,t}) = \rho \log(\hat{z}_{j,t-1}) + \mu_j + \lambda_t + e_{j,t}, \quad (4)$$

where μ_j 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 it does highlight the difficulty in empirically distinguishing productivity shocks from demand shocks.

Finally, we then define microeconomic uncertainty, $\sigma_{t-1}^{\hat{z}}$, as the cross-sectional dispersion of $e_{j,t}$ calculated on a yearly basis. This is shown in Figure 3 as the interquartile range (IQR) of this TFP shock within each year, displaying a clearly countercyclical behavior. This is particularly striking for the recent Great Recession of 2007-2009, which displays the highest value of TFP dispersion since the series begins in 1972.

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 0.064 which is highly significant (a t-statistic of 6.4). In the bottom panel we report that this S.D. of establishment TFP shocks also has a highly significant correlation with GDP growth of -0.45.¹¹ In columns (2) and (3) we examine the coefficient of skewness and kurtosis of TFP shocks over the cycle and interestingly find no significant correlations.¹² This suggests that recessions can be characterized at the

⁹The sampling issues arise both from the cyclical nature of exit and from the sample stratification rules for the Census, which rotates out smaller establishments at 5-yearly intervals. By restricting the sample to 25+ year lived establishments we eliminate cyclical frequency in sampling fluctuations.

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.

¹¹Several other recent papers have also reported similar findings of countercyclical increases in the variance of productivity shocks. Bachmann and Bayer (2014) 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 (2011) like our paper uses the U.S. Census data, but takes a different approach to sampling and estimating productivity, and finds a significant increase in the spread of productivity *levels* in recessions. Finally, Jurado, Ludvigson and Ng (2014) show increased dispersion for a wide range of US firm measures during downturns.

¹²This lack of significant correlation was robust in a number of experiments we ran. For example, if we drop the time trend and Census survey year controls the result in column (1) on the standard deviation remains

microeconomic level as a negative first-moment shock plus a positive second moment shock, with no shocks to higher moments. In column (4) we use an outlier-robust measure of cross-sectional dispersion, which is the IQR range of TFP shocks, and again find this rises significantly in recessions.¹³ In column (5) as another robustness test we use output growth, 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.¹⁴ For example, Figure A1 plots the correlation of plant TFP rankings between consecutive years. This shows that during recessions these rankings churn much more, as increased microeconomic variance leads plants to change their position within their industry-level TFP rankings more rapidly.

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 spanning all sectors of the economy, and providing 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.

One important caveat when using the variance of productivity ‘shocks’ to measure uncertainty is 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 (1982) in being a forecast error from an AR(1) equation rather than necessarily a shock to economic agents. We address this concern in two ways. First, in column (7) we examine the cross-sectional spread of stock returns, which reflects the volatility of news about firm performance, and again find this is countercyclical, echoing the prior results in Campbell et al. (2001).¹⁵ In fact we find that establishment-level shocks to TFP are significantly

highly significant at 0.062 (0.020), while the results in columns (2) and (3) on skewness and kurtosis remain insignificant at -0.250 (0.209) and -0.771 (2.782). We also experimented with changing the establishment selection rules (keeping those with 2+ or 38+ years rather than 25+ years) and again found the results robust, as shown in Appendix Table A1. Interestingly, Guvenen, Ozkan and Song (2014) find an increase in left-skewness for personal income growth during recessions, which may be absent in plant data because large negative draws lead plants to exit.

¹³Kehrig (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 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).

¹⁴For example, IQR of employment growth rates has a point estimate (standard error) of 0.051 (0.012), the IQR of TFP shocks measured using an industry-by-industry forecasting equation version of (4) has a point estimate (standard error) of 0.064 (0.020), using 2+ year samples for the S.D. of TFP shocks we find a point estimate (standard error) of 0.046 (0.014), using a balanced panel of 38+ year establishments we find a point estimate (standard error) of 0.075 (0.015), and using the IQR of TFP shocks measured after removing state-year means, and then applying (4) has a point estimate (standard error) of 0.061 (0.020). Finally, using the IQR of TFP shocks measured after removing firm-year means, and then applying (4) has a point estimate (standard error) of 0.028 (0.011).

¹⁵To remove the forecastable component of stock returns we also repeat column (7) first removing the quarter

correlated to their parent’s stock returns, so that, at least part of, these establishment TFP shocks are new information to the market.¹⁶

Second, we extend the TFP forecast regressions (4) to include many observables that are likely to be informative about future TFP changes. Adding these in the regression accounts for at least some of the superior information that the establishment might have over the econometrician, helping us in backing out true shocks to TFP (from the perspective of the establishments). Figure 4 reports the IQR of the TFP shocks for the baseline forecast regression, as well as for three other dispersion measures, where we sequentially add more variables to the forecasting regressions that are used to recover TFP shocks. First we add two extra lags in levels and polynomials of TFP, next we also include lags and polynomials of investment, and finally polynomials and lagged in multiple inputs including employment, energy and materials expenditure. As is clear from the figure, even when including forward looking establishment choices for investment and employment, the overall cyclical patterns of uncertainty are almost unchanged.

Finally, in column (8) we examine another measure of uncertainty, which is the cross-sectional spread of industry-level output growth rates, finding again that this is strongly countercyclical.

Hence, in summary *plant*-level, *firm*-level, and *industry*-level measures of volatility and uncertainty all appear to be strongly countercyclical, suggesting that microeconomic uncertainty rises in recessions.¹⁷

2.2 Industry Business Cycles and Uncertainty

In Table 2 we report another set of results which disaggregate down to the industry level, finding a very similar result that uncertainty is significantly higher during periods of slower growth. 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 the Census dataset means that 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.

by firm mean of firm returns. This controls for any quarterly factors - like size, market/book value, R&D intensity and leverage - that may influence expected stock returns (e.g. Bekaert et al. (2012)). The coefficient (standard error) on recession in these regressions is .019 (0.003), similar to the results obtained in column 7 of Table 1.

¹⁶We match public firms from the Census dataset to Compustat and regress the mean of a firm’s monthly stock returns at year $t+1$ on the sales weighted mean over the firm’s plants of TFP shocks calculated between t and $t+1$, including a full set of firm and year fixed effects. We find that the coefficient on the mean TFP shocks is highly significant with point estimate (standard error) of 0.0051 (0.0006).

¹⁷Vavra (2014) looks at changes in product-level prices (e.g. the price of a 2-litre bottle of Coke) and finds that these are also more dispersed during recessions.

Table 2 displays 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.¹⁸ 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 slowdowns in growth are associated with increases in the cross-sectional dispersion of shocks.

One immediate question is why within industry dispersion of shocks is higher during industry slowdowns. Maybe this is because industry slowdowns impact some types of establishments differently? To investigate this columns (2) to (9) run a series of estimations checking whether the increase in within industry dispersion is larger given 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 countercyclical 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 countercyclical in their dispersion? We again find nothing. The rest of the table reports similar results for the median and dispersion of plant size within each industry (measured by the number of employees, columns (4) and (5)), the median and dispersion of capital/labor ratios (columns (6) and (7)), and TFP and geographical dispersion interactions (columns (8) and (9)). In all of these we find insignificant coefficients on the interaction of industry growth with industry characteristics.

So, in summary, it appears that: first, the within-industry dispersion of establishment TFP shocks rises sharply when the industry growth rates slow down; and second, perhaps surprisingly, this relationship appears to be broadly robust across all industries.

2.3 Is Uncertainty a Cause or Effect of Slowdowns?

An obvious question regarding the relationship between uncertainty and the business cycle is the direction of causality. Does uncertainty drive the cycle, or do recessions drive increases in uncertainty, or does something else drive both? A recent literature has suggested a number of mechanisms for uncertainty to increase endogenously in recessions, so identifying the

¹⁸We 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).

direction of causation is clearly important in highlighting the extent to which countercyclical uncertainty is a shock driving cycles versus an endogenous mechanism amplifying cycles.¹⁹

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 obvious instrument exists at the macro level because as Kocherlakota (2010) 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.

The first instrument uses the Chinese accession to the WTO as a natural experiment, which led to the abolition of import quotas on Chinese textiles and apparel in 2005.²⁰ Since the Chinese accession to the WTO was agreed many years before, this instrument should not directly influence uncertainty itself. The second approach uses an instrumental variables approach which exploits movements in industry exchange rates to identify changes in industry growth, following Bertrand (2004).

In summary, and as Appendix Table A2 shows in detail, neither set of instrument finds any impact of firm-moment shocks on uncertainty. While these estimations have large standard errors, they do suggest that uncertainty shocks are not primarily driven by first-moment shocks. Instead, as the literature survey in the Bloom (2014) discusses, recessions appear to be initiated by a combination of negative first and positive second moment shocks, with ongoing amplification and propagation from further uncertainty induced by the recession.

2.4 Are Establishment-Level TFP Shocks a Good Proxy for Uncertainty?

The evidence we have provided for countercyclical aggregate and industry-level uncertainty relies heavily on using the dispersion of establishment-level TFP shocks as a measure of uncertainty. To 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.²¹

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 daily 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 daily stock

¹⁹See, for example, the papers on information collection by Van Nieuwerburgh and Veldkamp (2006) and Fagelbaum, Schaal and Tascherau-Dumouchel (2013), on experimentation in Bachmann and Moscarini (2011), on forecasting by Orlik and Veldkamp (2014), and on search by Petrosky-Nadeau (2013).

²⁰See Brambilla, Khandewal, and Schott (2010), Bloom, Draca, and Van Reenen (2011), and Appendix A.

²¹See, for example, Leahy and Whited (1996), Schwert (1989), Bloom, Bond, and Van Reenen (2007) and Panousi and Papanikolaou (2012).

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 ideal, 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 monthly returns rather than daily 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.²²

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 that 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 we find that this measure is also correlated with the average absolute size of establishment-level TFP shocks within the 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 indicators 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 countercyclicality of idiosyncratic (establishment, firm, and industry) uncertainty. Looking instead at macroeconomic uncertainty, there is already a growing literature providing evidence that this is also countercyclical, for example Schwert (1989), Campbell, Lettau, Malkiel, and Xu (2001), Engle and Rangel (2008) and Jurado, Ludvigsson and Ng (2014).

Rather than repeat this evidence here we simply include one additional model-specific empirical measure of aggregate uncertainty, which is the conditional heteroskedasticity of aggregate productivity A_t . This is estimated using a *GARCH*(1, 1) estimator on the Basu, Fernald, and Kimball (2006) data on quarterly TFP growth from 1972Q1 to 2010Q4. We find that conditional heteroskedasticity of TFP growth is strongly countercyclical, rising by 23% during recessions which is highly significant (a t-statistic of 5.27), with this series plotted in Appendix Figure A1.²³

²² As we did in column (7) of Table 1, to remove the forecastable component of stock returns we repeat columns 1 and 3 first removing the quarter by firm mean of firm returns. After doing this the coefficient (standard error) is very similar 0.324 (0.093) for column (1) and 0.387 (0.120) for column (3), mainly because the forecastable component of stock-returns explains a very small of total stock-returns.

²³ We also estimated a *GARCH*(1, 1) for monthly industrial production, including as many as twelve lags and

3 The General Equilibrium Model

We proceed by analyzing the quantitative impact of variation in uncertainty within a DSGE model. We consider an economy with heterogeneous firms that are subject to both first-moment 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.

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.²⁴ Specifically, a firm indexed by j produces output according to

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

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 allow the variance of innovations to the productivity processes, σ_t^A and σ_t^Z , to vary over time according to a two-state Markov chain.

3.1.2 Adjustment Costs

We allow for the presence of various types of convex and non-convex adjustment costs in capital and labor. As is well known in the literature, it is the presence of the non-convex adjustment costs that leads to a real options or wait-and-see effect of uncertainty shocks.

To be consistent with the existing evidence regarding the presence of multiple types of adjustment costs, we include the following forms of adjustment costs: in capital we allow

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.

²⁴ An alternative model has monopolistically competitive firms in which each firm produces a differentiated good. Note that the assumption of decreasing returns to scale allows us to pin down the firm's size.

for partial irreversibility, while in labor we allow for a fixed cost when changing the stock of labor, as well as hiring and firing costs.²⁵ We elaborate on these adjustment costs in what follows.

With respect to capital, we assume that a firm's capital stock evolves according to the standard law of motion

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

where δ_k is the rate of capital depreciation and $i_{j,t}$ denotes investment. The 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, given labor adjustment value $s_{j,t}$, is governed by

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

At each period a constant fraction δ_n of hours worked is exogenously destroyed due to retirement, illness, maternity leave, exogenous quits, etc.²⁶

We assume that whenever the firm chooses to 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. Note that these adjustment costs in labor imply that $n_{j,t-1}$ is a state variable for the firm.

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 , (4) aggregate productivity A , (5) macro uncertainty σ^A , (6) micro uncertainty σ^Z and (7) the joint distribution of idiosyncratic productivity and firm-level capital stocks and hours worked in the last period μ , which is defined for the space $S = R_+ \times R_+ \times R_+$.

The dynamic problem of the firm consists of choosing investment and hours to maximize

²⁵ See the literature focused on estimating labor and capital adjustment costs, including, Nickell (1986), Caballero and Engel (1999), Ramey and Shapiro (2001), Hall (2004), Cooper and Haltiwanger (2006), and Merz and Yashiv (2007). We incorporate all types of adjustment costs that have been estimated to be statistically significant at the 5% level in Bloom (2009).

²⁶ The assumption of exogenous separations is common in search models such as the Diamond-Mortensen-Pissarides model. For consistency with the prior business cycle literature we assume hiring adds to the labor force immediately and investment adds to the capital stock with a one period lag, although experimentation with these timing assumptions reveals they are not important for the results.

the present discounted value of future profits

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

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

$$\mu' = \Gamma(A, \sigma^A, \sigma^Z, \mu), \quad (9)$$

and the stochastic discount factor, m , which we discuss below in Section 3.4. $w(A, \sigma^A, \sigma^Z, \mu)$ denotes the wage rate in the economy while $AC^k(k, k')$ and $AC^n(n_{-1}, n)$ denote 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 investment 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

$$W(\phi, A, \sigma^A, \sigma^Z, \mu) = \max_{\{C, N, \psi'\}} \{U(C, N) + \beta \mathbb{E} [W(\phi', A', \sigma^{A'}, \sigma^{Z'}, \mu')]\}, \quad (10)$$

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

$$\begin{aligned} & C + \int q(k', n, z; A, \sigma^A, \sigma^Z, \mu) \psi'(dk' dndz) \\ & \leq w(A, \sigma^A, \sigma^Z, \mu)N + \int \rho(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu) \phi(dk dn_{-1} dz). \end{aligned} \quad (11)$$

Households receive labor income as well as the sum of dividends and the resale value of their investments priced at $\rho(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$. With these resources the household consumes and buys new shares at a price $q(k', n, z; A, \sigma^A, \sigma^Z, \mu)$ per share of the different firms in the economy. We denote by $C(\phi, A, \sigma^A, \sigma^Z, \mu)$, $N^s(\phi, A, \sigma^A, \sigma^Z, \mu)$, and $\Psi'(k', n, z; A, \sigma^A, \sigma^Z, \mu)$ 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 with idiosyncratic productivity z .

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 and policy functions solving (8) while W and $\{C, N^s, \Psi'\}$ are the value and policy functions solving (10). There is market clearing in the asset markets

$$\Psi'(k', n, z; A, \sigma^A, \sigma^Z, \mu) = \mu'(z, n, k'),$$

the goods market

$$\begin{aligned} & C(\phi, A, \mu) \\ = & \int_S \left[Azk^\alpha N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)^\nu - (K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu) - (1 - \delta_k)k) \right] \\ & \mu(dkdn_{-1}dz), \end{aligned}$$

and the labor market

$$N^s(\phi, A, \mu) = \int_S \left[N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu) \right] \mu(dkdn_{-1}dz).^{27}$$

Finally, the evolution of the joint distribution of z , k and n_{-1} 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 , along with the appropriate integration of firms' optimal choices of capital and hours worked given current state variables.

3.4 Sketch of the Numerical Solution

We briefly describe the solution algorithm, which heavily relies on the approach in Kahn and Thomas (2008) and Bachmann, Caballero and Engel (2013), with fuller details in Appendix B, and the full code available on-line.²⁸

The model can be simplified substantially if we combine the firm and household problems

²⁷Note that the distribution μ has inputs (z, n_{-1}, k) , so that the use of $\mu'(z, n, k')$ in the clearing condition above is an abuse of notation. More formally, we can write $\Psi'(k', n, z; A, \sigma^A, \sigma^Z, \mu) =$

$\int_{Q(k', n, z; A, \sigma^A, \sigma^Z, \mu)} \mu(dkdn_{-1})$ and

$Q(k', n, z; A, \sigma^A, \sigma^Z, \mu) = \{(k, n_{-1}) | k' = K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu), n = N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)\}.$

²⁸See <http://www.stanford.edu/~nbloom/RUBC.zip>

into a single dynamic optimization problem. From the household problem we get

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

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

where equation (12) is the standard optimality condition for labor supply and equation (13) is the standard expression for the stochastic discount factor. We assume that the momentary utility function for the household is separable across consumption and hours worked,

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

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

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

We also define the intertemporal price of consumption goods as $p(A, \sigma^Z, \sigma^A, \mu) \equiv U_C(C, N)$. This then allows us to redefine the firm's problem in terms of marginal utility, denoting the new value function as $\tilde{V} \equiv pV$. The firm problem can then be expressed as

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

To solve this problem we employ nonlinear techniques that build upon Krusell and Smith (1998). Again, detailed discussion of the algorithm is provided in Appendix B.

4 Calibration and Simulation

4.1 Parameter Values

This section motivates the choice of parameter values used in the simulations (see Table 5).

Frequency and Preferences We set the time period to equal a quarter. The household's discount rate, β , is set to match an annual interest rate of 5%. 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 (2013) we make the simplifying assumption that the Frisch 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

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

We set the parameter θ such that households spend a third of their time working in the non-stochastic steady state.

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 set to 35%. This estimate is based on Shimer (2005). We set the exponents on capital and labor in the firm's production function to be $\alpha = 0.25$ and $\nu = 0.5$, consistent with a capital cost share of $1/3$.

The existing literature provides a wide range of estimates for capital and labor adjustment costs.²⁹ 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. The resale loss of capital amounts to 34%. 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.

Aggregate and Idiosyncratic TFP Processes Productivity, both at the aggregate and the idiosyncratic level, is determined by AR(1) processes as specified in equations (2) and (3). The serial autocorrelation parameters ρ^A and ρ^Z are set to 0.95, similar to the quarterly value used by Kahn and Thomas (2008).

The Calibrated Process for Uncertainty Since we are interested in studying the effect of uncertainty shocks we assume that the uncertainty process is independent of the first-moment shocks. We assume for simplicity that the stochastic volatility processes, σ_t^A and σ_t^Z , follow a two-point Markov chain:

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

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

Since we cannot directly observe the stochastic process of uncertainty in the data, the calibration has to be guided by the impact of uncertainty on observable cross-sectional and

²⁹See, for example, Hayashi (1982), Nickell (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).

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 $\pi_{H,L}^{\sigma^Z}$. As a simplification to ease computational constraints we assume that a single process determines the economy's uncertainty regime, since empirically we see both microeconomic and macroeconomic uncertainty rising and falling together through the business cycle. This reduces the number of parameters to six: σ_L^A , σ_H^A , σ_L^Z , σ_H^Z , $\pi_{L,H}^{\sigma}$ and $\pi_{H,L}^{\sigma}$ since σ^A and σ^Z follow the same Markov process but with different levels of σ^A and σ^Z .³⁰

We calibrate these six parameters against eight moments from U.S. data. At the microeconomic level, we target the mean, standard deviation, skewness, and autocorrelation of the time series of the cross-sectional interquartile range of establishment TFP shocks computed from our annual Census sample covering 1972-2010. We display these moments in Table 6. At the macro level, we target the same four moments based on the time series of estimated heteroskedasticity in the annualized quarterly growth rate of the U.S. Solow residual, again covering 1972-2010.

