

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). 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 2011 (using the LBDNUM identifier).⁸ 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.⁹ 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 CUSIP8¹⁰.

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,

⁸The 2010 and 2011 ASM data became available only at late stages of the project. To avoid repeating extensive census disclosure analysis, in Tables 2 and 3 we use data only up to 2009. The 2011 data became available only recently, therefore the SMM estimation uses data only up to 2010.

⁹The access to CRSP and Compustat data sets is through WRDS: <https://wrds-web.wharton.upenn.edu/wrds/>.

¹⁰We 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).¹¹ 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 1997.¹² 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.¹³ 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.¹⁴ 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.¹⁵ 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.

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.¹⁶

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.

¹¹See: <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.

¹²The NBER-CES data are available only through 2009. 2010 industry-level data are therefore projected using an AR(4) regression for all external datasets.

¹³See <http://www.bls.gov/mfp/mprdload.htm>.

¹⁴See <http://www.bea.gov/national/FA2004/SelectTable.asp>.

¹⁵See <http://www.federalreserve.gov/releases/G17/Current/default.htm>.

¹⁶As 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.

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 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 $K_t = (1 - \delta_t) K_{t-1} + I_t$.¹⁷ 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

¹⁷The stock is reported for the end of period, therefore we use last period's stock with this period's depreciation and investment.

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

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^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:

$$\log(\hat{z}_{j,i,t}) = \rho \log(\hat{z}_{j,i,t-1}) + \mu_i + \lambda_t + e_{j,i,t}, \quad (21)$$

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 (21) used to back out the shocks. In particular we report the dispersion of TFP shocks, which were calculated by running (21) 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: 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 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.

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 (21)), 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.

B Online Appendix: Model Solution and Simulation

In this computational appendix we first lay out the broad solution algorithm for the model, which follows the generalized Krusell and Smith (1998) approach as laid out by Khan and Thomas (2008). We also provide information on the practical numerical choices made in the implementation of that solution method. Then, we discuss the calculation of the approximate impulse response or effects of an uncertainty shock in this economy, together with the various other experiments analyzed in the main text. We conclude with a discussion of the accuracy of the forecast rule used to predict the evolution of the approximate aggregate state in the economy as well as other related solution techniques for this class of models.

As we discuss below (see the section titled "Alternative Forecasting or Market-Clearing Assumptions") we consider two alternatives to the basic forecasting rule discussed below. For each of these methods we resolve the GE model until convergence of the forecasting rule is achieved. Importantly, as we discuss below the main results are robust across the different solution methods.

B.1 Solution Algorithm

The marginal-utility transformed Bellman equation describing the firm problem in the main text is reproduced below and results in a convenient problem with constant discounting at rate β but a transformation of period dividends by a marginal utility price p :

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

We can see that the aggregate state of the economy is $(A, \sigma^A, \sigma^Z, \mu)$ and as outlined in the main text household optimization implies the relationships

$$\begin{aligned} p(A, \sigma^A, \sigma^Z, \mu) &= C(A, \sigma^A, \sigma^Z, \mu)^{-\eta} \\ w(A, \sigma^A, \sigma^Z, \mu) &= \theta N(A, \sigma^A, \sigma^Z, \mu)^{\chi-1} C(A, \sigma^A, \sigma^Z, \mu)^\eta. \end{aligned}$$

The calibration we choose, with log consumption utility ($\eta = 1$) and an infinite Frisch elasticity of substitution ($\chi = 1$), implies that the equilibrium relationships are given by

$$\begin{aligned} p(A, \sigma^A, \sigma^Z, \mu) &= C(A, \sigma^A, \sigma^Z, \mu)^{-1} \\ w(A, \sigma^A, \sigma^Z, \mu) &= \theta C(A, \sigma^A, \sigma^Z, \mu) = \frac{\theta}{p(A, \sigma^A, \sigma^Z, \mu)}. \end{aligned}$$

With these choices wages w are a function of the marginal utility price p , so the evolution of the aggregate equilibrium of the economy is fully characterized by the following two mappings:

$$\begin{aligned} p &= \Gamma_p(A, \sigma^A, \sigma^Z, \mu) \\ \mu' &= \Gamma_\mu(A, \sigma^A, \sigma^Z, \mu). \end{aligned}$$

There are three related but distinct computational challenges involved in the characterization of the mappings Γ_p and Γ_μ . First, the cross-sectional distribution μ is generally intractable as a state variable. Second, the number of aggregate state variables (excluding the cross-sectional distribution μ) is large, since not only aggregate productivity A but also the levels of macro σ^A and micro σ^Z uncertainty are included. Third, the equilibrium price mapping Γ_p must be computed or approximated in some fashion.

We address each of these challenges in the following fashion. As noted in the main text, we assume that a single two-state Markov process $S \in \{0, 1\}$ for uncertainty governs the

evolution of micro and macro volatility across two discrete levels each, so that

$$\begin{aligned} S = 0 &\rightarrow \sigma^A = \sigma_{L,H}^A, \sigma^Z = \sigma_{L,H}^Z \\ S = 1 &\rightarrow \sigma^A = \sigma_{H,H}^A, \sigma^Z = \sigma_{H,H}^Z, \end{aligned}$$

where the transition matrix for S is governed by

$$\Pi^S = \begin{bmatrix} 1 - \pi_{L,H}^\sigma & \pi_{L,H}^\sigma \\ 1 - \pi_{H,H}^\sigma & \pi_{H,H}^\sigma \end{bmatrix}.$$

We then also approximate the intractable cross-sectional distribution μ in the aggregate state space with the current aggregate capital level $K = \int k(z, k, n_{-1}) d\mu$ as well as the lagged uncertainty state S_{-1} . The approximate aggregate state vector is now given by (A, S, S_{-1}, K) , which addresses the first and second computational challenges outlined above.

We are now in a position to define tractable approximations to the equilibrium mappings Γ_p and Γ_μ using the following log-linear rules $\hat{\Gamma}_p$ and $\hat{\Gamma}_K$:

$$\begin{aligned} \hat{\Gamma}_p : \log(\hat{p}) &= \alpha_p(A, S, S_{-1}) + \beta_p(A, S, S_{-1}) \log(K) \\ \hat{\Gamma}_K : \log(\hat{K}') &= \alpha_K(A, S, S_{-1}) + \beta_K(A, S, S_{-1}) \log(K). \end{aligned}$$

The approximations of the aggregate state space and the explicit forms chosen for the equilibrium mappings laid out above are assumptions which can and must be tested for internal accuracy. Below, we will calculate and discuss a range of accuracy statistics commonly used in the literature on heterogeneous agents models with aggregate uncertainty. Although this specification with lagged uncertainty S_{-1} governing coefficients in the forecast rule serves as our baseline, we also consider below two extensions to even larger forecast rule systems below, with additional endogenous aggregate state variables or a longer series of lagged uncertainty realizations included within the forecast rule. None of these extensions changes the impact of an uncertainty shock on economic output much at all relative to our baseline, uniformly causing a recession of around 2.5%.

However, we can now lay out the approximated firm Bellman equation as \hat{V} , where

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

Given this approximated equilibrium concept, we can now lay out the solution algorithm. First, initialize the equilibrium mappings or forecast rules $\hat{\Gamma}_p^{(1)}$ and $\hat{\Gamma}_K^{(1)}$ by guessing initial coefficients $(\alpha_p^{(1)}(A, S, S_{-1}), \beta_p^{(1)}(A, S, S_{-1}))$ and $(\alpha_K^{(1)}(A, S, S_{-1}), \beta_K^{(1)}(A, S, S_{-1}))$. Then, perform the following steps in iteration $q = 1, 2, \dots$ of the algorithm:

Step 1 - Firm Problem Solution

Solve the idiosyncratic firm problem laid out in the Bellman equation for \hat{V} above, conditional upon $\hat{\Gamma}_p^{(q)}$ and $\hat{\Gamma}_K^{(q)}$, resulting in an approximated firm value function $\hat{V}^{(q)}$.