In the model, we simulate 5000 firms for 5000 quarters and compute each of the eight moments discussed above, reporting the results in Table 6.³¹ The simulated data generally reproduces well the time series properties of volatility measures in the U.S. data. At the microeconomic level, the model implies a high level and standard deviation of cross-sectional dispersion, similar to the data, with somewhat lower serial correlation and skewness than in our Census sample. At the aggregate level, the model successfully produces lower mean levels of volatility and fluctuations in volatility for the Solow residual, both consistent with the U.S. data. We are also successful at reproducing the high persistence and skewness evident in U.S. aggregate volatility.³²

Based on our preferred calibration we find that periods of high uncertainty occur with a quarterly probability of 5%. The period of heightened uncertainty is quite persistent with a quarterly probability of 92% of staying in the high uncertainty state. Idiosyncratic volatility

³⁰That is, we assume one underlying two-state Markov chain $\tilde{S} \in \{L, H\}$ governing the timing of uncertainty shocks for both microeconomic and macroeconomic volatility. When $\tilde{S} = L$, microeconomic and macroeconomic volatility are equal to their low values σ_L^Z and σ_L^A . When $\tilde{S} = H$, microeconomic and macroeconomic volatility are equal to their high values σ_H^Z and σ_H^A .

³¹We discard the first 500 periods of this 5000-period simulation to eliminate the impact of initial conditions upon the business cycle statistics.

³²When computing the microeconomic moments in the model data, we first aggregate to annual frequency and then compute measured TFP for each firm, using a capital share of 1/3 and a labor share of 2/3 to approximate the factor share approach we use in the Census data. We also account for measurement error in our Census sample by adding noise to the underlying firm-level measures in the simulated data, assumed to have an equal variance to the underlying TFP shocks. This is consistent with the finding in Collard-Wexler (2011), and also with our own estimates for measurement error (see Appendix C for more details on our calculations).

To compute aggregate moments from the model data, we compute the aggregate annualized changes in the quarterly Solow residual, imposing labor and capital shares of 2/3 and 1/3. These shares are approximately equal to those used in the construction of John Fernald's quarterly measure in U.S. data.

is set to a low value of 3.9% and approximately triples in the heightened uncertainty state. Aggregate volatility is at a low of 0.58% and approximately doubles when an uncertainty shock hits.

4.2 Business Cycle Statistics

Table 7 illustrates that this calibration generates second-moment statistics that resemble their empirical counterparts in U.S. data.³³ As in the data, investment and hours commove with output. Output and consumption commove, although not as much as in the data. Investment is more volatile than output, while consumption is less volatile. Interestingly the model also generates a realistic volatility of hours relative to output, which has traditionally been difficult to achieve in RBC models.³⁴

5 The Effects of an Uncertainty Shock

As has been known since at least Scarf (1959), non-convex adjustment costs lead to Ss investment and hiring policy rules. Firms do not hire and invest until productivity reaches an upper threshold (the S in Ss) or fire and disinvest until productivity hits a lower threshold (the s in Ss). This is shown for labor in Figure 5, which plots the distribution of firms by their productivity/labor (Az/l) ratios after the micro and macro shocks have been drawn but before firms have adjusted. On the right is the firm-level hiring threshold (right black line) and on the left the firing threshold (left black line). Firms to the right of the hiring line will hire, firms to the left of the firing line will fire, and those in the middle will be inactive for the period.

An increase in uncertainty increases the returns to inaction, shown by the increased hiring threshold (right gray line) and reduced firing threshold (left gray line). When uncertainty is high firms become more cautious as labor adjustment costs make it expensive to make a hiring or firing mistake. Hence, the hiring and firing thresholds move out, increasing the range of inaction. This leads to a fall in net hiring, since the mass of firms is right-shifted due to labor attrition. A similar phenomena happens with capital, whereby increases in uncertainty reduce the amount of net investment.

5.1 Modelling a Pure Uncertainty Shock

To analyze the aggregate impact of uncertainty, we simulate 400 economies, where each economy has 20,000 firms. Each economy is of 100-quarter length. In the first ten periods

³³We simulate the model for 5000 firms over 5000 periods and compute the standard set of business cycle statistics.

³⁴See Rogerson (1988), Hansen (1985) and Benhabib, Rogerson, and Wright (1991).

of the economy we force uncertainty to be low. We then force uncertainty to jump in period 11. Uncertainty is then left to freely evolve according to its Markov transition process. We exclude the first 25 economies that are used to initialize the distribution over z , k , and n_{-1} , and we average the results over the remaining 375 simulated economies.

The impact on output is plotted in Figure 6, displaying a drop of just over 3% within one quarter, and then a recovery to close to level within one year. This is substantial and suggests that uncertainty shocks can be quantitatively important in driving business cycles.

These dynamics in output arise from three channels: labor, capital, and the misallocation of factors of production, shown in Figure 7. First, in the top-left panel we plot the time path of the aggregate labor force. When uncertainty increases most firms pause hiring, and the labor force starts to drop because workers are continuing to attrit from firms without being replaced. Importantly, in the model this rate of attrition is assumed to be constant over the cycle. This is consistent with Shimer (2005) and Hall (2005), which show that around three quarters of the movements in the volatility of unemployment are due to job-finding rates and not to the cyclicity of the destruction rate.

Second, in the top-right panel we plot the time path of investment, which drops rapidly due to the increase in uncertainty. Since investment falls but capital continues to depreciate, there is also a drop in the capital stock. Finally, in the bottom-left panel we show that productivity, measured as the Solow residual, also drops after the uncertainty shock.³⁵ This occurs because uncertainty increases the misallocation of factors across firms. In normal times, unproductive firms contract and productive firms expand, helping to maintain high levels of aggregate productivity. But when uncertainty is high, firms reduce expansion and contraction, shutting off much of this productivity-enhancing reallocation. This slow-down in reallocation manifests itself as a fall in measured aggregate TFP. The bottom-right panel of Figure 7 provides another measure of the reduction in reallocation in terms of an increase in misallocation. It plots the cross-sectional variation of the marginal product of labor (MPL) across plants in the model. This steady-state value is positive since adjustment costs prevent instant reallocation of factors across plants. But once the uncertainty shock hits the economy there is an increase of about 25% in this measure of misallocation, due to uncertainty impeding firms' ability to efficiently reallocate labor.³⁶

In the longer run, labor, investment, and TFP all start to recover to their steady state as the uncertainty shock is temporary. As uncertainty falls back firms start to hire and invest again to address their pent-up demand for labor and capital. But the pace of the

³⁵ The Solow residual is defined as $\log(Y_t) - \alpha \log(K_t) - \nu \log(L_t)$, where Y_t , K_t , and L_t are aggregate output, capital, and labor, and α and ν are the decreasing returns to scale production elasticities.

³⁶ Note that since the real options mechanism affects both capital and labor inputs, the increase in misallocation also occurs in capital as well, with the dispersion of the marginal product of capital (MPK) following a path qualitatively similar to that of misallocation in the MPL.

rebound is tempered by the desire of consumers to smooth consumption, leading to a slow gradual recovery. Hence, an uncertainty shock generates a short and sharp recession and a prolonged recovery.

In Figure 8 we plot the time profile of consumption. When the uncertainty shock occurs consumption jumps up in the first quarter, before subsequently falling back below trend for several 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, "freeing" up consumption. From period two onwards investment and labor rebound, so that consumption falls.

This overshoot in consumption may seem surprising, because uncertainty makes risk-averse consumers want to increase their precautionary savings. But offsetting this desire to increase precautionary savings is the fact that the returns to savings have now decreased. Consumers can only save through capital, and the returns to capital have also become more risky with the rise in uncertainty. Hence, on impact the uncertainty shock both temporarily reduces prices (increasing the attractiveness of consumption) and increases the riskiness of capital (reducing the attractiveness of saving), leading to an overshoot in consumption.

Clearly, this rise in consumption at the start of recessions is an unattractive feature of a pure uncertainty shock model of business cycles. Several options exist, however, to try and address this. One is to allow consumers to save in other technologies besides capital, for example, in foreign assets. This is the approach Fernandez-Villaverde, Guerron-Quintana, Rubio-Ramirez, and Uribe (2011) take in modelling risk shocks in small open economies. In an open economy model a domestic uncertainty shock induces agents to increase their savings abroad (capital flight). In our closed model this is not possible, but extending the model to allow a foreign sector would make this feasible although computationally more intensive. Another option would be to use utility functions such as those in Greenwood, Hercowitz, and Huffman (1998). Due to the complementarity between consumption and hours in such preference structures, they should reduce the overshoot in consumption. As we discuss below, another option is to augment our uncertainty shock with a small first-moment shock. Adding a negative first-moment shock reduces the supply of goods during a recession, eliminating the initial consumption overshoot.

5.2 First-Moment and Second-Moment Shocks

The evidence in Section 2 suggests that recessions are periods of both first- and second-moment shocks. So, to generate an empirically more realistic simulation we consider the combination of an uncertainty shock *and* a -2% first moment shock. Specifically, we consider an economy where both a negative first-moment and a positive second-moment shock

hit the economy at the same period, namely period one in Figure 9.³⁷

As Figure 9 suggests, this additional shock magnifies the drop in output and investment. The size of these fluctuations are now large enough to account for the drop and rebound in output even during the most recent recession. The addition of the first-moment shock also leads to a fall in consumption on impact.

Hence, both empirically and in the simulation we find that recessions appear to be well characterized as a combination of a negative first-moment shock and a positive second-moment shock. Having a first-moment component of the shock helps to fit the time series of consumption. Having a second-moment component of the shock helps to reduce the size of the first-moment shock necessary to lead to recessions, while also generating a U-shaped path of output over the business cycle.

6 Decomposing the Impact of Uncertainty

How do the effects of uncertainty shocks differ across a General Equilibrium (GE) framework and Partial Equilibrium (PE)? To address this question we plot in Figure 10 the impact of an uncertainty shock in three different economies. The black line (\times symbols) depicts again the effects of an uncertainty shock in our GE model economy, the red line (+ symbols) depicts the same response but with PE only (all prices and wages are held constant and the consumer side of the economy ignored), while the blue line (o symbols) depicts the effects of an uncertainty shock in PE economy with no adjustment costs at all.

As the blue line (o symbols) suggests, when there are no adjustment costs of any type in a PE economy, output actually increases following an uncertainty shock. The reason for this result is related to the Oi (1961), Hartman (1972) and Abel (1983) effect, whereby a higher variance of productivity increases investment, hiring, and output because the optimal capital and labor choices are convex in productivity.³⁸

As the red line (+ symbols) suggests, the addition of adjustment costs to the PE setup dramatically changes the effect of an uncertainty shock. Now, on impact there is a fall in aggregate output. The reason is the increase in uncertainty moves firms' labor and capital Ss bands out, temporarily pausing hiring and investment. If all firms pause hiring and investment, aggregate labor and capital drop due to labor attrition and capital depreciation. But this pause is short-lived, as once uncertainty drops back firms start to hire and invest

³⁷See Appendix B for a description of the algorithm.

³⁸To be precise, if $Y = AK^aL^b$ with $a + b < 1$, the per period rental cost of capital is r , and the wage rate is w , then 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, 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 Oi-Hartman-Abel effect after Oi (1961), Hartman (1972), and Abel (1983).

again. So in the medium-run the Oi-Hartman-Abel effect dominates and output rises above its long-run trend.³⁹

While these forces are also present in the baseline GE and adjustment cost economy, the curvature in the utility function (i.e. the endogenous movement in the interest rate) moderates the rebound and overshoot.⁴⁰ The overshoot in the PE economy requires big movements in investment and labor, which would lead to excessively large swings in consumption. The curvature in utility slows down the rebound of the GE economy, generating a smoother and more persistent output cycle.

Intriguingly, in the first period, however, GE has very little impact on output because the Ss bands have moved so far out that there is almost no density of firms near the hiring or investment threshold to respond to prices. Hence, the short-run robustness of the impact of uncertainty shocks to GE suggested by Bloom (2009) seems to be present, while the medium-run sensitivity to GE highlighted by Thomas (2002) and Khan and Thomas (2003, 2008) is also present.

7 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 experiments as a means of documenting the effects of such policies in times of heightened uncertainty. It is also worth noting that this ignores the direct impact of policy on uncertainty, as this is too complex to include in this model.⁴¹

The policy experiment we consider is 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 one quarter and financed through a lump-sum tax on households. 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 attempt to identify the effect of uncertainty on policy effectiveness.

Figure 11 depicts these experiments. The black line (\times symbols) is the impact of the

³⁹ Interestingly, this overshoot is actually more persistent than with PE and no adjustment costs because the introduction of adjustment costs smooths the overshoot and rebound, similar to the results on the smoothing effects of microeconomic rigidities in, for example, Caballero and Engel (1993).

⁴⁰ This is similar to the difference between PE and GE economies' convergence to the steady state when starting with a capital stock below the steady state. A PE economy converges within one period to the steady state value. In a GE economy, due to the curvature in utility, the convergence is slower. This effect is also present in small open economy models that are de facto PE economies. In such models, absent adjustment costs, the behavior of investment is very volatile.

⁴¹ For a discussion and literature review of the impact of policy on uncertainty see, for example, Baker, Bloom, and Davis (2012).

1% wage subsidy in an economy that is not hit by an uncertainty shock. Not surprisingly, the artificially reduced wage stimulates hiring and increases output which then gradually returns to its long run trend. The red line (+ symbols) shows the net impact of the 1% wage subsidy applied at the same time as the uncertainty shock hits. The presence of uncertainty reduces the effects of the wage policy by 26% on impact (by 46% over the following four quarters). The reason is that as soon as uncertainty rises, the S s thresholds jump out, so most firms are far away from their hiring and investment thresholds, making them less responsive to any policy stimulus.

Finally, the blue line (o symbols) in Figure 11 represents the impact of the policy when implemented four quarters after the uncertainty shock. Now, the effectiveness of the policy is above its baseline value during periods of low uncertainty, with output increasing by 9% more on impact than during normal times (and by 15% more than normal times over four quarters). The reason is that as the distribution of firm-level TFP fans out towards the thresholds there is now more density near the hiring threshold, so the economy is more responsive to wage subsidies than usual.

Overall, this highlights how uncertainty shocks lead to time-varying policy effectiveness. At the instant an uncertainty shock hits, policy is not as effective relative to normal times, while once uncertainty starts to drop down, policy becomes more effective. Hence, uncertainty shocks not only impact the economy directly, but also indirectly change the response of the economy to any potential reactive stabilization policy.

8 Conclusions

Uncertainty has received substantial attention as a potential factor shaping the depth and duration of the Great Recession. We formally model this as a shock helping to drive business cycles. The first part of the paper uses Census microdata to show that measured uncertainty is indeed 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.

The second part of the paper then builds a DSGE model with heterogeneous firms, time-varying uncertainty, and adjustment costs to quantify the impact of these second-moment shocks. We find that they typically lead to drops of about 3% in GDP, with a sharp drop and slow recovery. This suggests that uncertainty could play an important role in driving business cycles. We also find that because uncertainty makes firms cautious but also makes firm-level outcomes more volatile, it substantially changes the response of the economy to

stimulative policy.

In future research we would like to extend the theoretical work by including trade flows, modelling a more sophisticated policymaking process under uncertainty, and allowing for a distinction between durable and nondurable consumption, with real options effects on durable consumer behavior.

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.
- Arellano, C., Y. Bai, P. Kehoe, 2012, "Financial Markets and Fluctuations in Uncertainty", Federal Reserve Bank of Minneapolis working paper.
- Bachmann, R. and Bayer, R. (2014), "Investment Dispersion and the Business Cycle", *American Economic Review*, 104(4), 1392-1416.
- Bachmann, R., Caballero, R.J. and Engel, E.M.R.A. (2013), "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model", *American Economic Journal: Macroeconomics*, 5(4), 29-67.
- Bachmann, R. and Moscarini, G. (2011), "Business Cycles and Endogenous Uncertainty", Yale mimeo.
- Baker, S., Bloom, N. and Davis, S. (2012), "Measuring economics policy uncertainty", Stanford mimeo.
- Baker, Scott and Nicholas Bloom (2011). "Does uncertainty drive business cycles? using disasters as natural experiments". NBER Working Paper 19475.
- Bartelsman, Eric J., Becker, Randy A. and Gray Wayne B. "NBER-CES Manufacturing Industry Database", <http://www.nber.org/nberces/nbprod96.htm>
- Basu, S. and Bundick, B. (2011), "Uncertainty Shocks in a Model of Effective Demand", Boston College WP 774.
- Basu, S., Fernald, J. and Kimball, M.S. (2006), "Are Technology Improvements Contractionary?", *American Economic Review*, 96(5), 1418-1448.
- Bekaert, Geert, Robert J. Hodrick, and Xiaoyan Zhang. (2012). "Aggregate Idiosyncratic Volatility". *Journal of Financial and Quantitative Analysis*, 47 (6):1155-1185.
- Benhabib, Jess, Rogerson, Richard, and Wright, Randall (1991), "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy*, 99(6), 1166-1187.
- Bernanke, B. (1983), "Irreversibility, Uncertainty and Cyclical Investment", *Quarterly Journal of Economics*, 98, 85-106.
- Bertola, G. and Caballero, R.J. (1994), "Irreversibility and Aggregate Investment", *Review of Economic Studies*, 61, 223-246.
- 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.
- Bloom, N. (2009), "The Impact of Uncertainty Shocks", *Econometrica*, May 77, pp. 623-685..
- Bloom, N., Bond, S. and Van Reenen, J. (2007), "Uncertainty and Investment Dynamics", *Review of Economic Studies*, April, 74, pp. 391-415..
- 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 (1993), "Microeconomic adjustment hazards and aggregate dynamics", *Quarterly Journal of economics*, 108, 359-383.

Caballero, R.J. and Engel, E.M.R.A (1999), "Explaining Investment Dynamics in U.S. Manufacturing: a Generalized Ss 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. (2014), "Financial factors in economic fluctuations", forthcoming *American Economic Review*.

Collard-Wexler, Allan (2011), "Productivity Dispersion and Plant Selection in the Ready-Mix Concrete Industry", working paper.

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

Cooper, Russell, Haltiwanger, John, and Willis, Jonathan L. (2004), "Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications," NBER Working Paper 10297.

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

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.

Fajgelbaum, Pablo, Schaal, Eduardo and Taschereau-Dumouchel, Mathieu, (2014), "Uncertainty traps", NBER Working Paper No. 19773.

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.

Fernandez-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramirez, J. and Uribe, M. (2011), "Risk matters: the real effects of volatility shocks." *The American Economic Review* 101(6), 2530– 2561.

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 Syverson, 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.

- Greenwood, Jeremy, Hercowitz, Zvi, and Huffman, Gregory W. (1988), "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78(3), 402-417
- Gourio, Francois and Anil K. Kashyap. 2007. Investment Spikes: New Facts and a General Equilibrium Exploration. *Journal of Monetary Economics* 54:1-22.
- Güvenen, F, Ozkan, S. and Song, J., (2014), "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122(3), 621-660.
- Hall, Robert E. (2004), "Measuring Factor Adjustment Costs," *Quarterly Journal of Economics*, 119(3), 899-927.
- Hall, Robert E. (2005), "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," *NBER Macroeconomics Annual*, 101-137.
- Hansen, Gary D. (1985), "Indivisible labor and the business cycle," *Journal of Monetary Economics*, 16(3), 309-327.
- Hartman, R. (1972), "The effect of price and cost uncertainty on Investment", *Journal of Economic Theory*, 5, 258-266.
- Hassler, J. (1996), "Variations in Risk and Fluctuations in Demand - a Theoretical Model", *Journal of Economic Dynamics and Control*, 20, 1115-1143.
- Hayashi, Fumio (1982), "Tobin's Marginal q and Average q : A Neoclassical Interpretation," *Econometrica*, 50(1), 213-224.
- Hopenhayn, H and Rogerson, R. (1993), "Job turnover and policy evaluation: a general equilibrium analysis", *Journal of Political Economy*, vol 101, pp 915-938.
- House, C.L. (2014), "Fixed Costs and Long-Lived Investments", *Journal of Monetary Economics* 68, 86-100.
- Jurado, Kyle, Ludvigson, Sydney and Ng, Serena, (2014), "Measuring uncertainty", forthcoming *American Economic Review*
- Kehrig, Matthias, (2011), "The cyclical productivity dispersion", working paper.
- Khan, A. and Thomas, J.K. (2003), "Nonconvex factor adjustment in equilibrium business cycle models: do nonlinearities matter?" *Journal on Monetary Econometrics*, 50, 331-360.
- 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. (2013), "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity," *Journal of Political Economy*, 121, 1055-1107.
- King, R.G. and Rebelo, S.T. (1999), "Resuscitating Real Business Cycles", in *Handbook of Macroeconomics*, John B. Taylor and Michael Woodford (eds.), Elsevier.
- Knotek, Edward S. and Khan, Shujaat (2011), "How do households respond to uncertainty shocks?" *Federal Reserve Bank of Kansas City Economic Review*, Second Quarter, 63-92.
- 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
- Leahy, J. and Whited, T. (1996) "The Effects of Uncertainty on Investment: Some Stylized Facts", *Journal of Money Credit and Banking*, 28, 64-83.
- Lillien, D.M. (1982), "Sectoral Shifts and Cyclical Unemployment", *Journal of Political Economy*, 90(4), 777-793.
- 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.
- Oi, Walter (1961) "The Desirability of Price Instability Under Perfect Competition," *Econometrica*, 29(1), 58-64.
- Orlik, Anna and Veldkamp, Anna (2014), "Understanding uncertainty shocks and the role of the black swan", NYU mimeo.
- Panousi, V. and Papanikolaou, D. (2012), "Investment, Idiosyncratic Risk, and Ownership", *Journal of Finance*, 67(3).
- Petrosky-Nadeau, N. and Wasmer, E. (2013), "The Cyclical Volatility of Labor Markets under Frictional Financial Markets", *American Economic Journal: Macroeconomics*, 5(1).
- Pierce, Justin R., and Schott, Peter K.. (2009), "Concording U.S. Harmonized System Categories Over Time", NBER working paper 14837.
- 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.
- Rogerson, Richard (1988), "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, 21(1), 3-16.
- Romer, C. (1990) "The Great Crash and the Onset of the Great Depression.", *Quarterly Journal of Economics*, 105, 597-624
- Scarf, Herbert (1959) "The optimality of policies in the dynamic inventory problem," in *Mathematical Methods in the Social Sciences*, ed. Kenneth J. Arrow, Stanford University Press.
- Schott, Peter K. (2008), "The relative sophistication of Chinese exports", *Economic Policy*, 23(53), 5-49
- Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?", *Journal of Finance*, 44, 1115-1153.
- Shapiro, Matthew D. (1986), "Investment, Output, and Cost of Capital," *Brookings Papers on*

Economic Activity, 111-152.

Shimer, Robert (2005), "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," American Economic Review, 95(1), 25-49.

Stock, James H. and Watson, Mark W. (2012), "Disentangling the Channels of the 2007-2009 Recession," working paper.

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

Tauchen, George (1986) "Finite State Markov Chain Approximations to Univariate and Vector Autoregressions," Economics Letters, 20(2), 151-155.

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

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

Vavra, Joseph (2014), "Inflation dynamics and time-varying uncertainty: new-evidence and an Ss interpretation", Quarterly Journal of Economics, 129(1).

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 Online Appendix: Census Uncertainty Data

We use data from the Census of Manufactures (CM) and the Annual Survey of Manufactures (ASM) from the U.S. 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 first-moment 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 2010 (using the LBDNUM identifier).⁴² 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 a 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, cost of materials, 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.⁴³ 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 (1972-2010), we assigned each FIRMID after 2001 to the last observed CUSIP8 and before 1976 to the first observed CUSIP8⁴⁴.

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,

⁴²The 2010 ASM data became available only very recently. To avoid repeating extensive census disclosure analysis, in Tables 2, 3 and 4 we use data only up to 2009.

⁴³The access to CRSP and Compustat data sets is through WRDS: <https://wrds-web.wharton.upenn.edu/wrds/>.

⁴⁴We do this assignment for 2002-2005, since the bridge has many missing matches for these years.

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), and capital expenditures (PIINV).⁴⁵ It is also the main source for industry-level total cost of inputs for labor (PAY). The total cost variable is 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 1997.⁴⁶ We complete our set of price indices using FRB capital investment deflators, with separate deflators for equipment and structures, kindly provided to us by Randy Becker.

The BLS multifactor productivity database is used for constructing data on industry-level cost of capital and capital depreciation.⁴⁷ 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 an 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 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.⁴⁸ 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.⁴⁹ The data is at a monthly frequency and is provided at NAICS 3-digit to 6-digit 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 4-digit SIC87 codes.⁵⁰ We use the Cost, Insurance, and Freight (CIF) definition of imports by industry from this data set when we construct weighted industry-level exchange rate indices which are used as instruments for first-moment 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 U.S. dollar.⁵¹ We focus on G-20 countries as these have large diversified economies so should have exchange rate movements which are exogenous 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 problematic. We obtain price deflators for the 15 countries from the OECD website.⁵²

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

⁴⁵ See: <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.

⁴⁶ The NBER-CES data are available only through 2009. 2010 industry-level data are therefore projected using an AR(4) regression for all external datasets.

⁴⁷ See <http://www.bls.gov/mfp/mprdownload.htm>.

⁴⁸ See <http://www.bea.gov/national/FA2004/SelectTable.asp>.

⁴⁹ See <http://www.federalreserve.gov/releases/G17/Current/default.htm>.

⁵⁰ See: 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, Becker, and Gray (2000).

⁵¹ See <http://www.imfstatistics.org/IMF/imfbrowser.aspx?branch=ROOT>.

⁵² See <http://stats.oecd.org/index.aspx?querytype=view&queryname=221>.

(see the data appendix in their paper for more details on construction of the quotas).

A.2 Sample Selection

We have five main establishment 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 establishments which are not used in Census tabulation (TABBED=N), establishments with missing or nonpositive data on total value of shipments (TVS) and establishments with missing values for LBDNUM, value added (VA), labor inputs or investment. We also require each establishment to have at least one record of capital stock (at any year). This sample consists of 211,939 establishments and 1,365,759 establishment-year observations.

The second sample, which is our baseline sample, keeps establishments which appear for at least 25 years between 1972 and 2009. This sample consists of 15,673 establishments and 453,704 establishment-year observations.⁵³

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

The fourth sample uses a balanced panel of establishments which were active for all years between 1972 and 2009. This sample consists of 3,449 establishments and 127,182 establishment-year observations.

Our last sample (used in Figures 1 and 2), is based on the first sample, but includes only establishments that were active in 2005, 2006, 2008 and 2009.

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 for at least 25 years.

A.3 Variable 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

⁵³ As the 2010 ASM data became available only very recently, whenever 2010 data is used we keep the sample of establishments unchanged. For example, we choose establishments that were active for 25 years between 1972 and 2009, but use data for these establishments also from 2010.

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 establishment-level series of capital stocks using the plant's initial level of capital stock, the establishment-level investment data and industry-level depreciation rates. 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 then apply the formula $K_t = (1 - \delta_t) K_{t-1} + I_t$.⁵⁴ This procedure is done separately for structures and for equipment. However, starting in 1997, the CM does not separately report 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 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 $c_{i,t}^N$ 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 2-digit industry level. To obtain capital income at 4-digit industry level we apply the ratio of capital income to capital input calculated using BLS data to the 4-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:

$$\log(\hat{z}_{j,i,t}) = \rho \log(\hat{z}_{j,i,t-1}) + \mu_i + \lambda_t + e_{j,i,t}, \quad (19)$$

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

A.3.5 Microeconomic 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 i 's growth rate 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: In 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,

⁵⁴ The stock is reported for the end of period, therefore we use last period's stock with this period's depreciation and investment.

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 $(x_{i,t+4} - x_{i,t}) / (0.5 \times x_{i,t+4} + 0.5 \times x_{i,t})$. We use the industrial production index constructed by the FRB to calculate the cross-sectional dispersion of industry production growth $(x_{i,t+12} - x_{i,t}) / (0.5 \times x_{i,t+12} + 0.5 \times x_{i,t})$ at the monthly level.

Firm Level: In 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 it is calculated 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-level 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 (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 a firm's annual growth calculated at a quarterly rate $(x_{i,t+4} - x_{i,t}) / (0.5 \times x_{i,t+4} + 0.5 \times x_{i,t})$.

Industry Level: For industry level measures of volatility we use the standard deviation over a year of an industry's annual growth calculated at a monthly rate $(x_{i,t+12} - x_{i,t}) / (0.5 \times x_{i,t+12} + 0.5 \times x_{i,t})$ 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 the mean or the median plant's real sales growth rates within industry year. Industry characteristics 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 the 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 First-Moment 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 first-moment 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 first-moment shock to this industry in the U.S. 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 4-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.⁵⁵ 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-year 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 Eurozone, 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 prior to 1999 using the December 31, 1998 exchange rate.⁵⁶ We then use the data from Peter K. 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) definition of imports. 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

$$w_{i,c} = \frac{CIF_{i,c}}{\sum_c CIF_{i,c}}.$$

B Online Appendix: Model Solution and Simulation

In this appendix, we first lay out our algorithm for numerically solving the model and then discuss our approach to simulating the response of the economy to various shocks. The full code is also available at <http://www.stanford.edu/~nbloom/RUBC.zip>.

B.1 Solving the Model

To numerically solve for the recursive general equilibrium of the model, we follow Kahn and Thomas

⁵⁵ For detailed description of the quotas system and its effect on China’s export see Bloom, Draca, and Van Reenen (2011) and Brambilla, Khandewal, and Schott (2010).

⁵⁶ This is available in an ECB press release from December 31, 1998, available at http://www.ecb.int/press/pr/date/1998/html/pr981231_2.en.html.

(2008) by using an adaptation of the algorithm laid out in Krusell and Smith (1998).

Household Optimality Conditions: To tractably compute general equilibrium, we combine household optimization with an individual firm's problem. Household optimization implies standard expressions for labor supply in terms of the wage w and for a household's stochastic discount factor m , given by:

$$w = -\frac{U_N(C, N)}{U_C(C, N)}, m = \beta \frac{U_C(C', N')}{U_C(C, N)}.$$

The flow utility function for a household takes the following separable form across consumption C and hours worked N :

$$U(C, N) = \frac{C^{1-\eta}}{1-\eta} - \theta \frac{N^\chi}{\chi},$$

so that the wage rate w and the stochastic discount factor m relevant for the firm are given by

$$w = \frac{\theta N^{\chi-1}}{C^{-\eta}}, m = \beta \frac{C'^{-\eta}}{C^{-\eta}}.$$

Transformed Firm Problem: We define the intertemporal price of consumption $p \equiv U_C(C, N)$, equal to the marginal utility of consumption for households. We then redefine a firm's value function V in terms of marginal utility and consider the new firm value function $\tilde{V} \equiv pV$. A new and entirely equivalent firm recursive optimization problem can then be expressed as

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

where discounting now occurs at the fixed rate β .

Note that as discussed in the calibration section in the text, we assume that the transition between high and low uncertainty states occurs with identical timing at the idiosyncratic and aggregate levels. We therefore replace the aggregate and idiosyncratic volatility states σ^A and σ^Z above with $\tilde{S} \in \{L, H\}$, where \tilde{S} is the two-state Markov process for uncertainty.

Also, note that based on our discussion of the household problem, in equilibrium the wage rate depends on the marginal utility of consumption and on labor input in the current period. For computational convenience and following Kahn and Thomas (2008), we further simplify this relationship by calibrating the model to match a unit intertemporal elasticity of substitution ($\eta = 1$) and an infinite Frisch elasticity of labor supply ($\chi = 1$). Together these parameter values imply that $w = \frac{\theta}{C^{-1}} = \frac{\theta}{p}$, or that the wage rate is a function of the intertemporal price of consumption only.

Forecast and Price Rules: The aggregate states of the economy are economy-wide productivity A , uncertainty \tilde{S} , and the distribution of individual firm capital stocks, labor input, and idiosyncratic productivity $\mu(k, n_{-1}, z)$. However, the distribution μ is a numerically intractable state. We therefore follow Krusell and Smith (1998) by replacing the distribution μ in the firm's problem with a small set Ω of moments of the distribution. Further, in the expectations embedded in the firm's optimization we replace the law of motion for this distribution $\mu' = \Gamma(A, \tilde{S}, \mu)$, with a simplified forecast rule $\hat{\Omega}'$ for next period's moments Ω' . Finally, we also assume that during optimization, firms make use of an approximate pricing rule \hat{p} to predict the intertemporal price p . Given our parameter calibration, the pricing rule \hat{p} also predicts the equilibrium wage rate, implying that an additional wage rule is unnecessary.

Algorithm: Armed with the transformed firm problem and notation defined above, we now discuss the solution algorithm in detail. Essentially, numerical solution of the model involves calculation of pricing and forecast rules consistent with the simulated behavior of the model given firm optimization and the evolution of the exogenous stochastic processes. First, we choose specific log-linear approximate pricing and forecast rules to be used by firms:

$$\begin{aligned}\log \hat{p} &= \beta_0^p(A, \tilde{S}) + \beta_1^p(A, \tilde{S})K \\ \log \hat{K}' &= \beta_0^K(A, \tilde{S}) + \beta_1^K(A, \tilde{S})K.\end{aligned}$$

As evident from these formulas, in our application we replace the set of moments Ω with the aggregate capital stock K , and we allow for general dependence of the coefficients β^p and β^K on some of the aggregate states of the economy. Given the explicit functional form above for the forecasting and pricing rules, we now state the algorithm for solving the model:

1. Initialize the forecasting and price rules $\hat{K}'^{(0)}$ and $\hat{p}^{(0)}$ based on initial guesses for the coefficients β^p and β^K .
2. Given rules $\hat{K}'^{(l)}$ and $\hat{p}^{(l)}$, solve a discretized version of the dynamic optimization problem of an individual firm using value function iteration to determine $\tilde{V}^{(l)}$.
3. Simulate the economy for a large number of periods, computing equilibrium prices p in each period without the use of the pricing rule \hat{p} .
4. Based on the simulated data, update the rules $\hat{K}'^{(l)}$ and $\hat{p}^{(l)}$ to obtain $\hat{K}'^{(l+1)}$ and $\hat{p}^{(l+1)}$.
5. If the rules $\hat{K}'^{(l)}$ and $\hat{p}^{(l)}$ are sufficiently close to the rules $\hat{K}'^{(l+1)}$ and $\hat{p}^{(l+1)}$, stop. If the rules have not converged, return to step 2.

In step 2, we compute $\tilde{V}^{(l)}$ using value function iteration on a discretized state space. For each of the endogenous state variables idiosyncratic capital, idiosyncratic labor, and aggregate capital k, n_{-1} , and K , we assign a log-linear grid with sizes n_k, n_n , and n_K , respectively. Firms assume that the equilibrium price within a period is given exactly by the rule $\hat{p}^{(l)}$ and anticipate next period's value assuming that the forecasting rule $\hat{K}'^{(l)}$ predicts K' exactly. In order to approximate the expectations operator applied over the exogenous stochastic processes for idiosyncratic productivity z , aggregate productivity A , and uncertainty \tilde{S} , we first compute discrete approximations to the productivity processes, following Tauchen (1986). The resulting discretized stochastic processes for z and A vary over grids of size n_z and n_A . Using the transition matrices of the discretized processes for z and A , together with the transition matrix for the two-state Markov process for uncertainty \tilde{S} itself, we compute discrete approximations to the expectations in the firm problem.

In step 3, we simulate the behavior of a large number n_F of firms for a large number of quarters T , computing discrete approximations to integrals over the firm distribution in our theoretical model. In the first period, we initialize firm idiosyncratic capital and labor states to points in the middle of their respective grids, discarding the first T_{erg} periods of our simulation to eliminate the impact of this initialization choice.⁵⁷ In this simulation step 3 it is important to note that we do not use the pricing rule \hat{p} . Instead, given simulated series for uncertainty and aggregate productivity, along with n_F distinct simulated series of idiosyncratic productivity, in each period we solve for the equilibrium consumption price p in that period using golden section search. For each guess of p in the golden section search algorithm, we reoptimize firm investment and hiring policies based on that guess for p and our value function solution $\tilde{V}^{(l)}$. More precisely, given a guess for p , we solve the optimization problem

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

for each firm. Given firm adjustment policies, output, and investment resulting from this optimization problem, we then compute consumption $C(p)$ implied by the goods market clearing condition

$$C(p) = \int \left(y + i - AC^k - AC^n \right) \mu(dkdn_{-1}dz).$$

⁵⁷ We choose the particular values $n_k = 110$ (idiosyncratic capital), $n_n = 45$ (idiosyncratic labor), $n_z = 5$ (idiosyncratic productivity), $n_A = 5$ (aggregate productivity), $n_K = 25$ (aggregate capital), $T = 5,000$ (total number of periods), $T_{erg} = 500$ (number of periods discarded), and $n_F = 5,000$ (number of firms).

A given value of p is the equilibrium price, and the golden section search routine stops, if p satisfies the fixed point condition $p = U_C(C(p), N(p)) = U_C(C(p))^{58}$, or in other words if the goods market clearing value of consumption given p implies a value for the marginal utility of consumption equal to p itself.

In step 4, updating the pricing and forecasting rules to $\hat{K}^{(l+1)}$ and $\hat{p}^{(l+1)}$ based on the simulated data is equivalent to updating the coefficients $\beta^p(A, \tilde{S})$ and $\beta^K(A, \tilde{S})$. For each combination of aggregate productivity and uncertainty values (A, \tilde{S}) on their discrete grids, we collect the subset of periods in which (A, \tilde{S}) was realized. Then, on this subset of periods, we estimate $\beta^p(A, \tilde{S})$ and $\beta^K(A, \tilde{S})$ using ordinary least squares regression.

The converged forecasting and pricing rules \hat{K}' and \hat{p} from the above algorithm are quite successful in explaining aggregate outcomes, with a mean R^2 of 0.99 in the updating regressions for β^K and a mean R^2 of 0.89 in the updating regressions for β^p . Regardless of the conditioning aggregate state (A, \tilde{S}) , the signs of all coefficients are also reasonable: higher values of the aggregate capital stock today are predictive of higher capital tomorrow as well as lower intertemporal prices or marginal utility today.

We use MATLAB to implement the solution technique outlined above. The code used to produce all of the results in this paper, as well as a document explaining the code, is available on Nicholas Bloom's website at the following address: <http://www.stanford.edu/~nbloom/RUBC.zip>.