Step 2 - Unconditional Simulation of the Model

Based upon the approximated firm value $\hat{V}^{(q)}$, simulate the economy unconditionally for T periods, without imposing adherence to the assumed equilibrium price mapping $\hat{\Gamma}_p^{(q)}$.

Step 3 - Equilibrium Mapping Update

Based upon the simulated data from the model computed in Step 2, update the forecast mappings to obtain $\hat{\Gamma}_p^{(q+1)}$ and $\hat{\Gamma}_K^{(q+1)}$.

Step 4 - Test for Convergence

If the approximate mappings have converged, i.e. if the difference between $\hat{\Gamma}_p^{(q+1)}$ and $\hat{\Gamma}_K^{(q+1)}$ is smaller than some tolerance ε_Γ according to some predetermined criterion, exit the algo-

rithm. If the mappings have not converged, return to Step 1 for the $q + 1$ -th iteration.

The practical numerical implementation of each of the Steps 1-4 laid out above requires some explanation in more detail. We discuss each step in turn now, noting that a pack containing all of the code for this paper can be found on Nicholas Bloom's website.

Firm Problem Solution

We discretize the state space of the idiosyncratic firm problem. For the endogenous idiosyncratic variables k and n , we choose log-linear grids of size $n_k = 91$ and $n_n = 37$ closed with respect to depreciation. For the exogenous productivity processes z and A , we discretize each process following a straightforward generalization of Tauchen (1986) to the case of time-varying volatility, resulting in processes with $n_z = n_A = 5$ discrete productivity points. Given the discretization of the firm problem, we compute $\hat{V}^{(q)}$ using Howard policy acceleration with 200 value function update steps within each policy loop until the policy functions converge to within some prescribed tolerance. This routine is described in, for example, Judd (1998). Here, the discrete nature of the firm problem allows for exact convergence of firm policies. Throughout the problem, continuation values are computed using linear interpolation in the forecast value of aggregate capital \hat{K} implied by the mapping $\hat{\Gamma}_K^{(q)}$.

Unconditional Simulation of the Model & Convexified Market Clearing

We simulate the model using a fixed set of $T = 5000$ periods of exogenous aggregate productivity and uncertainty realizations (A_t, S_t) , $t = 1, \dots, T$ which follow the discrete Markov chain structures discussed above and are drawn once and for all outside the solution algorithm's outer loop (Steps 1-4).

Within the solution loop, we follow the nonstochastic or histogram-based approach to tracking the cross-sectional distribution proposed by Young (2010). This simulation approach avoids the Monte Carlo sampling error associated with the simulation of individual firms and is faster in practice. In particular, we track a histogram $\hat{\mu}_t$ of weights on individual points (k, n_{-1}, z) in the firm-level discretized state space from period to period. Let the policy functions in period t be given by $n_t(k, n_{-1}, z)$ and $k'_t(k, n_{-1}, z)$ and the transition matrix over idiosyncratic productivity in period t be given by $\Pi^Z(z, z'; S_t)$. If we consider discretized triplets $(k, n_{-1}, z)_i$, $i = 1, \dots, n_k n_n n_z$ then we have that μ_{t+1} is determined by

$$\mu_{t+1}((k', n, z')_j) = \sum_{(k, n, z)_i} \mu_t((k, n, z)_i) \Pi^Z(z_i, z'_j; S_t) \mathbb{I}(k'_j = k'_t((k, n, z)_i), n_j = n_t((k, n, z)_i)).$$

The calculation of the individual firm policy functions k'_t and n_t in period t must be consistent with market-clearing as well as individual firm optimization, however, in the sense that the simulated consumption level C_t and hence market-clearing marginal utility price $p_t = \frac{1}{C_t}$ must be generated by the approximate cross-sectional distribution μ_t as well as the firm policy rules. To guarantee that this occurs, in each period we must iterate over a market clearing price p , using the continuation value $\hat{V}^{(q)}$ but discarding the forecast price level $\hat{p}_t^{(q)}$. For any guess \tilde{p} , we re-optimize the right hand side of the firm Bellman equation given \tilde{p} , i.e. we compute firm policy functions k'_t and n_t to solve

$$\max_{k', n} \left\{ \begin{array}{l} \tilde{p} \left(y - \frac{\theta}{\tilde{p}} n - i - AC^k - AC^n \right) \\ + \beta \mathbb{E} \left[\hat{V}^{(q)}(k', n, z'; A', S', S, \hat{K}^{(q)'}) \right] \end{array} \right\}.$$

Market clearing occurs when the consumption level implied by the price, $C(\tilde{p})$ is equal to the reciprocal of that price, $\frac{1}{\tilde{p}}$. However, as mentioned above we employ a discretization method to the solution of individual firm problems. Although the discretization method is both fast and robust, the resulting excess demand function $e(\tilde{p})$ may contain discontinuities associated with some positive mass of firms discretely shifting capital or investment policies in response to a small price shift. Any such discontinuities in the resulting excess demand function would remove our ability to guarantee arbitrarily precise market clearing. We directly address this issue in a manner which implies internally consistent aggregate dynamics, through "convexification" of firm policies and hence the underlying excess demand function. For

each period in our simulation of the model, we employ the following process to guarantee market clearing.

Step 1 - Precompute certain values over a price grid

Utilize a pre-determined grid of marginal utility price candidates of size N_p $\{\tilde{p}_i\}_{i=1}^{N_p}$. For each individual candidate price level, recompute firm policy functions as discussed above, to obtain values $C(\tilde{p}_i)$ and $K'(\tilde{p}_i)$ defined as

$$C(\tilde{p}_i) = \sum_{(k,n-1,z)_i} \mu_t((k,n-1,z)_i) (y - k' + (1 - \delta_k)k - AC^k - AC^n)$$

$$K'(\tilde{p}_i) = \sum_{(k,n-1,z)_i} \mu_t((k,n-1,z)_i) k'$$

We also label the value of the cross-sectional distribution over values of firm capital, labor, and idiosyncratic productivity in the next period that would prevail given the firm policies implied by the candidate value \tilde{p}_i today as $\mu'(\tilde{p}_i)$.

Step 2 - Construct the implied convexified excess demand function

Given the set of candidate price and implied consumption values $\{\tilde{p}_i, C(\tilde{p}_i)\}_{i=1}^{N_p}$, linearly interpolate the consumption function to obtain the continuous piecewise linear approximating consumption function $\tilde{C}(\tilde{p})$ for aggregate consumption, which is defined for any candidate input price \tilde{p} including those outside the initial grid. Then, define the convexified excess demand function

$$e(\tilde{p}) = \frac{1}{\tilde{p}} - \tilde{C}(\tilde{p}),$$

which is both defined for arbitrary positive candidate input prices \tilde{p} . This convexified excess demand function is also continuous over its entire domain because both of the functions $\frac{1}{\tilde{p}}$ and $\tilde{C}(\tilde{p})$ are continuous. In practice, the excess demand function computed in this manner is also strictly decreasing. Intuitively, the process of linear interpolation of implied consumption values over a pre-determined grid to compute excess demand “convexifies” the market clearing process.

Step 3 - Clear markets with arbitrary accuracy

Given a continuous, strictly decreasing excess demand function, we use a robust hybrid bisection/inverse quadratic interpolation algorithm in \tilde{p} in each period to solve the clearing equation

$$\tilde{e}(\tilde{p}) = 0$$

to any desired arbitrary accuracy ε_p .