B.2 Simulating an Uncertainty Shock

Having solved the model following Kahn and Thomas (2008), as laid out in the previous subsection, we compute the responses of this nonlinear model to an uncertainty shock by averaging over the responses to such a shock in a large number of economies, each with a different set of realizations for the driving productivity and uncertainty processes.

As in the solution algorithm, we approximate the continuum of firms in our model with a large but finite number of firms (20,000), at all times computing discrete approximations to the integrals in our theoretical model. Then, we initialize the labor and capital states of each firm to a point in the middle of the idiosyncratic labor and capital grids. Finally, we simulate 400 economies of 25-year, or equivalently 100-quarter, length. In each economy, we impose low uncertainty for 10 periods. In the next quarter, we impose an increase in uncertainty from its low level to its high level, the uncertainty shock itself. Afterwards, in each economy, we allow the duration of the high uncertainty state to be determined by the natural evolution of the two-state Markov process for uncertainty.

For convenience, we serially simulate the 400 distinct economies of 100-quarter length using a string of 40,000 quarters. To eliminate the effects of our initial choice of capital and labor grid points for each firm on our results, we discard the first 25 economies (2500 periods), using the remaining 375 economies to compute the responses discussed in the text. Then, for each period in each economy, we compute all aggregate or cross-sectional quantities of interest, for example output, labor, investment, or labor misallocation, using the distribution of firm states and policies in that period and economy. To compute levels of the series of interest in an "average" economy in a given period, we compute the cross-economy average of that series' level in that period. More concretely, if the value of the series X in period t and simulated economy e is given by X_{et} , the average level upon which our figures are based is given by

$$\bar{X}_t = \frac{1}{375} \sum_{e=1}^{375} X_{et}.$$

The figures in the text further transform the resulting cross-economy average levels \bar{X}_t to percent deviations of these series from their values in period 0, i.e. the period immediately preceding the uncertainty shock. Therefore, we plot \hat{x}_t , where

$$\hat{x}_t = 100 \log(\bar{X}_t / \bar{X}_0).$$

⁵⁸ The second equality holds here because we calibrate $\chi = 1$ to achieve an infinite Frisch elasticity of labor supply.

B.3 Alternative Experiments

We perform two alternative experiments in addition to simulating a baseline uncertainty shock. First, we compute the response of the model to an uncertainty shock coincident with a shock to the level of aggregate productivity. Second, we compute the response of the economy to a wage bill subsidy in normal times, coincident with an uncertainty shock, and several periods after an uncertainty shock. We describe both experiments in turn.

B.3.1 An Uncertainty Shock Coincident with a Productivity Shock

In order to compute the response of the model to an uncertainty shock occurring in the same quarter as an adverse aggregate productivity shock, we simulate a large number of economies with different realizations of the model's exogenous processes, averaging over their response to coincident uncertainty and productivity shocks. The number of firms (20,000), total number of economies (400), periods per economy (100), number of initial discarded economies (25), and timing conventions are identical to those used to simulate a baseline uncertainty shock, with the exception that we now also impose an aggregate productivity shock as well.

We construct the productivity shock to yield an average reduction in aggregate productivity across simulated economies of 2%. Given our discretized grid $\tilde{A} = (\tilde{A}_1, \dots, \tilde{A}_{n_A})$ for exogenous aggregate productivity, such a productivity shock is not in general equivalent to a deterministic reduction of aggregate productivity by a fixed number of grid points in each economy. Instead, we convexify the drop in \tilde{A} in period 1, the period of the joint uncertainty and productivity shock, by randomly imposing a drop in aggregate productivity to its lowest grid point. More precisely, for each economy e in period 1, we draw both a uniformly distributed shock $\xi_e \in (0, 1)$ along with an unrestricted candidate value $A_e^* \in \tilde{A}$ for A_{e1} drawn from the unrestricted discretized Markov process for A . We set

$$A_{e1} = \begin{cases} \tilde{A}_1, & \text{if } \xi < \bar{\xi} \\ A_e^*, & \text{if } \xi \geq \bar{\xi} \end{cases},$$

where $\bar{\xi}$ was chosen to guarantee that the cross-economy aggregate productivity average declines by 2%. Equivalently, $\bar{\xi}$ was chosen to guarantee that

$$\frac{1}{375} \sum_{e=1}^{375} \log(A_{e1}/A_{eo}) = -0.02.$$

Since the probability of a drop in the aggregate productivity process to its lowest grid point is increasing in $\bar{\xi}$, the cross-economy average above is decreasing in $\bar{\xi}$. Based on this relationship, we determine $\bar{\xi}$ by iteratively guessing values of $\bar{\xi}$ until the equation above is satisfied.

B.3.2 A Wage Bill Subsidy Policy Experiment

We consider the simulated impact on output of a stimulative wage bill subsidy of one-quarter duration in three context: in normal times, concurrent with an uncertainty shock, and after an uncertainty shock. The number of firms (20,000), total number of economies (400), periods per economy (100), and number of initial discarded economies (25), are identical to those used to simulate the impact of a baseline uncertainty shock.

A Wage Bill Subsidy in Normal Times: To compute the impact of a wage bill subsidy in normal, i.e. low uncertainty, times, we first simulate exogenous aggregate and idiosyncratic productivity shocks for all 400 economies, imposing low uncertainty throughout. Then in period 1 for each economy, the same quarter in which we also impose an uncertainty shock in alternative simulations, we implement a wage bill subsidy of 1%. In period 1 firms face marginal wage rates equal to (100-1)% of the equilibrium wage rate, with the subsidy financed through a lump-sum tax on households. For aggregate output Y , we compute the percent deviation each period from the quarter preceding the subsidy of the level of the relevant series averaged across economies. For later reference, we refer to this series of percent output deviations as $\hat{y}_{sub,t}$.

A Wage Bill Subsidy Concurrent with an Uncertainty Shock: To compute the impact of a wage bill subsidy contemporaneous with an uncertainty shock, we first simulate the aggregate productivity, idiosyncratic productivity, and uncertainty processes exactly as in the case of a baseline uncertainty shock. That is, we impose low uncertainty for 10 periods, then in a quarter labeled period 1 for convenience we shock the economy with a transition to high uncertainty, afterwards allowing the level of uncertainty to vary according to its calibrated Markov chain. Productivity series at the aggregate and idiosyncratic levels evolve naturally, given the series for uncertainty. In period 1 we also implement a wage bill subsidy of 1%. For aggregate output Y , we define $\hat{y}_{unc+currsub,t}$ to be the percent deviation from period 0 of the average economy in period t , given the concurrent uncertainty shock and wage bill subsidy in period 1.

A Wage Bill Subsidy after an Uncertainty Shock: To compute the impact of a wage bill subsidy after, as opposed to concurrent with, an uncertainty shock, we again simulate all exogenous processes exactly as in the case of a baseline uncertainty shock. This uncertainty shock occurs in period 1. Then, in period 3, we implement a 1% wage bill subsidy. For aggregate output Y , we define $\hat{y}_{unc+aftersub,t}$ to be the percentage deviation from period 0 of the average economy in period t , given the uncertainty shock in period 1 and the wage bill subsidy in period 3.

After computing each of the three responses $\hat{y}_{sub,t}$, $\hat{y}_{unc+currsub,t}$, and $\hat{y}_{unc+aftersub,t}$ described above, together with the model's output response to a baseline uncertainty shock \hat{y}_t , we plot three series: $\hat{y}_{sub,t}$ (the impact of a wage bill subsidy in normal times), $\hat{y}_{unc+currsub,t} - \hat{y}_t$ (the impact of a wage bill subsidy with an uncertainty shock), and $\hat{y}_{unc+aftersub,t} - \hat{y}_t$ (the impact of a wage bill subsidy after an uncertainty shock).

C Online Appendix: Measurement Error Calculations

In this appendix we show how we use OLS and IV estimates of the AR coefficients in the TFP forecast regressions to calculate the share of measurement error in $\log(\text{TFP})$. Suppose that (4) is observed with error (omitting the j subscripts)

$$\log(\hat{z}_t^*) = \log(\hat{z}_t) + e_t.$$

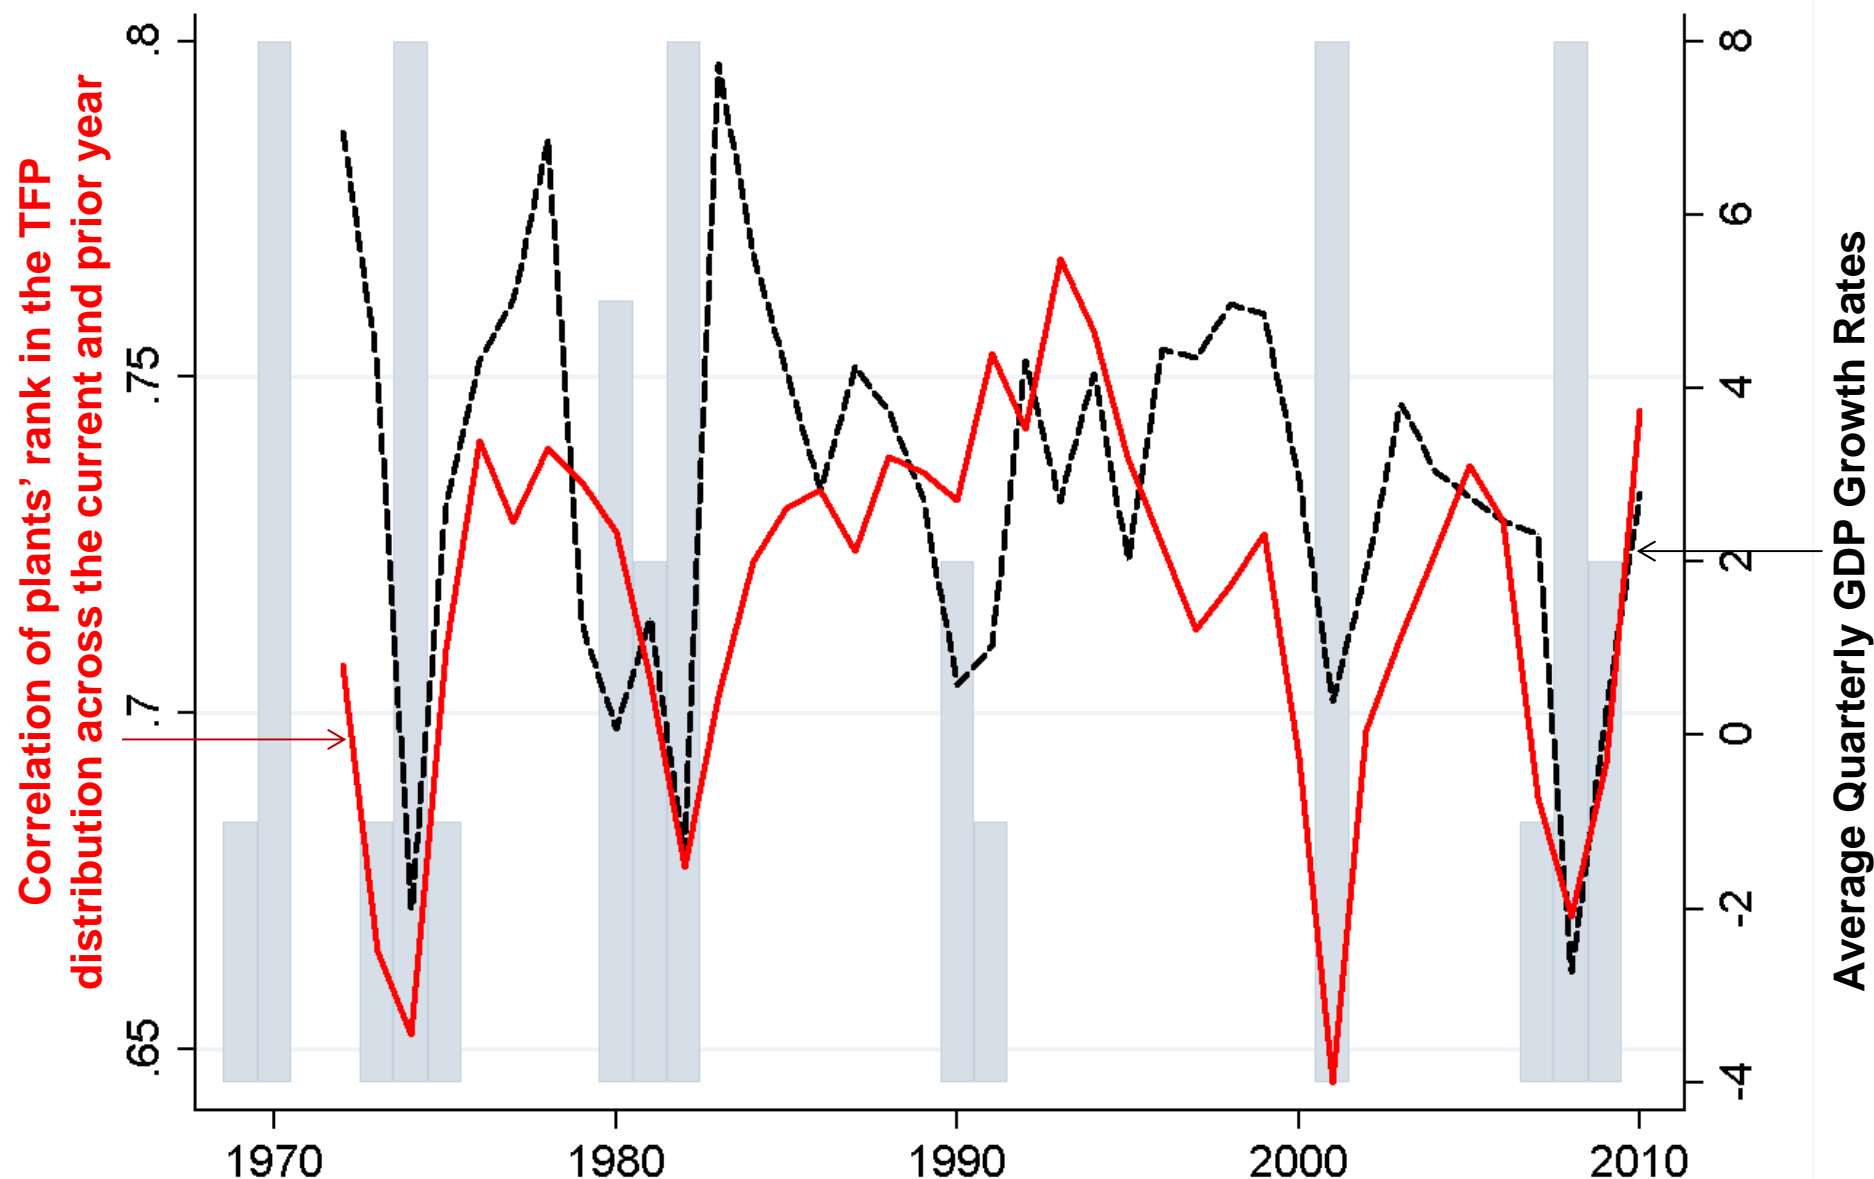
If the measurement (e_t) error is i.i.d, estimating (4) using $\log(\hat{z}_{t-2}^*)$ to instrument for $\log(\hat{z}_{t-1}^*)$ we would obtain a consistent estimate for ρ (Call this estimate ρ_{IV1}). The OLS estimate for (4) is inconsistent but the bias is a function of the measurement error in TFP:

$$\rho_{OLS} = \frac{1}{\rho_{IV1}} \left(1 + \frac{\sigma_e^2}{\text{var}(\log(\hat{z}_t))} \right)$$

We therefore use ρ_{OLS} together with ρ_{IV1} to obtain an estimate for $\sqrt{\frac{\sigma_e^2}{\text{var}(\log(\hat{z}_t))}}$, which is the share of measurement error in total standard deviation of TFP. The results for ρ_{OLS} and ρ_{IV1} are reported in columns (1) and (2) of Table A2 respectively. These estimates yield a measurement error share of 37.4%.

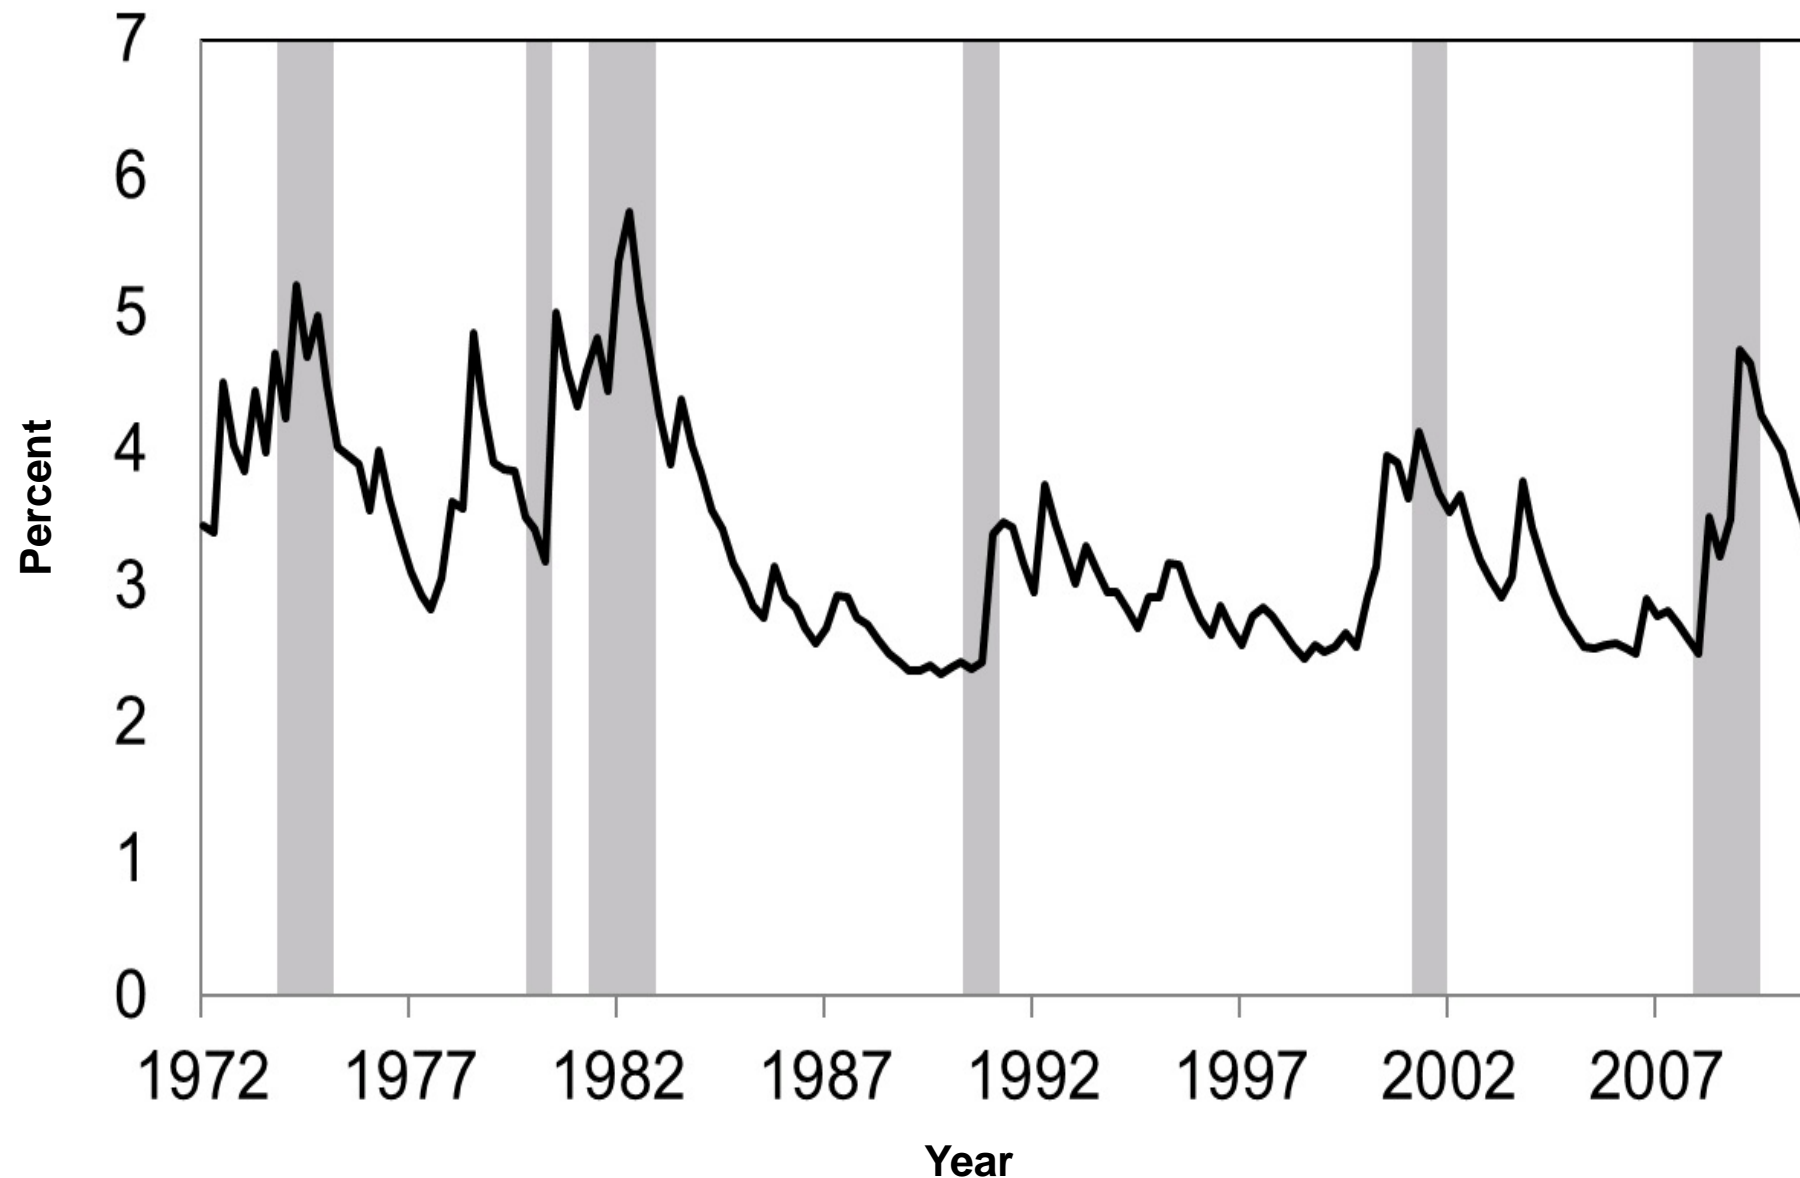
Suppose now that there is some serial correlation in measurement error and that we can write $\text{cov}(e_t, e_{t-1}) = \sigma_{t,t-1}$. As before, define ρ_{IV1} to be the estimate for ρ from an IV regression where $\log(\hat{z}_{t-2}^*)$ is used to instrument for $\log(\hat{z}_{t-1}^*)$. Define ρ_{IV2} to be the estimate for ρ from an IV regression where $\log(\hat{z}_{t-3}^*)$ is used to instrument for $\log(\hat{z}_{t-1}^*)$. Then we can combine the estimates for ρ_{OLS} with ρ_{IV1} and ρ_{IV2} to obtain estimates for measurement error share as well as for $\sigma_{t,t-1}$. For this specification, our estimates imply measurement error share are of 45.4%.

Figure A2: Recessions increase turbulence: plant rankings in the TFP distribution churn more in recessions



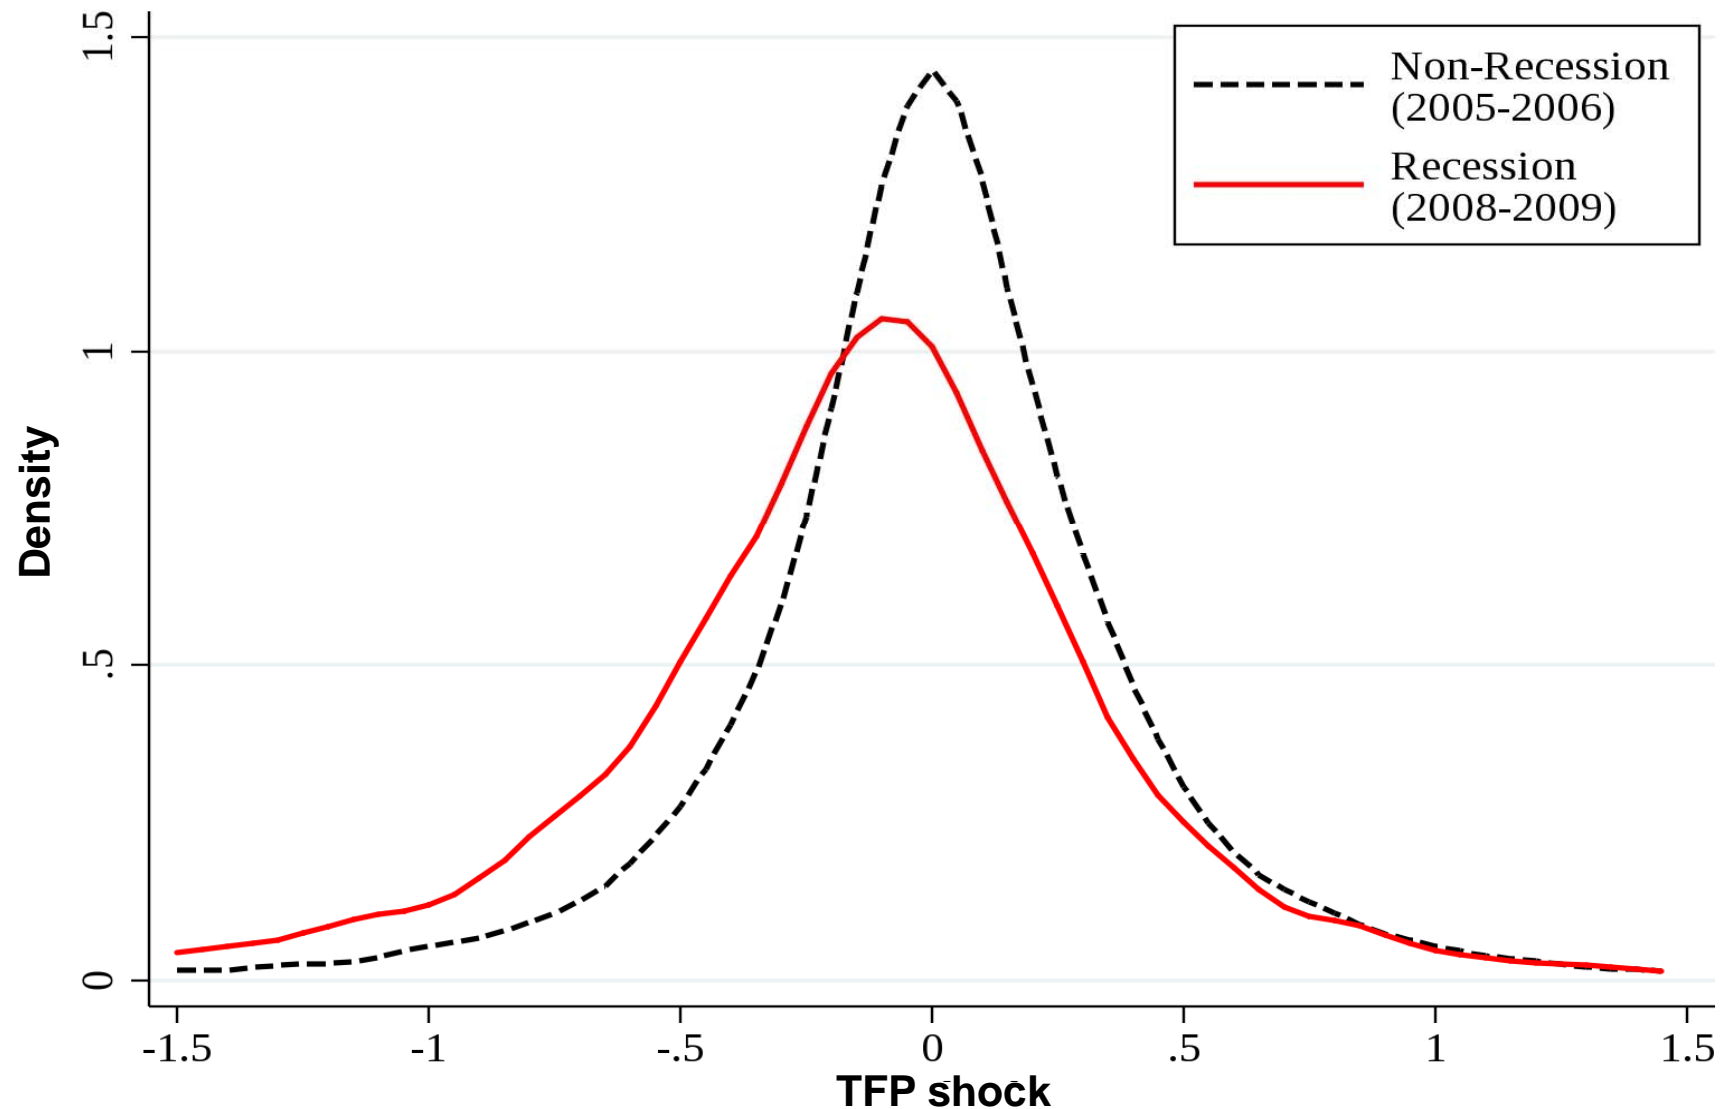
Notes: Constructed from the Census of Manufacturers and the Annual Survey of Manufacturing establishments, using establishments with 25+ years to address sample selection. Grey shaded columns are share of quarters in recession within a year. Plants' rank in the TFP distribution is their decile within the industry and year TFP ranking.

Figure A2: TFP growth GARCH (1,1) volatility



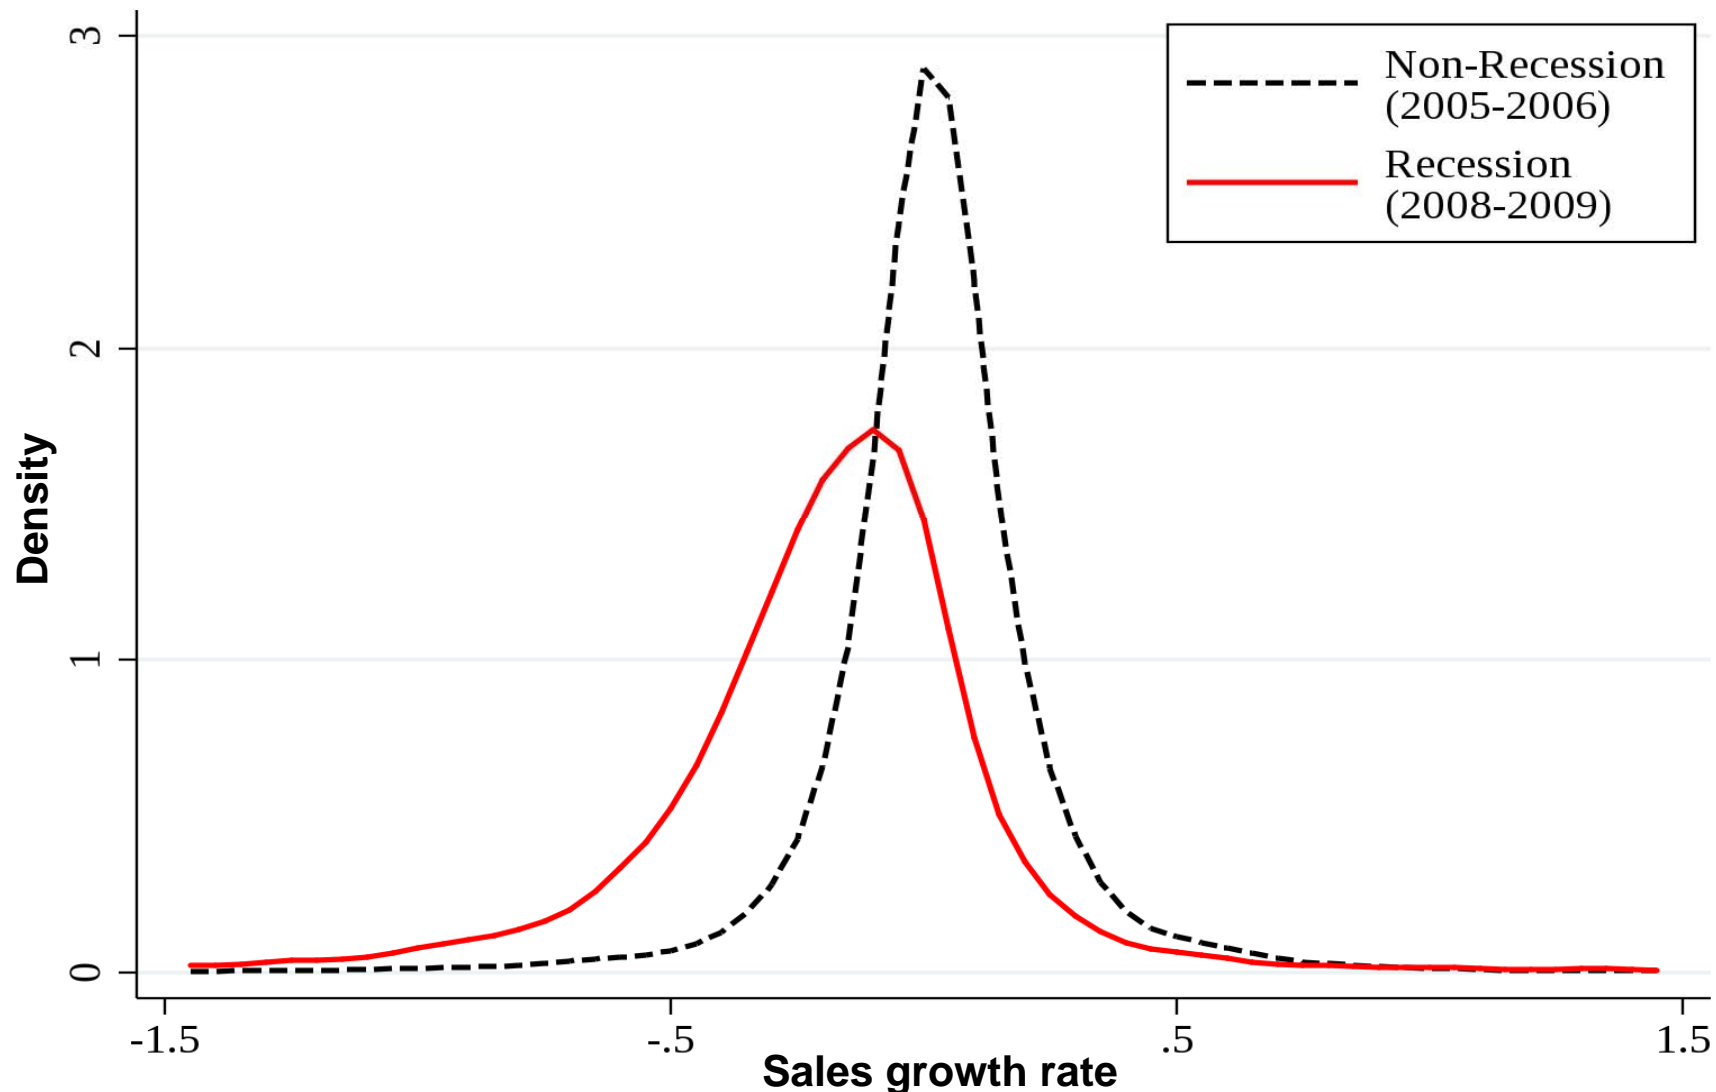
Notes: The conditional heteroskedasticity series above is estimated using a GARCH(1,1) model for the change in the aggregate Solow residual in US quarterly data from 1972-2010, as available on John Fernald's website on May 17, 2012 (series dTFP). The recession bars refer to standard NBER dates.

Figure 1: The variance of establishment-level TFP shocks increased by 76% in the Great Recession