Step 4 - Update the firm distribution and aggregate transitions

With a market clearing price p^* in hand for the current period t , we have now cleared an approximation to the excess demand function. However, we must then update the underlying firm distribution and aggregate transitions in a manner which is consistent with the construction of the approximated excess demand function. In particular, we do so by first computing the underlying linear-interpolation weights in the approximation $\tilde{C}(\tilde{p})$ at the point p^* . In other words, we compute the value

$$\omega(p^*) = \frac{p^* - \tilde{p}_{i^*-1}}{\tilde{p}_{i^*} - \tilde{p}_{i^*-1}},$$

where $[\tilde{p}_{i^*-1}, \tilde{p}_{i^*}]$ is the nearest bracketing interval for p^* on the grid $\{\tilde{p}_i\}_{i=1}^{N_p}$. Note that $\omega(p^*)$ defined in this manner lies in the interval $[0, 1]$, and also note by the definition of linear interpolation that $\tilde{C}(p^*) = (1 - \omega(p^*))C(\tilde{p}_{i^*-1}) + \omega(p^*)C(\tilde{p}_{i^*})$.

For each endpoint of the interval \tilde{p}_{i^*-1} and \tilde{p}_{i^*} , we already have in hand values of capital next period $K'(\tilde{p}_{i^*-1})$ and $K'(\tilde{p}_{i^*})$ as well as distributions next period $\mu'(\tilde{p}_{i^*-1})$ and $\mu'(\tilde{p}_{i^*})$ which would prevail at the candidate prices. We simply update the cross-sectional

distribution μ_{t+1} and aggregate capital K_{t+1} for the next period as

$$\begin{aligned} K_{t+1} &= (1 - \omega(p^*))K'(\tilde{p}_{i^*-1}) + \omega(p^*)K'(\tilde{p}_{i^*}) \\ \mu_{t+1} &= (1 - \omega(p^*))\mu'(\tilde{p}_{i^*-1}) + \omega(p^*)\mu'(\tilde{p}_{i^*}). \end{aligned}$$

In practice, when applying this market clearing algorithm, we use a clearing error tolerance of $\varepsilon_p = 0.0001$, therefore requiring that the maximum percentage clearing error cannot rise above 0.01% in any period. We also utilize a grid of size $N_p = 25$, after ensuring that a grid of this size delivers results which do not change meaningfully at higher grid densities. The resulting path of aggregate consumption and implied clearing errors - all of course within the required percentage band of less than 0.01% - are plotted over a representative portion of the unconditional simulation of our model in Figure B1.

One final computational choice deserves mention here. The simulation step, and in particular the determination of the market-clearing price p_t in each period, is the most time-consuming portion of the solution algorithm, especially since firm policies and aggregate variables must be pre-computed along a grid of N_p candidate price values. We have found that substantial speed gains can be obtained, with little change in the resulting simulated aggregate series, if the re-optimization of firm policies conditional upon a price for \tilde{p} in period t is only computed for states above a certain weight $\varepsilon_{dist} = 0.0001$ in the histogram μ_t . Thereafter, only those firm-level states $(k, n_{-1}, z)_i$ with weight above ε_{dist} are used in the calculation of the market clearing price p_t and the evolution of aggregates in the economy from t to $t + 1$.

Equilibrium Mapping Update

At this point in the solution algorithm in iteration q we have obtained a series of prices and capital stocks (p_t, K_t) , $t = 1, \dots, T$ together with exogenous aggregate series (A_t, S_t, S_{t-1}) . Recall that we set $T = 5000$. These simulated series are conditioned upon the equilibrium mappings $\hat{\Gamma}_p^{(q)}$ and $\hat{\Gamma}_k^{(q)}$. To update the equilibrium mappings, which are simply lists of coefficient pairs $(\alpha_p^{(q)}((A, S, S_{-1})_i), \beta_p^{(q)}((A, S, S_{-1})_i))$ for each discrete triplet $(A, S, S_{-1})_i$, we first discard the $T_{erg} = 500$ initial periods in the simulation to remove the influence of initial conditions. For each set of values $(A, S, S_{-1})_i$, we collect the subset of periods $t \in \{T_{erg} + 1, \dots, T\}$ with those particular exogenous aggregate states. We then update the mapping coefficients via the following OLS regressions on that subset of simulated data:

$$\begin{aligned} \log(p_t) &= \alpha_p((A, S, S_{-1})_i) + \beta_p((A, S, S_{-1})_i) \log(K_t) \\ \log(K_{t+1}) &= \alpha_K((A, S, S_{-1})_i) + \beta_K((A, S, S_{-1})_i) \log(K_t). \end{aligned}$$

After collecting the estimated coefficients we obtain updated mappings $\hat{\Gamma}_p^{(q+1)}$ and $\hat{\Gamma}_K^{(q+1)}$.

Test for Convergence

At this point within the solution algorithm, there are several potential criteria which could in principle be used to determine final convergence of the equilibrium mappings and therefore the model solution. One option is to declare convergence if the maximum absolute difference between any two corresponding coefficients is less than some tolerance $\varepsilon_{mapping}$.

However, we have found that there are substantial speed gains, with little difference in the resulting simulated series, if instead convergence is defined based upon the accuracy of the forecast system itself. In particular, a commonly accepted practice in the literature is to define the internal accuracy of a forecast mapping based upon the maximum Den Haan (2010) statistics. These statistics for both capital K and price p , which we label DH_k^{max} and DH_p^{max} , respectively, are the maximum absolute log difference across the full model simulation of the simulated series (p_t, K_t) and their dynamically forecasted counterparts (p_t^{DH}, K_t^{DH}) . “Dynamic forecasts” are simply the result of repeated application of the equilibrium mappings $\hat{\Gamma}_p$ and $\hat{\Gamma}_K$ using previously predicted values as explanatory or right-hand-side variables in the forward substitution. Such an approach allows for the accumulation of prediction error within the system to a more stringent degree than would result from a one-period or “static” evaluation of prediction errors. We conclude that the forecast mapping has converged, and therefore that the model has been solved, when the change

in the model's accuracy statistics is less than a prescribed tolerance $\varepsilon_{mapping} = 0.01\%$, i.e. when

$$\max\{|DH_p^{max,(q+1)} - DH_p^{max,(q)}|, |DH_K^{max,(q+1)} - DH_K^{max,(q)}|\} < \varepsilon_{mapping}.$$

B.2 Conditional Response Calculations

In this subsection we describe the calculation of the response of the economy to an uncertainty shock, an uncertainty shock coincident with an aggregate productivity shock, and a policy or wage subsidy experiment.

The Effects of an Uncertainty Shock

Armed with the solution to the model following the algorithm outlined above, we compute the conditional response of the economy to an uncertainty shock by simulating $N = 2500$ independent economies of length $T_{IRF} = 100$. We designate $T_{shock} = 50$ as the shock period. In each economy i , we simulate the model as normal for periods $t = 1, \dots, T_{shock} - 1$ starting from some initial conditions and following the procedure for within-period market clearing outlined above. In period T_{shock} , we impose high uncertainty for economy i , i.e. $S_{iT_{shock}} = 1$. Thereafter, each economy i evolves normally for periods $t = T_{shock} + 1, \dots, T_{IRF}$.

Given any aggregate series of interest X , with simulated value X_{it} in economy i , period t , we define the period t response of the economy to an uncertainty shock in period T_{shock} , \hat{X}_t as

$$\hat{X}_t = 100 \log \left(\frac{\bar{X}_t}{\bar{X}_{T_{shock}-1}} \right),$$

where \bar{X}_t is the cross-economy average level in period t : $\frac{1}{N} \sum_i X_{it}$. The notation \hat{X}_t for the percentage deviations of a series from its pre-shock level will be used throughout this subsection in the context of various experiments and shocks. Also, note that for the purposes of labelling the figures in the main text, we normalize $T_{shock} = 1$. The initial conditions we use to start each simulation include a low uncertainty state, the median aggregate productivity state, and the cross-sectional distribution from a representative period in the unconditional simulation of the model.

The Effects of an Uncertainty Shock and First-Moment Shock

To simulate the effect of an uncertainty shock coincident with a first-moment shock to exogenous aggregate productivity A , we follow the same basic procedure as above, simulating independent economies $i = 1, \dots, N$ for quarters $t = 1, \dots, T_{IRF}$, where $N = 2500$ and $T_{IRF} = 100$. For each economy the aggregates evolve normally until a shock period $T_{shock} = 50$. In the shock period, we impose a high uncertainty state for all economies, i.e. we set $S_{iT_{shock}} = 1$ for all i . However, we also wish to impose a negative aggregate productivity shock in period T_{shock} equal to -2% on average. To operationalize this, we choose a threshold probability $\bar{\xi}$ and then draw independent uniform random variables $\xi_i \sim U(0, 1)$ for each economy i . With probability $\bar{\xi}$, i.e. if $\xi_i \leq \bar{\xi}$, we set the discretized aggregate productivity state $A_{iT_{shock}}$ equal to the lowest grid point value. With probability $1 - \bar{\xi}$, i.e. if $\xi_i > \bar{\xi}$, we allow the aggregate productivity process to evolve normally for economy i in period T_{shock} . For all economies, post-shock periods $t = T_{shock} + 1, \dots, T_{IRF}$ evolve normally. By iterating on the value of $\bar{\xi}$, we choose a probability of a first-moment shock which guarantees that, on average, a -2% shock obtains:

$$\hat{A}_{T_{shock}} = 100 \log \left(\frac{\bar{A}_{T_{shock}}}{\bar{A}_{T_{shock}-1}} \right) = -2.$$

Note that the plotted series of shock responses \hat{X}_t for individual aggregate series of interest X are defined exactly as in the baseline uncertainty shock case above, again normalizing the shock period to $T_{shock} = 1$ for plotting purposes.

The Effects of a Wage Subsidy Policy Experiment

The effect of a wage subsidy policy experiment require the simulation of three additional experiments, each with the same structure of $i = 1, \dots, N$ independent economies with

quarters $t = 1, \dots, T_{IRF}$ where $N = 2500$ and $T_{IRF} = 100$. All shocks and policies are imposed in period $T_{shock} = 50$, with normal evolution until that period.

The first additional experiment we run is a wage subsidy with no uncertainty shock. In this case in period T_{shock} we impose an unanticipated wage bill subsidy so that the wage rate faced by firms is equal to $(100 - 1)\%$ of the equilibrium wage rate w . This subsidy is financed with lump-sum taxes on the representative household. Afterwards each economy evolves normally, and we denote the resulting percentage deviations from the pre-shock period of this post-subsidy path by $\hat{X}_t^{subsidy,normal}$.

The second additional experiment we run is a wage subsidy concurrent with an uncertainty shock. In this case, in period T_{shock} we impose the same unanticipated wage bill subsidy but also a high uncertainty state $S_{T_{shock}} = 1$. Afterwards each economy evolves normally, and we denote the resulting post-subsidy path of this simulation by $\hat{X}_t^{subsidy,unc}$.

The third additional experiment we run is a simulation of “normal times,” where no wage subsidy and no uncertainty shock is imposed in period T_{shock} . We denote the resulting path of this economy as \hat{X}_t^{normal} .

If we denote the baseline response of the economy to an uncertainty shock as \hat{X}^{unc} , then we can now define the two objects reported in the main text. The effective of a wage subsidy in normal times is simply

$$\hat{X}_t^{subsidy,normal} - \hat{X}_t^{normal},$$

while the effect of a wage subsidy with an uncertainty shock is

$$\hat{X}_t^{subsidy,unc} - \hat{X}_t^{unc}.$$

An Alternative Approach to Computing the Effect of Uncertainty

In nonlinear models there is not a universally defined notion of a conditional response or impulse response. As an alternative to our baseline method of computing the impact of uncertainty on the economy described above, we also have checked and report in Figure B4 the results from an alternative procedure constructed for nonlinear models based on Koop, et al. (1996). That procedure works as follows. Draw $i = 1, \dots, N$ sets of exogenous random draws for the uncertainty and aggregate productivity processes, each of which include $t = 1, \dots, T_{IRF}$ periods. Then, for each set of draws or economies i , simulate two versions, a *shock* and *no shock* economy. In the *shock* economy, simulate all macro aggregates X_{it}^{shock} for each period $t = 1, \dots, T_{shock} - 1$ as normal. Then, in period T_{shock} impose high uncertainty. For all periods $t = T_{shock} + 1, \dots, T_{IRF}$, simulate the economy as normal. Then, for version *no shock*, simulate all macro aggregates $X_{it}^{noshock}$ unconditionally without any restrictions. The only difference in the exogenous shocks in the *shock* and *noshock* economies is the imposition of the single high uncertainty state in period T_{shock} . The resulting effect of an uncertainty shock on the aggregate X is given by the cross-economy average percentage difference between the shock and no shock versions:

$$\hat{X}_t^{Koop} = \frac{1}{N} \sum_{i=1}^N \log \left(X_{it}^{shock} / X_{it}^{noshock} \right).$$

By construction, this response is equal to zero before T_{shock} . The figures plotted in the text normalize T_{shock} to 1 for labelling and report $100\hat{X}_t^{Koop}$. Note that N , T_{IRF} , and T_{shock} are identical to our baseline version. Figure B4 reports the response of output, labor, investment, and consumption to an uncertainty shock, computed in both the baseline and the alternative fashion described here, labeled “simulation differencing” in the figure. The response of the economy to an uncertainty shock is distinct but does not vary qualitatively across the two alternative notions of impulse responses.

A Denser Grid to Solve the Model

In our baseline solution to the model we, use $n_z = n_A = 5$ discrete productivity points for both micro productivity z and macro productivity A . To explore the impact of the grid density on our conclusions about the impact of uncertainty on the macroeconomy, we also solved the model with a denser grid of $n_z = n_A = 9$ grids points for each productivity process. The move expands the size the numerical grid by a factor of $81/25 = 3.24$. This

of course entails additional costs in the solution of the firm problem, but the unconditional simulation of the model had to be extended from $T = 5,000$ quarters to $T = 15,000$ quarters in order to ensure that the forecast rule for each configuration of aggregate productivity and uncertainty states had enough density to be accurately estimated. Figure B5 reports the resulting response of the economy to an uncertainty shock, computed using our baseline impulse response simulation method. The response of the economy to an uncertainty shock is distinct but does not vary qualitatively with the larger grid. Ideally the response could be computed with an even denser grid, although current computational limits render such an approach infeasible.

B.3 Internal Accuracy of the Approximate Equilibrium Mappings

In this subsection we first report basic accuracy statistics emphasized by the literature on heterogeneous agents business cycle models for evaluation of the approximate equilibrium mappings or forecast rules $\hat{\Gamma}_K$ and $\hat{\Gamma}_p$. We also introduce and examine the results implied by alternative forecasting systems with more complexity than our baseline approach. Finally, we conclude with a brief discussion of alternative solution algorithms to Krusell Smith which have been proposed in the literature on heterogeneous agents business cycle models.