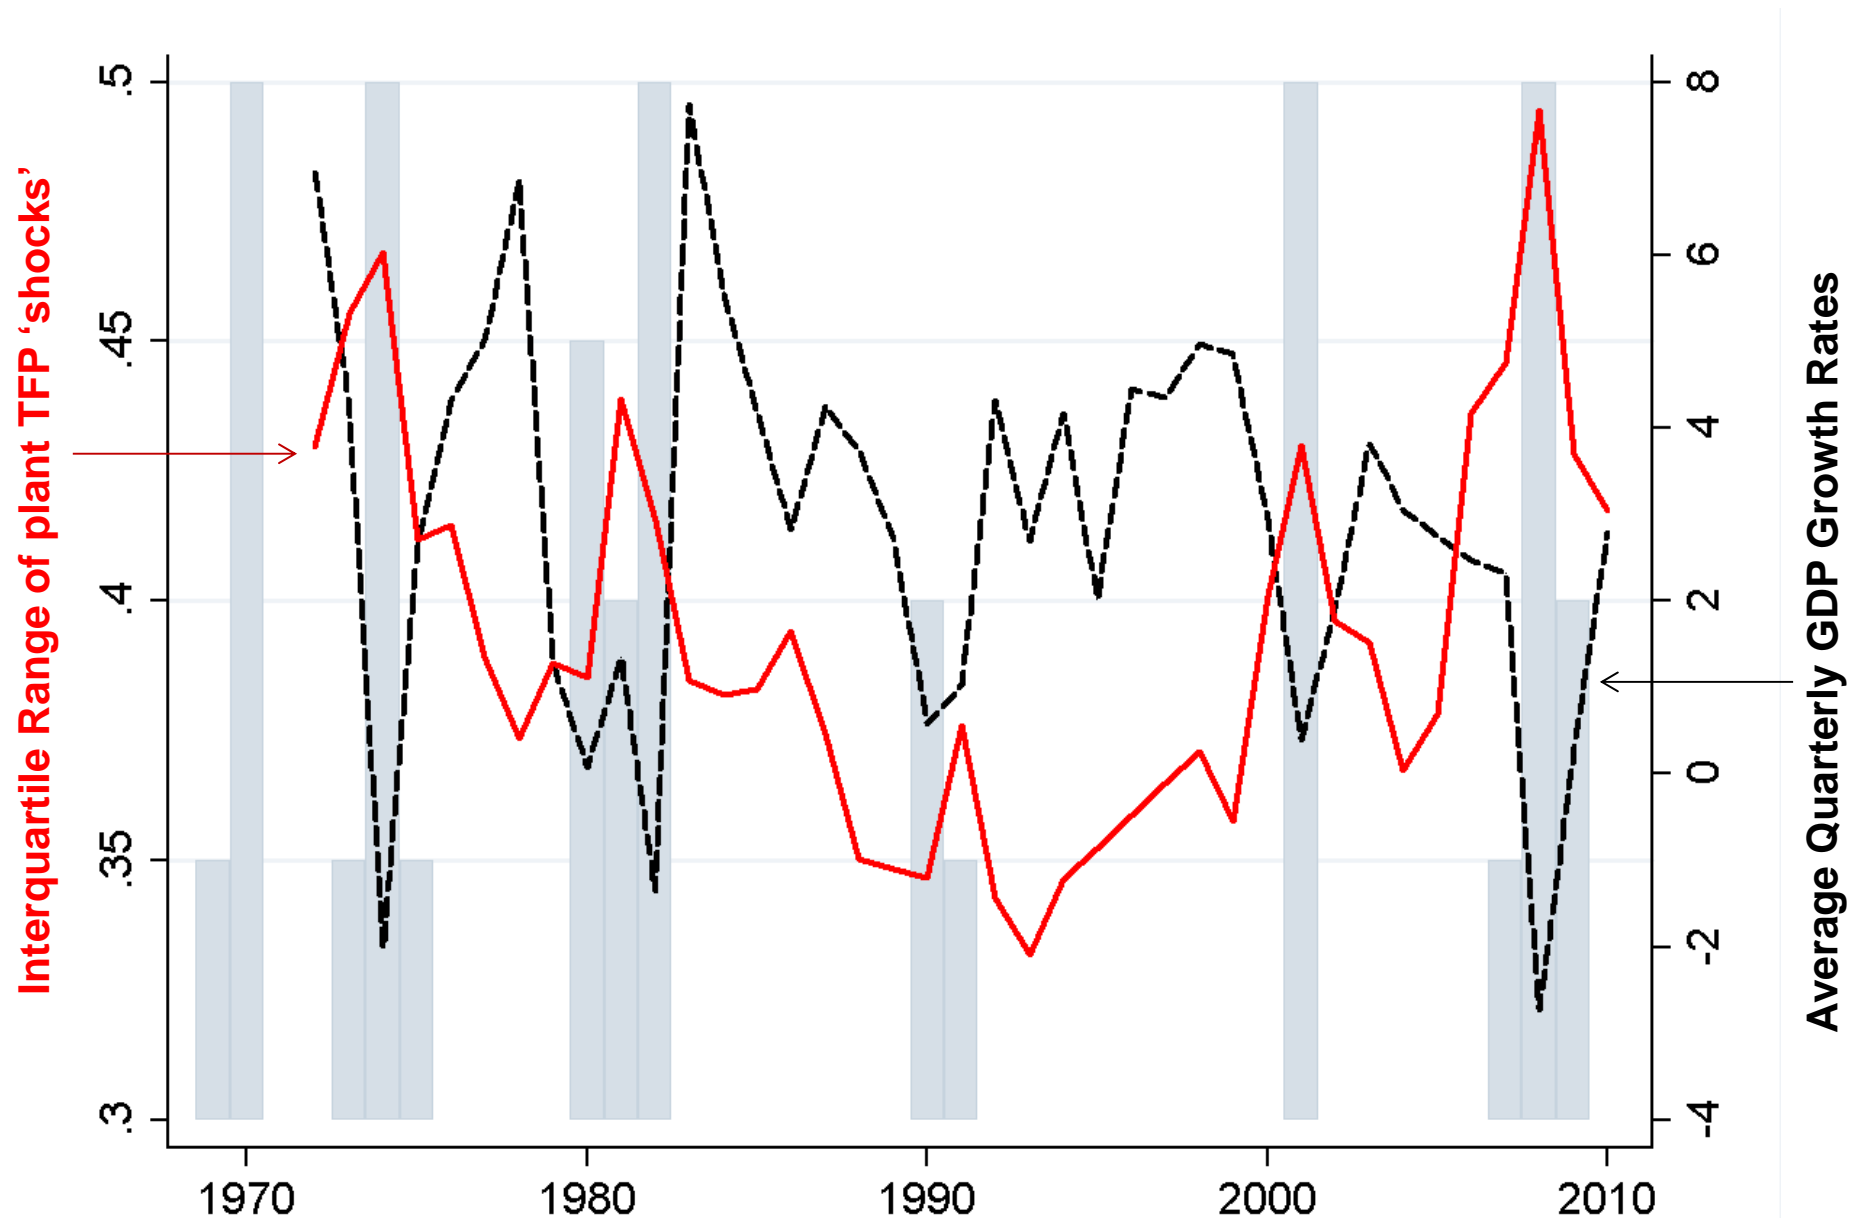
Notes: Constructed from the Census of Manufactures and the Annual Survey of Manufactures using a balanced panel of 15,752 establishments active in 2005-06 and 2008-09. Moments of the distribution for non-recession (recession) years are: mean 0 (-0.166), variance 0.198 (0.349), coefficient of skewness -1.060 (-1.340) and kurtosis 15.01 (11.96). The year 2007 is omitted because according to the NBER the recession began in December 2007, so 2007 is not a clean “before” or “during” recession year.

Figure 2: The variance of establishment-level sales growth rates increased by 152% in the Great Recession



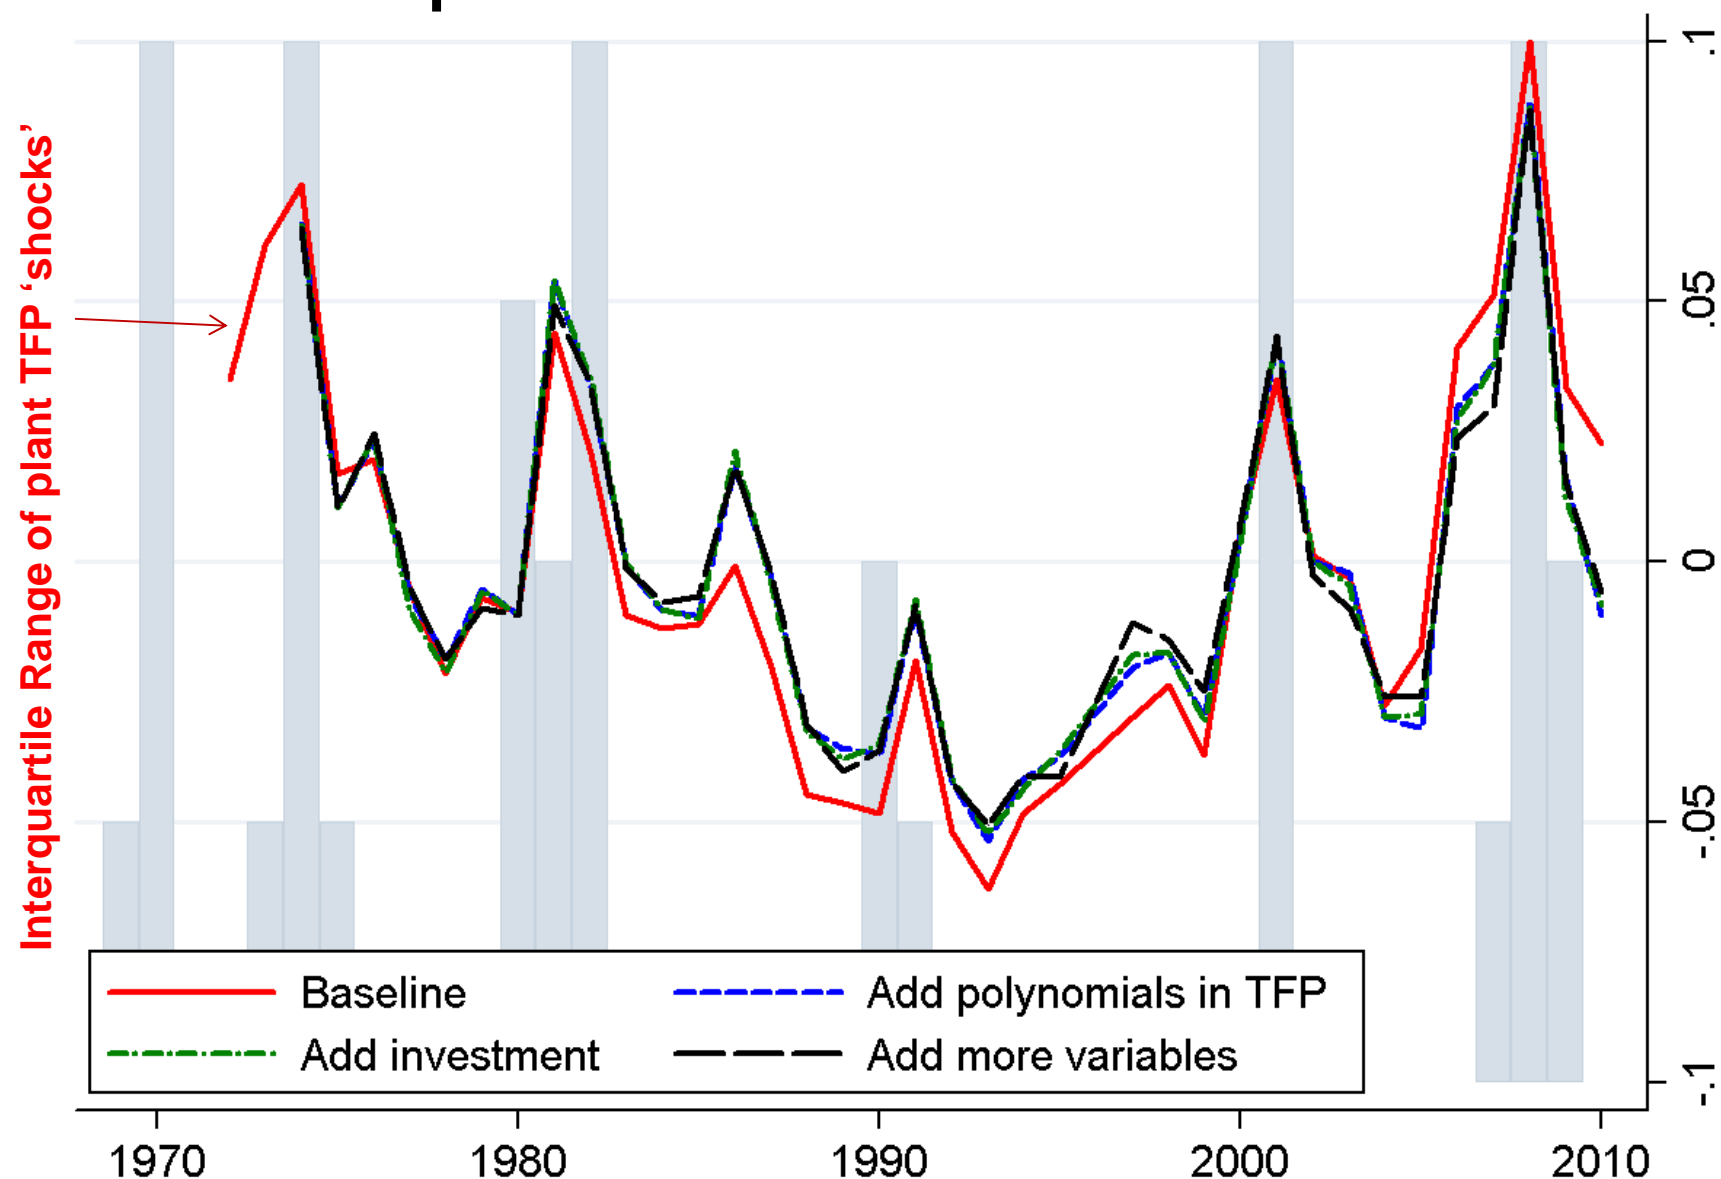
Notes: Constructed from the Census of Manufactures and the Annual Survey of Manufactures using a balanced panel of 15,752 establishments active in 2005-06 and 2008-09. Moments of the distribution for non-recession (recession) years are: mean 0.026 (-0.191), variance 0.052 (0.131), coefficient of skewness 0.164 (-0.330) and kurtosis 13.07 (7.66). The year 2007 is omitted because according to the NBER the recession began in December 2007, so 2007 is not a clean “before” or “during” recession year.

Figure 3: TFP ‘shocks’ are more dispersed in recessions



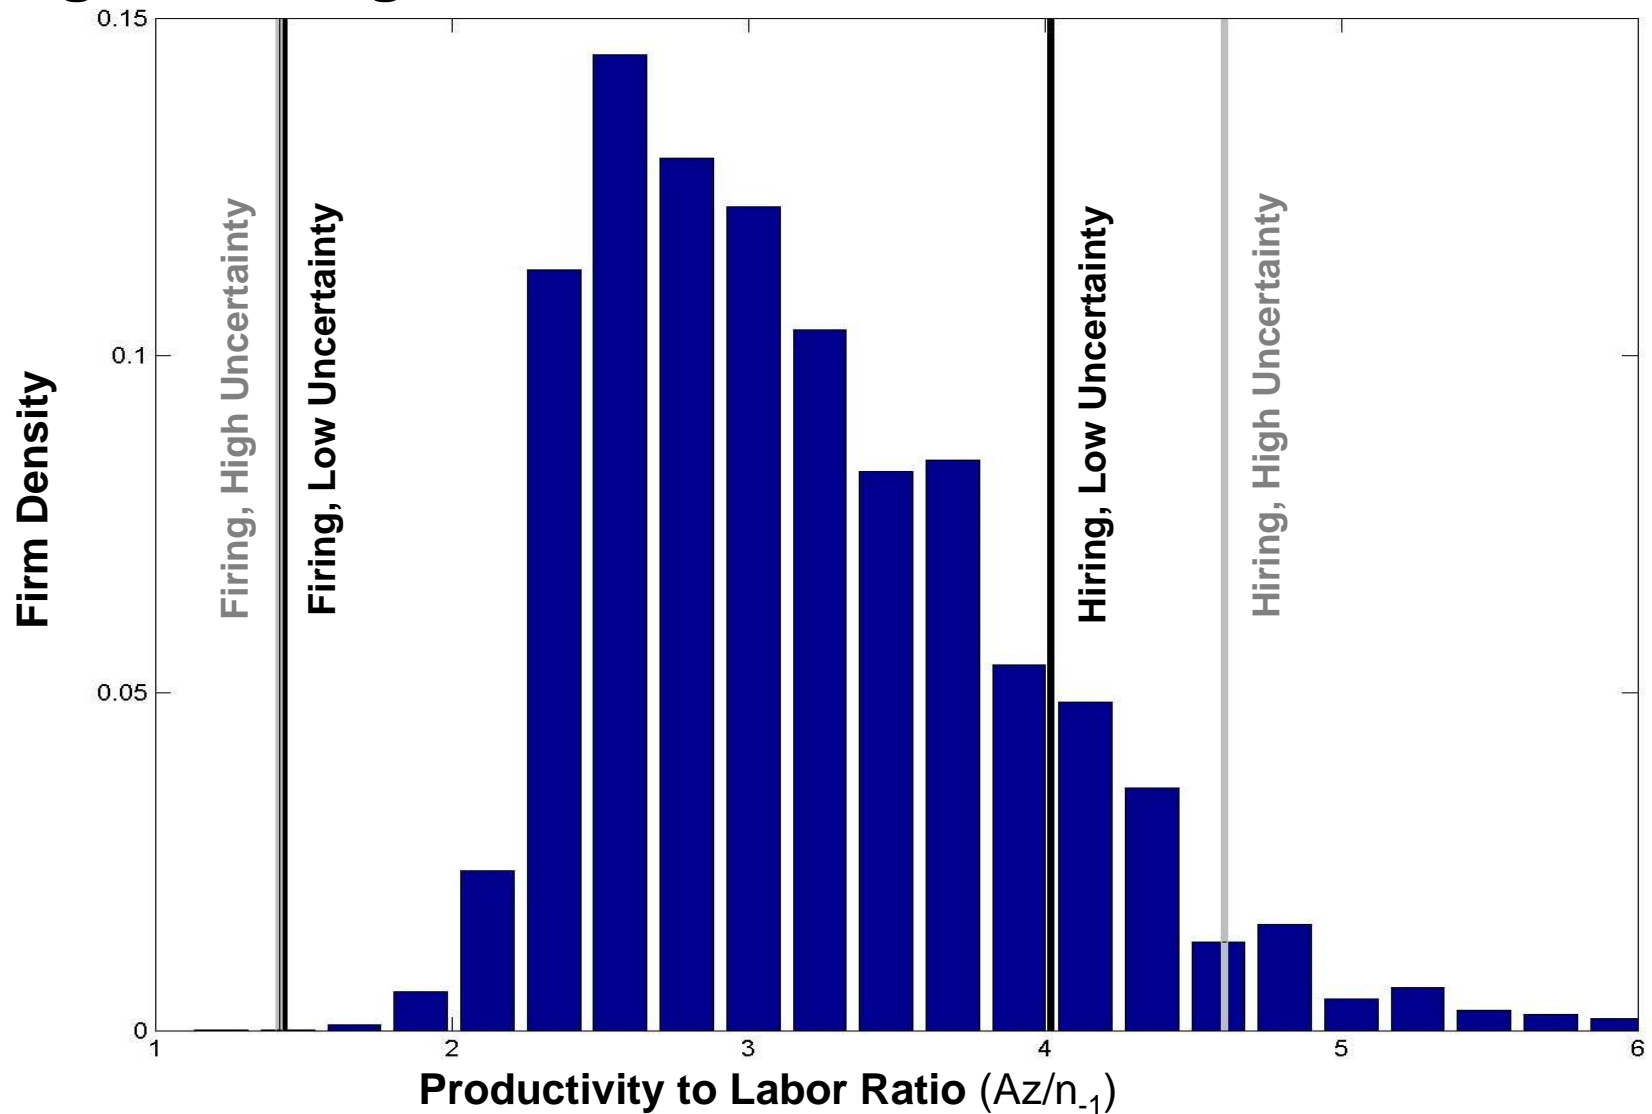
Notes: Constructed from the Census of Manufactures and the Annual Survey of Manufactures establishments, using establishments with 25+ years to address sample selection. Grey shaded columns are share of quarters in recession within a year.