Static Accuracy Statistics

Recall that the prediction rules for current-period market-clearing prices $\hat{\Gamma}_p$ and next period's aggregate capital state take the following form:

$$\begin{aligned}\log(\hat{p}_t) &= \alpha_p(A_t, S_t, S_{t-1}) + \beta_p(A_t, S_t, S_{t-1}) \log(K_t) \\ \log(\hat{K}_{t+1}) &= \alpha_K(A_t, S_t, S_{t-1}) + \beta_K(A_t, S_t, S_{t-1}) \log(K_t).\end{aligned}$$

Using the standard terminology of time series forecasting, we refer to the predicted price and capital levels \hat{p}_t and \hat{K}_{t+1} as “static” or one-period ahead forecasts, since their calculation exploits information from the current period t , namely K_t and the state (A_t, S_t, S_{t-1}) . For each discretized triplet (A, S, S_{-1}) of exogenous aggregate states, Table B1 reports the accuracy of these static forecasts for our baseline model based on two common metrics: the R^2 of each prediction regression and the percentage root mean squared error (RMSE) of the predictions. As the table shows, by these metrics the prediction rules are generally quite accurate with R^2 near 1 and RMSEs never rising above approximately 0.5%.

Alternative Forecasting or Market-Clearing Assumptions

The dynamics of consumption C , and hence the dynamics of the market-clearing marginal utility price $p = \frac{1}{C}$, crucially influence the behavior of investment in our general equilibrium economy. A high price today relative to the future signals that investment is costly, and vice-versa, resulting in smoother investment paths relative to a partial equilibrium structure. Given the importance of the price p for determining investment dynamics, we now compare the behavior of output, together with realized and forecast prices, for four different economies that encompass different forecast rules or market-clearing assumptions for the aggregate price level p . Dynamic forecast accuracy statistics for each approach are reported in Table B2, and Figure B2 plots the response of output, consumption, and the implied forecast errors for each strategy after an uncertainty shock.

1) Baseline Economy

This economy serves as the baseline for calculations of the uncertainty impulse responses presented in the main text. The forecast rule with lagged uncertainty included and the market clearing algorithm based on convexified excess demand are as described above.

2) Extra Uncertainty Lags Economy

This economy uses an identical market clearing algorithm as the baseline. However, the forecast levels of price and future capital are based on the generalized log-linear forecast rules

$$\begin{aligned}\log(\hat{p}_t) &= \alpha_p(A_t, S_t, S_{t-1}, \dots, S_{t-k}) + \beta_p(A_t, S_t, S_{t-1}, \dots, S_{t-k}) \log(K_t) \\ \log(\hat{K}_{t+1}) &= \alpha_K(A_t, S_t, S_{t-1}, \dots, S_{t-k}) + \beta_K(A_t, S_t, S_{t-1}, \dots, S_{t-k}) \log(K_t).\end{aligned}$$

In other words, the forecast rule coefficients are allowed to depend upon a larger set of conditional variables with extra lags of uncertainty beyond S_{t-1} . In the results reported in Table B2 and Figure B2, the extra uncertainty lags results are computed with $k = 3$, i.e. with a total of 4 values or one year of uncertainty realizations S_t, \dots, S_{t-3} conditioned into the forecast rule. Because this procedure results in a smaller subsample for the regression update of each forecast rule, we also expand the size of the unconditional simulation of the model to 20,000 quarters rather than the 5,000 quarters used for the baseline model, continuing to discard the first 500 quarters to cleanse the simulation of the impact of initial conditions.

3) *Extra Forecast Moment Economy*

This economy adds an additional endogenous moment, M_t to the forecast rules for marginal utility. The addition requires expansion of the baseline forecast rule system to include an additional explanatory variable on the right hand side as well as another equation, taking the resulting form

$$\begin{aligned}\log(\hat{p}_t) &= \alpha_p(A_t, S_t, S_{t-1}) + \beta_p(A_t, S_t, S_{t-1}) \log(K_t) + \gamma_p(A_t, S_t, S_{t-1}) M_t \\ \log(\hat{K}_{t+1}) &= \alpha_K(A_t, S_t, S_{t-1}) + \beta_K(A_t, S_t, S_{t-1}) \log(K_t) + \gamma_K(A_t, S_t, S_{t-1}) M_t \\ \hat{M}_{t+1} &= \alpha_M(A_t, S_t, S_{t-1}) + \beta_M(A_t, S_t, S_{t-1}) \log(K_t) + \gamma_M(A_t, S_t, S_{t-1}) M_t\end{aligned}$$

After the forecast system itself is expanded in the manner above, all value functions in the solution of the model must take another aggregate input M_t , but the computational strategy and market clearing approach based on convexified aggregate demand remain otherwise identical. We considered a number of candidate values for inclusion as the extra additional endogenous moment, including the second moments of idiosyncratic capital and labor inputs across firms. Our chosen moment, the predetermined beginning of period t /end of period $t - 1$ cross-sectional covariance $M_t = Cov(\log k_{t-1}, \log z_{t-1})$, performs well among our candidates because it proxies for allocative efficiency in the economy as a measure of the alignment of idiosyncratic profitability and capital inputs across firms. After an uncertainty shock M_t declines steadily, rising after uncertainty subsides. Within our forecast structure, the extra covariance moment M_t thus serves as a continuous proxy for “time since an uncertainty shock,” a useful piece of information given the rich dynamics of our economy. Because of the size of the expanded forecast rule in this case, we use an expanded unconditional simulation of 20,000 quarter length rather than the 5,000 quarters used to update the more parsimonious baseline forecast system.

4) *Nonconvexified Clearing*

As a comparison for economies 1)-3), this fourth alternative is based on approximate market clearing choosing prices in each period to come as close as possible to clearing the nonconvexified or raw excess demand function. As noted above, potential discontinuities due to discrete firm choices embedded in this excess demand function imply that markets may not be cleared to arbitrary accuracy, although the computational burden is substantially lower. Although the clearing algorithm is distinct, this economy uses a forecast system identical to the baseline economy.

How do each of these alternative computational strategies compare in practice? First, we evaluate their conditional patterns and accuracy after an uncertainty shock. The left panel of Figure B2 plots the response of output in each economy to an uncertainty shock, the middle panel plots the response of consumption, and the right panel plots the response of forecast errors, (the difference between the paths for realized prices p_t and forecast prices \hat{p}_t). Crucially, each economy considered here delivers an immediate drop in output of around 2.5% in the wake of an uncertainty shock, followed by a quick recovery. Each of the economies relying upon arbitrarily accurate market clearing, i.e. the baseline economy (x symbols) the extra uncertainty lags economy (* symbols), and the economy with an extra forecast rule moment (diamonds), recovers more strongly than the economy relying on nonconvexified market clearing (o symbols). In every economy, the combination of rising consumption levels and hence high interest rates from period 2 onwards - plotted in the middle panel - with continued misallocation of inputs depresses investment after an

initial recovery. The result is a second decline in output as discussed in the main text, a shared conclusion across these computational strategies. The right panel plots the path of forecast errors for prices after an uncertainty shock across methods, and the methods providing the most accurate conditional paths of prices after the shock are the models with extra lags of uncertainty or an extra moment in the forecast rule. These computational structures essentially eliminate any mismatch between forecast and realized prices in the wake of an uncertainty shock, and the left panel of the figure reveals that the basic pattern of output after an uncertainty shock in the more parsimonious baseline economy matches the computationally costly extensions quite well.

Dynamic Accuracy Statistics

As Den Haan (2010) notes, one-period ahead or “static” metrics like the R2 of the forecast regressions reported in Table B1 are not very stringent tests of the accuracy of a forecasting system in a Krusell-Smith approach. Instead, he recommends using “dynamic” forecasts and evaluating these for accuracy. The procedure for producing dynamic forecasts such as these is to use s -period ahead forecasts as an input into $s+1$ -period ahead forecasts, iterating forwards to any desired horizon conditioning only on the realized values of exogenous aggregate series. Forecast errors can accumulate over time in a manner which is obscured by the one-period ahead accuracy metrics. In this section, we discuss the performance of our model using the dynamic forecast accuracy metrics implied by Den Haan.

In our model general equilibrium impacts a firm’s input problem only through the value of the marginal utility price p . Different from Krusell and Smith (1998)’s economy with a representative firm, the forecasted value of aggregate capital K does not independently influence prices in the economy outside of its influence on the forecast for p . So to fully characterize the accuracy of a forecast rule in our model, it suffices to examine the accuracy of the implied forecasts for the marginal utility or price p . Therefore, Table B2 provides a comprehensive set of Den Haan accuracy statistics for p at various horizons. Each block of rows in Table B2 reports the mean and maximum errors in the dynamic forecasts in each of our four computational versions of the model: our baseline, an economy with extra lags in the forecast rule, an economy with an extra moment in the forecast rule, and a version of our baseline with nonconvexified market clearing. Each column reports the errors at a different set of horizons, starting at 3 years (12 periods ahead given the quarterly periodicity of the model) and moving to 12 years (48 model periods or quarters).

For concrete interpretation, note that the number at the far left of the top row, 0.63, indicates that the average price forecasting error of firms in this model at a horizon of 12 quarters or 3 years is 0.63%. The magnitude of error is difficult to interpret on its own. However, by our reading, this level of accuracy is in fact comparable to dynamic forecast errors from other exercises solving firm-level investment models with similarly rich aggregate state spaces. See, for example, the computational appendix to Khan and Thomas (2013), reporting dynamic forecast errors of around 0.8% in a model with both productivity and financial shocks. By contrast, note that papers which report somewhat smaller dynamic forecast errors, such as Terry (2015) at around 0.1% for a version of the baseline Khan and Thomas (2008) model, only do so for models with a single aggregate shock and univariate optimization problems at the firm level.

Comparing across methods and horizons in Table B2, we see that the forecasting system with an extra lag of uncertainty performs almost uniformly better than the baseline forecasting system, although not by a large margin. This improvement is unsurprising given the elimination of meaningful swings in the price forecasting error in this economy after an uncertainty shock documented in Figure B2. Interestingly, the economy with an extra moment in the forecast rule performs comparably or worse than the baseline economy, depending on the exact metric used. The direct implication is that additional forecast uncertainty from the accumulated forecast errors in the prediction rule for the extra covariance moment itself degrade the forecasting effectiveness of the expanded system. Also, note that although the economy with nonconvexified clearing appears to perform best at various horizons, it is not strictly comparable to the other methods. The use of a less accurate clearing algorithm implies that this model’s price series is less volatile than the price series for the other

computational strategies purely for numerical reasons.

Taken as a whole, we feel that the good performance of the baseline model using the accuracy metrics in Table B2 is strong enough to justify its continued use, especially in light of the almost constant immediate impact of uncertainty on output revealed by Figure B2 across each alternative forecasting strategy and its smaller computational burden.

Alternatives to Krusell Smith

We conclude with a discussion of some of the alternative solution methods proposed in the literature for heterogeneous agents business cycle models. Our analysis uses the Khan and Thomas (2008) extension of Krusell and Smith (1998). More precisely, we rely upon an approximation of the full cross-sectional distribution μ by the first moment K as well as log-linear forecast rules for predicting equilibrium prices p and the evolution of aggregate capital K . The repeated simulation and forecast rule update steps required by this solution algorithm are quite computationally intensive. In principle, use of an alternative technique such as the method of parameterized distributions (Algan, Allais, and Den Haan 2008, 2010) or the explicit aggregation method (Den Haan and Rendahl 2010) might be more computationally convenient and/or yield gains in terms of the internal accuracy of the approximate equilibrium mappings.

However, the results in Terry (2015) suggest that these alternative algorithms in the context of the closely related Khan and Thomas (2008) model do not yield substantial accuracy improvements over the extended Krusell and Smith (1998) approach. By contrast, that paper reports that both the method of parametrized distributions and the explicit aggregation method yield virtually unchanged quantitative implications for the business cycle in the context of the benchmark lumpy capital adjustment model, but that the internal accuracy of the equilibrium mappings for each method are slightly degraded relative to the Krusell and Smith (1998) algorithm. Furthermore, although both alternative methods yield a substantial reduction in model solution time relative to the Krusell and Smith approach, unconditional simulation time for the model remains costly in the case of the explicit aggregation method and in fact increases for the method of parametrized distributions. Simulation time is particularly critical for our computational approach, because the structural estimation strategy we use relies upon calculation of simulated moments capturing the behavior of micro and macro uncertainty. Therefore, an increase in model solution speed without an increase in simulation speed would yield only a small gain in practice.

The results in Terry (2015) are of course specific to the context of the Khan and Thomas (2008) model rather than our baseline model. However, our model environment is quite similar to the Khan and Thomas (2008) structure, differing primarily through the inclusion of second-moment or uncertainty shocks. Based on this similarity, we conclude that our reliance upon the adapted Krusell and Smith (1998) approach is a reasonable choice in the context of our heterogeneous firms model. Terry (2015) also considers the performance of a conceptually distinct solution technique, the projection plus perturbation method proposed by Reiter (2009). That first-order perturbation approach would not capture the impact of variations over time in aggregate uncertainty in our model, since it relies upon a linearization of the model's equilibrium around a nonstochastic steady state.

C Online Appendix: Simulated of Moments Estimation

In this section we lay out the SMM approach used for the structural estimation of the uncertainty process in the baseline model. We begin by defining the estimator and noting its well known asymptotic properties. We then discuss in more detail the practical choices made in the implementation of the estimation procedure, both empirically and in the model.

Overview of the Estimator

Our dataset X consists of microeconomic and macroeconomic proxies for uncertainty computed from data covering 1972 – 2010. In each year t , an observation consists of

$$X_t = (IQR_{t+1}, \sigma_{t+1,1}, \sigma_{t+1,2}, \sigma_{t+1,3}, \sigma_{t+1,4}).$$

Above, the variable IQR_t is the year t cross-sectional interquartile range of TFP shocks constructed from the establishment-level data in the Census of Manufactures and the Annual Survey of Manufacturers and plotted on the left hand scale in Figure 3 in the main text. The variable $\sigma_{t,j}$ is the quarterly estimated GARCH(1,1) heteroskedasticity of the series dTFP from John Fernald’s website, i.e. the annualized change in the aggregate quarterly Solow residual, in year t and quarter j in US data. The macroeconomic series of estimated conditional heteroskedasticity is plotted in Figure A2. Stacking the quarterly observations for the macroeconomic uncertainty proxy $\sigma_{t,j}$ together with the annual microeconomic uncertainty proxy IQR_t in the same year within the vector X_t allows us to account for the frequency difference in the data underlying moment calculation.

Our estimator is based on the $r \times 1 = 8 \times 1$ moment vector including the mean, standard deviation, skewness, and serial correlation of the micro and macro components of X_t . The parameter vector θ is the $q \times 1 = 6 \times 1$ vector $(\sigma_L^A, \frac{\sigma_H^A}{\sigma_L^A}, \sigma_L^Z, \frac{\sigma_H^Z}{\sigma_L^Z}, \pi_{L,H}^\sigma, \pi_{H,H}^\sigma)'$. Note that since $r > q$, this is overidentified SMM.