Figure 4: Robustness test: different measures of TFP ‘shocks’ are all more dispersed in recessions



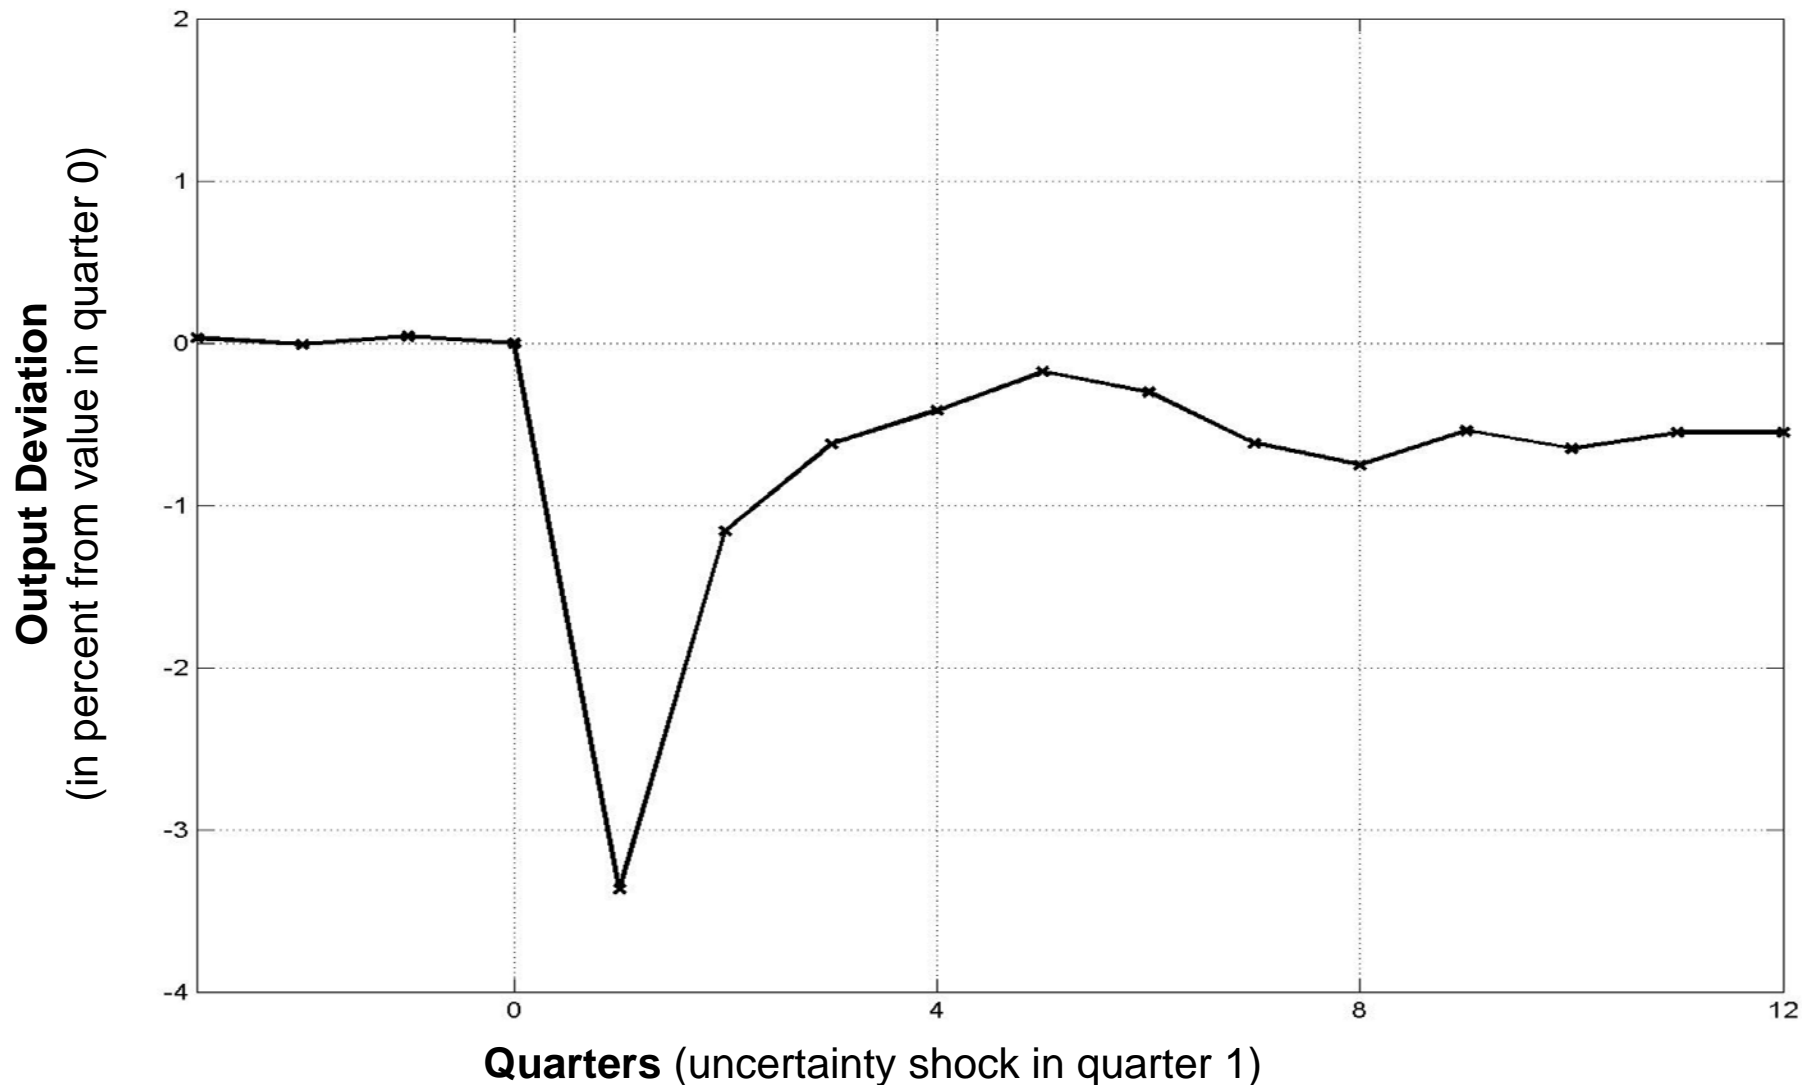
Notes: Constructed from the Census of Manufactures and the Annual Survey of Manufactures establishments, using establishments with 25+ years to address sample selection. Grey shaded columns are share of quarters in recession within a year.

Figure 5: The impact of an increase in uncertainty on the hiring and firing thresholds



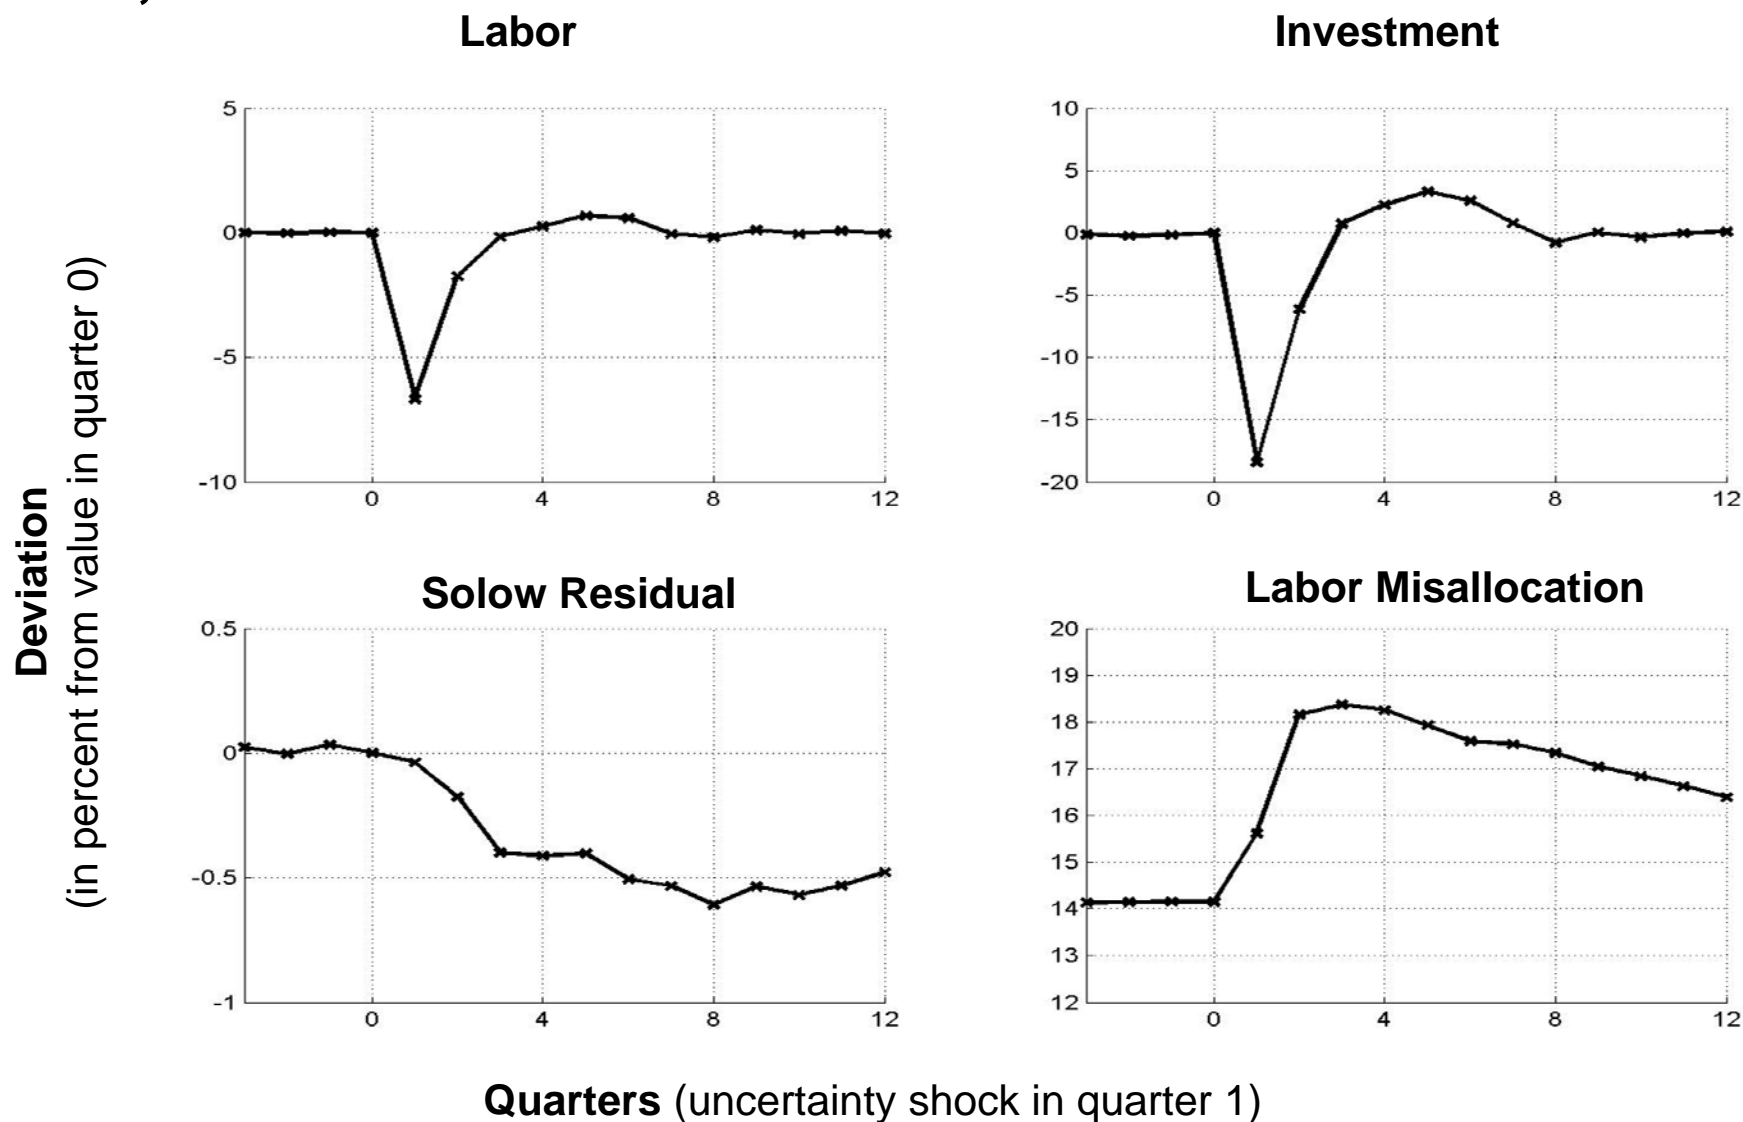
Notes: To produce the thresholds and histogram we first simulate 400 economies of 100-quarter length with 20,000 firms, imposing one uncertainty shock in each economy. Then, we form a sample of all firms with idiosyncratic capital within 3 grid points of the modal idiosyncratic capital value. Then, we form output-weighted mean hiring and firing thresholds before and during uncertainty shocks. Finally, we plot a histogram of productivity/labor ratios in our sample of firms.

Figure 6: An uncertainty shock causes an output drop of just over 3% and a recovery to almost level within 1 year



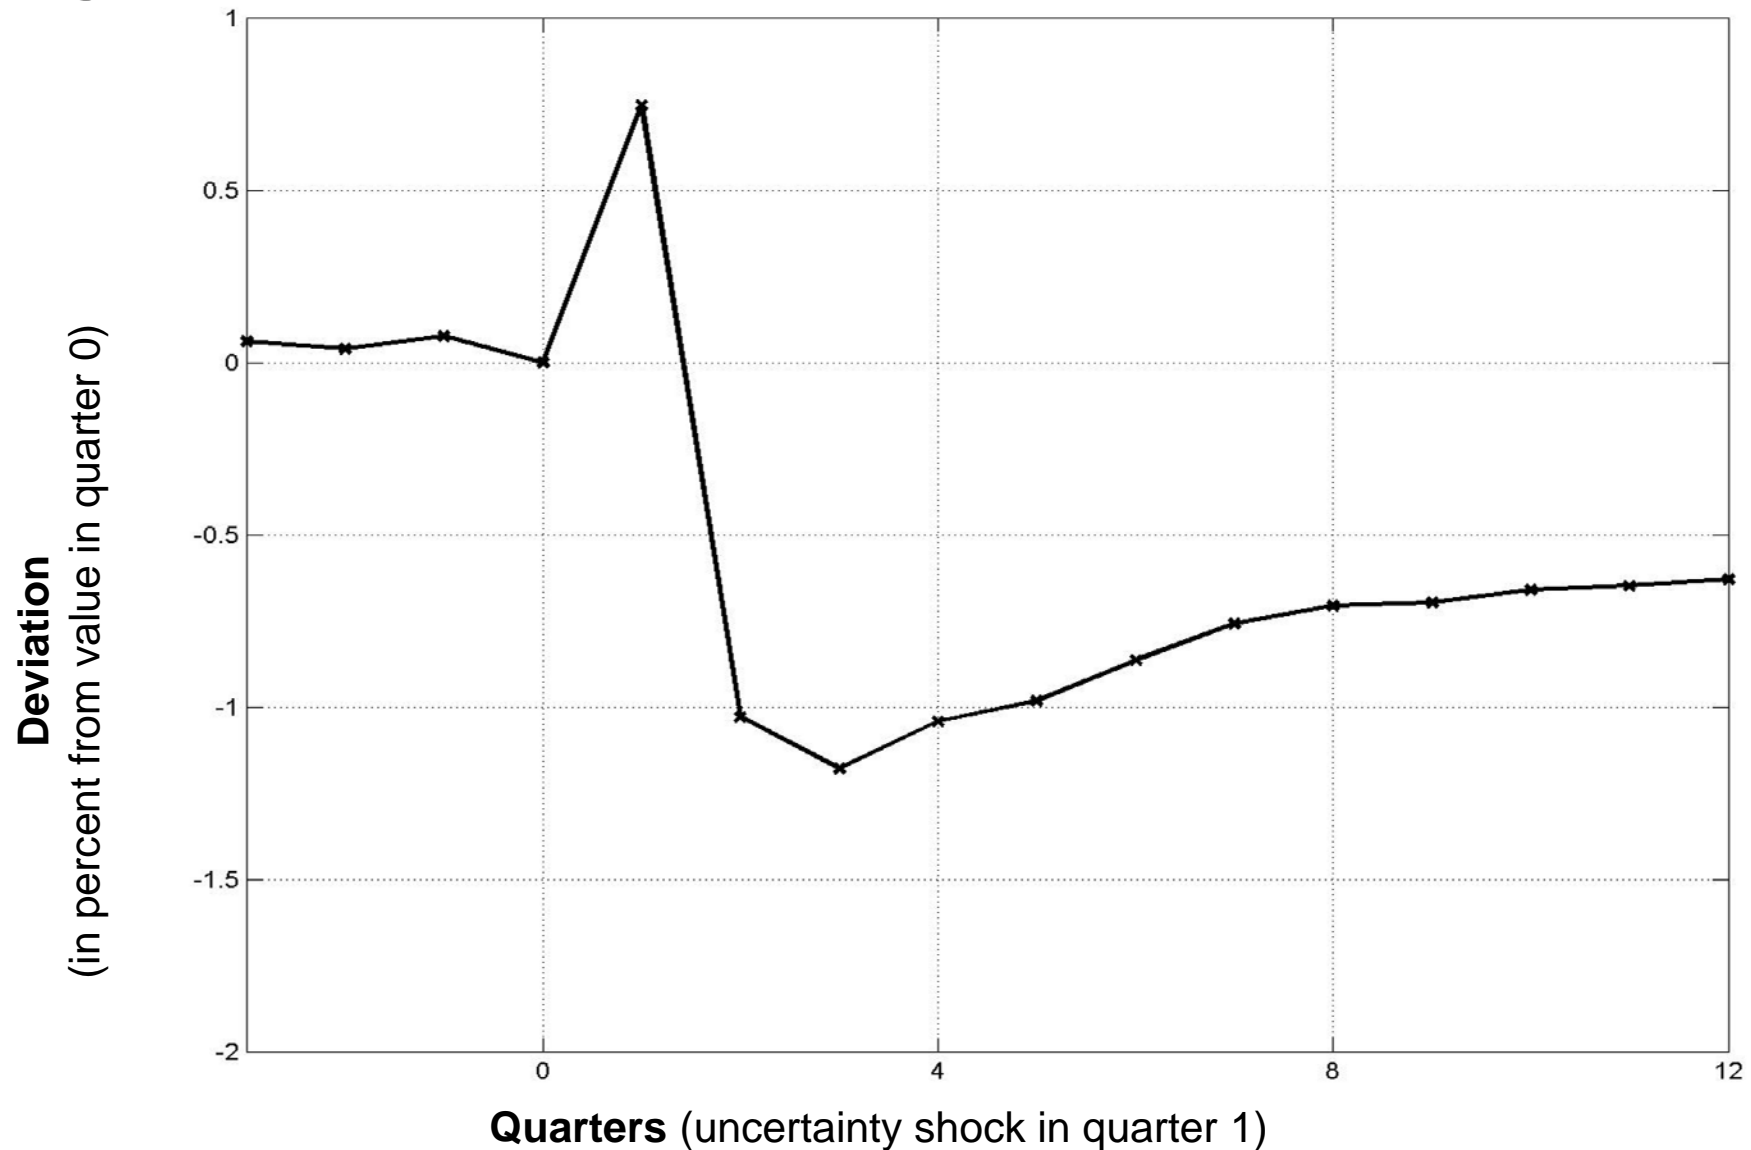
Notes: Results simulated for 400 separate economies with 20,000 firms each. We first impose an uncertainty shock in quarter 1, afterwards allowing the uncertainty process to evolve naturally. Then we compute the resulting average aggregate output level across economies for each quarter. Finally, we plot percent deviation of the cross-economy average from its value in quarter 0. Small fluctuations from period 5 onwards are due to echo effects arising from lumpy adjustment costs.