We minimize the scalar objective defined from the moments $g_S(\theta)$ delivered by a model simulation with length ST , where T is the length of our simulated dataset in years. Let the vector $g(X)$ be the set of moments computed from the data. Then our estimator $\hat{\theta}$ is equal to

$$\hat{\theta} = \arg \min_{\theta} (g_S(\theta) - g(X))' W (g_S(\theta) - g(X)), \quad (22)$$

where $W = \text{diag}(1/g(X))^2$, an $r \times r$ symmetric matrix. Therefore, in practice to estimate the model we minimize the sum squared percentage deviations of the model and data moments.

Subject to regularity conditions, we have that by standard SMM arguments laid out in, for example, Lee and Ingram (1991), our SMM estimator is consistent, i.e. $\hat{\theta} \rightarrow_p \theta$, where θ is the population parameter vector. Furthermore, $\hat{\theta}$ is asymptotically normal with

$$\sqrt{T} (\hat{\theta} - \theta) \rightarrow_d N(0, \Omega).$$

The $q \times q$ asymptotic covariance matrix Ω also follows standard SMM formulas and is given by

$$\Omega = \left(1 + \frac{1}{S}\right) (A'WA)^{-1} A'W\Sigma WA (A'WA)^{-1'} \quad (23)$$

where Σ is the $r \times r$ asymptotic covariance matrix of the moment vector $g(X)$ and A is the $r \times q$ Jacobian of the moments with respect to the parameter vector θ .

Practical Empirical Implementation

Our estimate of the covariance matrix of the moments $\hat{\Sigma}$ allows for arbitrary stationary serial correlation across years in the underlying series X_t using the block bootstrap procedure with random block length due to Politis and Romano (1994). This procedure also allows for an arbitrary stationary covariance structure across microeconomic and macroeconomic

moments across years and within years in our sample of data. We compute 50,000 bootstrap replications and choose a mean block length of $4 \times T^{1/3}$ years following the discussion in Politis and Romano (1994).

We compute model moments based on a simulation of 5000 quarters, discarding the initial 500 quarters. Given our conversion to annual frequency and the length of our sample, the simulation multiple parameter S in the asymptotic standard errors is given by $S = 1124/38 \approx 29.58$. The resulting inflation in standard errors due to simulation variability is $1 + \frac{1}{S} \approx 1.03$.

To minimize the SMM objective function and compute $\hat{\theta}$ we use particle swarm optimization, which is a robust global stochastic optimization routine. We also numerically differentiate the model moment function $g_S(\theta)$ at the estimated parameters $\hat{\theta}$ to obtain \hat{A} by averaging over forward-difference approximations starting from the estimated parameters $\hat{\theta}$ with step sizes 0.4, 0.5, and 0.6%. With $\hat{\Sigma}$ and \hat{A} in hand, the estimated asymptotic covariance matrix for $\hat{\theta}$ is given by

$$\hat{\Omega} = \left(1 + \frac{1}{S}\right) \left(\hat{A}'W\hat{A}\right)^{-1} \hat{A}'W\hat{\Sigma}W\hat{A} \left(\hat{A}'W\hat{A}\right)^{-1'}. \quad (24)$$

Note that there are multiple ways in which the moment covariance matrix $\hat{\Sigma}$ could in principle be estimated, so in unreported results we have also considered the calculation of an estimate of Σ based on a long simulation of 250 replication economies within the model itself. We then recalculated our moment vector for each economy and used this set of data to compute an alternative estimate $\hat{\Sigma}$. The resulting moment covariance matrix is similar to the estimate we describe above, never changing inference of the underlying structural parameters meaningfully. However, the alternative approach typically increases the precision of our parameter estimates, so for conservatism we chose to rely on the results from our empirical stationary bootstrap procedure.

Measurement Error in the Data

In this subsection 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 $\log(\text{TFP})$ in our Census micro data sample. The resulting estimate is a useful input into our procedure for computing model equivalents of data moments described in the next subsection. 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 - which is quite likely given that ASM respondents are shown prior years values when filling in the current year - so that $\text{cov}(e_t, e_{t-1}) = \sigma_{t,t-1}$. 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%.

Practical Model Implementation

The SMM estimation routine requires repeatedly computing the vector of model moments $g_S(\theta)$ given different values of the parameter vector θ . The macroeconomic moments are straightforward to compute, once we have solved the model in general equilibrium following the solution algorithm outlined above. We unconditionally simulate the model for 5000 quarters and discarding the first 500 periods, using the computationally efficient version of our model with lagged uncertainty in the forecast rule and clearing the non-convexified excess demand function. We form a quarterly aggregate TFP series equal in quarter s to

$$TFP_s = \log(Y_s) - \frac{1}{3} \log(K_s) - \frac{2}{3} \log(N_s),$$

where Y_s , K_s , and N_s are the aggregate output, capital, and labor inputs in the economy. The one-third/two-third weighting of capital and labor is consistent with standard values in the macroeconomic literature, and the use of s rather than t is to avoid notational confusion with the annual-frequency microeconomic dispersion series computed below. We compute the annualized change in $dTFP_s = 400(TFP_s - TFP_{s-1})$, estimating the conditional heteroskedasticity for $dTFP$ in each quarter using a GARCH(1,1) model. The mean, standard deviation, skewness, and serial correlation of the resulting series $\hat{\sigma}_s$ form the macroeconomic block of the model moments $g_S(\theta)$.

The microeconomic moments are based on the cross-sectional dispersion of innovation in regressions of micro-level TFP on lagged values. To run equivalent regressions in the model requires the simulation of a panel of individual firms. We simulate 1000 individual firms for the same 5000 period overall simulation of the model, again discarding the first 500 quarters of data. We must convert the quarterly simulated firm data to annual frequency accounting for the data timing in the underlying Census data sample. We compute firm j 's capital stock in year t k_{jt} as the firm's fourth-quarter capital stock. The labor input n_{jt} is the first-quarter value. Annual output y_{jt} is cumulated over the year. The micro-level TFP measure is a Solow residual given by

$$\log(\hat{z}_{jt}) = \log(y_{jt}) - \frac{1}{3} \log(k_{jt}) - \frac{2}{3} \log(n_{jt}).$$

In the subsection above we laid out evidence, concordant with Collard-Wexler (2011), that measurement accounts for around half of the observed TFP variation in our Census sample. To account for this type of measurement error, we first compute the unconditional standard deviation σ_z of the measured TFP values $\log(\hat{z}_{jt})$ above. Then we draw a set of independent measurement error shocks $\xi_{jt} \sim N(0, \sigma_z)$ and compute measurement-error adjusted tfp values \hat{z}_{jt}^* given by

$$\log(\hat{z}_{jt}^*) = \log(\hat{z}_{jt}) + \xi_{jt}.$$

We then run the following panel regression of measured simulated tfp on lagged measured tfp and firm and year fixed effects

$$\log(\hat{z}_{jt}^*) = \mu_j + \lambda_t + \rho \log(\hat{z}_{jt-1}^*) + \varepsilon_{jt}.$$

The cross-sectional interquartile range $I\hat{Q}R_t$ of the residual innovations to micro-level TFP, $\hat{\varepsilon}_{jt+1}$, forms the micro-level uncertainty proxy for year t in our simulated data. The microeconomic block of the model moments $g_S(\theta)$ is simply given by the mean, standard deviation, skewness, and serial correlation of the series $I\hat{Q}R_t$.