Figure 7: Labor and investment drop and rebound, TFP slowly falls, and misallocation rises



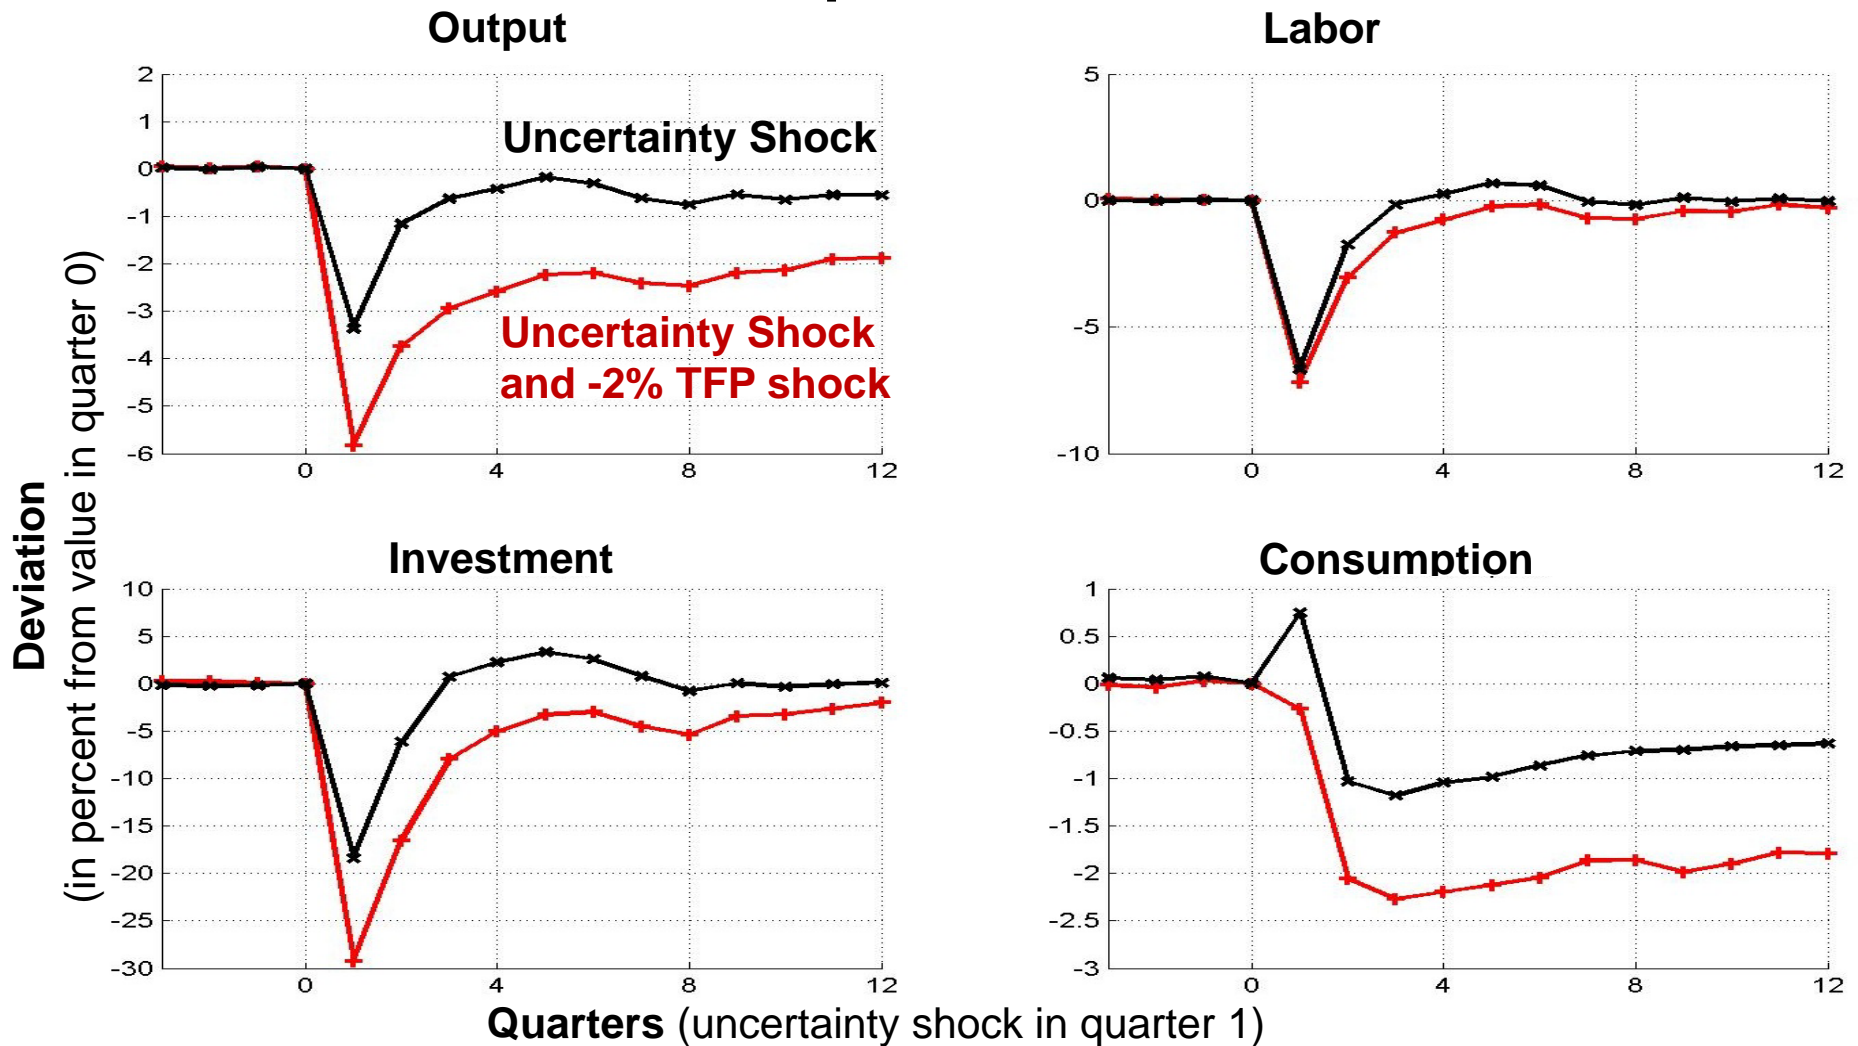
Notes: Results simulated for 400 separate economies with 20,000 firms each. We first impose an uncertainty shock in quarter 1, afterwards allowing the uncertainty process to evolve naturally. Then we compute the resulting average aggregate series level across economies for each quarter. Finally, for labor, investment, and the Solow Residual, we plot percent deviation of the cross-economy average from its value in quarter 0. For labor misallocation we simply plot the cross-economy average in percent.

Figure 8: Consumption overshoots and then drops



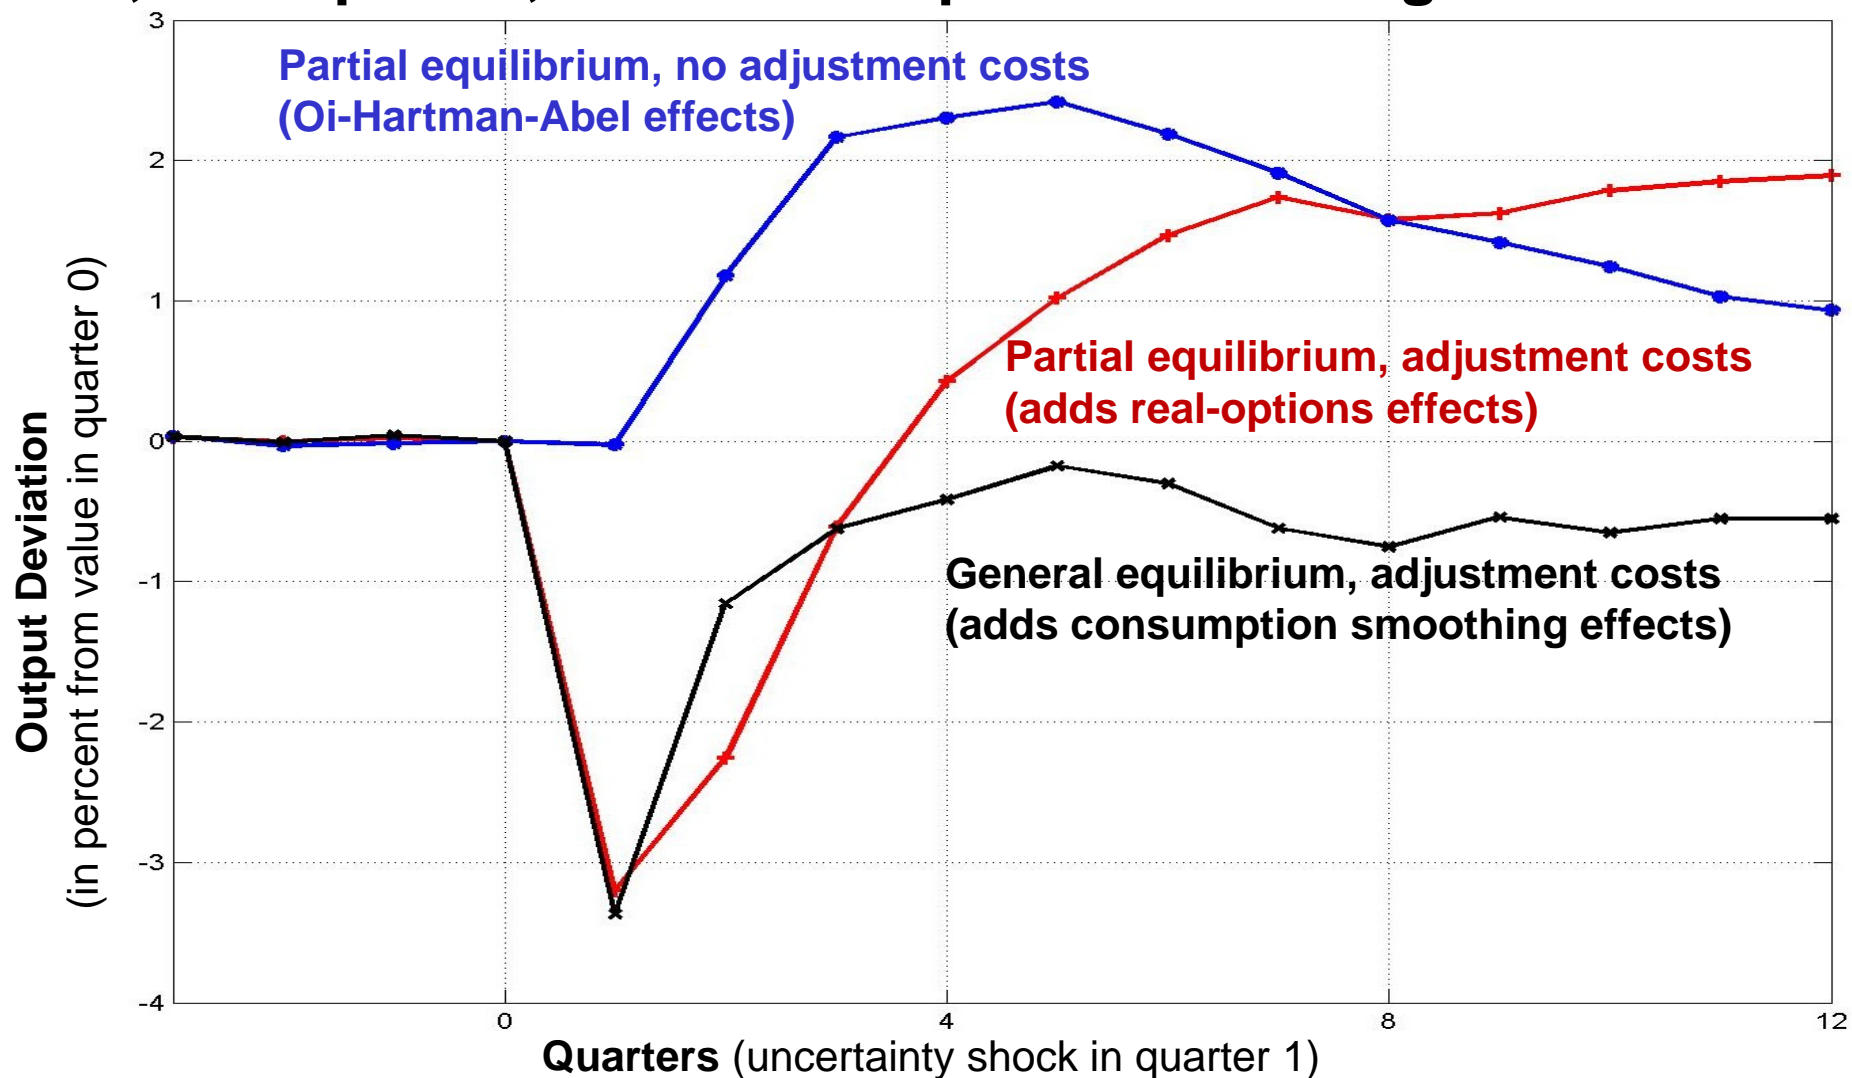
Notes: Results simulated for 400 separate economies with 20,000 firms each. We first impose an uncertainty shock in quarter 1, afterwards allowing the uncertainty process to evolve naturally. Then we compute the resulting average aggregate consumption level across economies for each quarter. Finally, we plot percent deviation of the cross-economy average from its value in quarter 0.

Figure 9: Adding a -2% first-moment shock increases the drop and eliminates the consumption overshoot



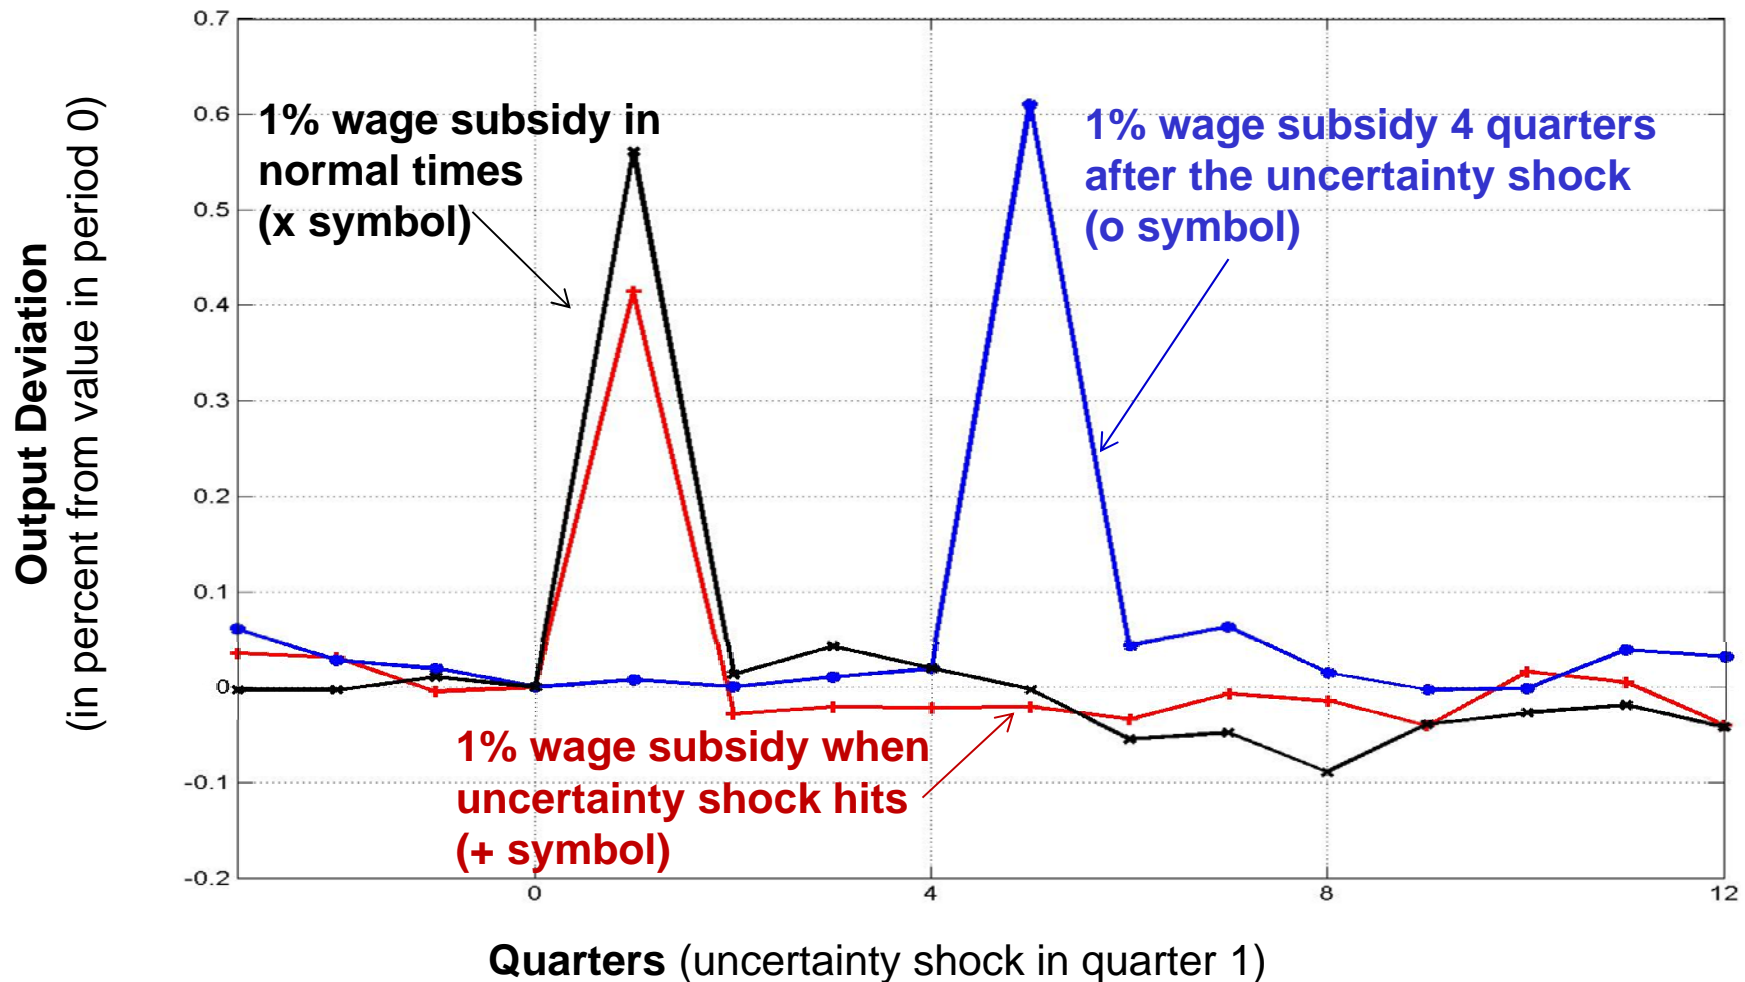
Notes: Simulated for 400 separate economies with 20,000 firms each. We first impose a shock in quarter 1, afterwards allowing the uncertainty process to evolve naturally. Then we compute the resulting average aggregate series level across economies for each quarter. Finally, we plot percent deviation of the cross-economy average from its value in quarter 0. For the baseline uncertainty shock in quarter 1 we increase the level of uncertainty (x symbols). For the joint uncertainty and negative TFP shock (+ symbols), in quarter 1 we proceed as in the baseline case but also impose a negative shock to aggregate productivity with cross-economy average equal to -2%.

Figure 10: Splitting the uncertainty impact into Oi-Hartman-Abel, real options, and consumption smoothing effects



Notes: Results simulated for 400 separate economies with 20,000 firms each. We first impose an uncertainty shock in quarter 1, afterwards allowing the uncertainty process to evolve naturally. Then we compute the resulting average aggregate output level across economies for each quarter. Finally, we plot percent deviation of the cross-economy average from its value in quarter 0. The general equilibrium response is the baseline simulation. Partial equilibrium simulations impose fixed output prices and consumption, with or without adjustment costs for capital and labor inputs.

Figure 11: Policy is less effective immediately after an uncertainty shock but more effective four quarters later



Notes: Simulated for 400 separate economies with 20,000 firms each. For a subsidy in normal times, we allow all exogenous processes to evolve naturally and provide a one-quarter 1% wage bill subsidy in quarter 1, plotting the percent deviation of the average output level in a given period from quarter 0. In normal times, the choice of quarter for a subsidy does not affect the output response. For a subsidy in quarter x paired with an uncertainty shock in quarter 1, we impose an uncertainty shock in quarter 1 together with a wage bill subsidy of 1% in quarter x. We plot the difference between the percent deviation of average output from quarter 0 in the wage bill and subsidy case and the percent deviation of average output from quarter 0 in the case of an uncertainty shock in quarter 1 with no wage bill subsidy.