Underlying Uncertainty in the Model versus the Uncertainty Proxy

The micro uncertainty proxy used for model estimation, IQR_t , differs from the underlying volatility of micro productivity shocks σ_t^Z in the model, in multiple important ways. These differences are crucial to keep in mind when comparing the moderate empirical variability of the uncertainty proxy in Figure 3 with the high estimated variation in underlying volatility upon impact of an uncertainty shock.

First, IQR_t is measured at annual frequency rather than the quarterly frequency of σ_t^Z . Second, as in the Census data sample itself, the uncertainty proxy IQR_t is based on establishment-level Solow residuals computed using temporally mismatched data drawn

from different quarters in the year (labor, capital) or summed throughout the year (output).¹⁸ Third, the series IQR_t reflects substantial, and empirically plausible, measurement error in the micro data.

Figure C1 plots several decomposed series which unpack the contribution of each of these steps. In the left hand panel, we show a representative 120-quarter period drawn from the unconditional simulation of the model, plotting four series at quarterly frequency.

The series labelled “Quarterly” is the interquartile range of the underlying micro productivity shocks in the model in a given quarter. This reflects the underlying fundamental uncertainty concept in the model.

The second series, labelled “Annual,” is the interquartile range of a normal distribution with standard deviation equal to the standard deviation of the sum of the quarterly productivity shocks within a year. This uncertainty series is constant within the four quarters of a year.

The third series, labelled “Mismatch,” is the interquartile range of measured TFP innovations computed from an annual panel regression on the Solow residuals $\log(\hat{z}_{jt})$ from above. This uncertainty series therefore reflects the contribution of mismatch to measured dispersion, but not the contribution of measurement error, and it does not vary within the year.

Finally, the series labeled “Measurement Error” is simply equal to the annual model uncertainty proxy IQR_t defined above. It is the interquartile range of measured TFP innovation computed from an annual panel regression on the mis-measured Solow residuals $\log(\hat{z}_{jt}^*)$ and also does not vary within a year.

As is evident from the left hand side of Figure C1, each annual uncertainty measure naturally has a higher level than the quarterly concept. Furthermore, the Mismatch and Measurement Error series fluctuate slightly even if underlying uncertainty does not change, because quarterly productivity shocks throughout a given year lead to input and output responses within firms that are measured at different times and do not wash out in the Solow residual productivity measurement. Finally, the accounting for measurement error of course leads to a higher measured level of dispersion in TFP shocks.

The right hand side of Figure C1 sheds light on the variability of each uncertainty concept relative to its mean level. The first four bar heights are equal to the coefficient of variation of each series from the left hand side, computed over the full unconditional simulation of the model. The fifth bar height reflects the coefficient of variation of our micro uncertainty proxy from the data, i.e. the interquartile range of TFP shocks in the Census of Manufactures sample plotted in Figure 3. Reflecting the large estimated increase in uncertainty upon impact of an uncertainty shock, the quarterly series has a high coefficient of variation of around 75%. The annualization process does little to change this pattern. However, the temporal mismatch of TFP measurement within the year leads to a large reduction in the coefficient of variation to about 28%. Measurement error leads to a higher mean level of the uncertainty proxy and therefore a low coefficient of variation of around the 12% seen in the data.

Evidently, large estimated underlying uncertainty jumps of around 310% upon impact of an uncertainty shock feed into the more muted variability of our uncertainty proxy because of the temporal mismatch of inputs and outputs as well as the large contribution of measurement error.

¹⁸Note that the use of a constant returns to scale specification of costs shares within the productivity measurement, while underlying productivity in the model feeds into a decreasing returns to scale production function, is also a form of misspecification that contributes to differences between fundamental and underlying shock dispersion. We have found that this distinction makes little difference in practice for the cross-sectional dispersion of measured TFP shocks in the model. Furthermore, the symmetric constant returns treatment of the model and data moments in this respect is appropriate from an econometric standpoint.

D Online Appendix: A Representative Firm Model

In this section we lay out the structure of a simple and purely illustrative representative agent and firm model, an adaptation of Brock and Mirman (1972), which we use in the main text to investigate the declining investment path in the medium term after an uncertainty shock. We also lay out a purely illustrative calibration of this model and compute the effect of a capital destruction shock in this framework.

D.1 Model Details

Output is produced using a single capital input K_t with one-period time to build in investment. Production is subject to stochastic productivity shocks A_t . A representative household has preferences over consumption in terms of final output. The equilibrium of this neoclassical economy can be derived as a solution to the following planner problem:

$$\begin{aligned} \max_{\{I_t\}_t} \mathbb{E} \sum_{t=0}^{\infty} \rho^t \log(C_t) \\ C_t + I_t &= A_t K_t^\alpha \\ K_{t+1} &= I_t + (1 - \delta)K_t \\ \log(A_{t+1}) &= \rho_A \log(A_t) + \sigma_A \varepsilon_{t+1} \end{aligned}$$

Model Calibration and Solution

We choose illustrative parameters for our model, which we solve numerically at an annual frequency. The capital depreciation rate is $\delta = 0.09$, the curvature of the aggregate technology is $\alpha = 1/3$, the subjective discount rate is $\rho = 0.95$, the autocorrelation of aggregate productivity in logs is 0.9, and the standard deviation of aggregate productivity shocks is $\sigma_A = 0.025$. We solve the model using discretization of the aggregate capital state with $n_K = 350$ grid points, Howard policy iteration, and a discretization of the aggregate productivity process A_t following Tauchen (1986) with $n_A = 25$ grid points.

A Capital Destruction Shock Experiment

We compute the effects of an unanticipated capital destruction shock in this economy. In broad terms, this experiment is meant to correspond to the medium-term aftermath of an uncertainty shock, after uncertainty levels have subsided to a large degree but the economy is left with a smaller capital stock which must be rebuilt through investment.

To implement this experiment, we simulate 75000 independent economies of 100-year length. Note that the drastically reduced computational burden in this simple model allows for a much larger number of independent simulations than in our full heterogeneous agents structure with uncertainty shocks in the main text. In period 25 in each economy, we impose an exogenous and unanticipated movement to a capital grid point K^* , afterwards allowing each economy to evolve normally. We choose the capital grid point K^* , lower than the mean of the capital stock in the ergodic distribution of the model, to ensure that on average this shock results in an approximately -10% reduction in aggregate capital relative to its pre-shock mean.

The Figure D1 in this appendix normalizes the shock period to 1 for plotting purposes and shows, for each indicated variable X , the percentage deviation of the cross-economy mean of X from its pre-shock level, i.e.

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

Here we have that $\bar{X}_t = \frac{1}{75000} \sum_i X_{it}$ is the cross-economy mean of the aggregate X across all economies $i = 1, \dots, 75000$ in period t .

Appendix Table A1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent Variable:	S.D. of log(TFP) shock	Skewness of log(TFP) shock	Kurtosis of log(TFP) shock	IQR of log(TFP) shock	IQR of output growth	S.D. of log(TFP) shock	Skewness of log(TFP) shock	Kurtosis of log(TFP) shock	IQR of log(TFP) shock	IQR of output growth
Sample:	Establishments (manufacturing) in the sample 2 years or more					Establishments (manufacturing) in the sample 38 years				
Recession	0.046*** (0.014)	-0.257* (0.154)	0.696 (1.242)	0.039*** (0.015)	0.078*** (0.019)	0.075*** (0.015)	-0.233 (0.309)	-2.542 (3.340)	0.067*** (0.021)	0.078*** (0.020)
Mean of Dep. Variable:	0.494	-1.454	18.406	0.384	0.226	0.500	-1.541	22.188	0.392	0.185
Corr. with GDP growth	-0.366**	0.091	0.029	-0.380**	-0.504***	-0.461***	0.174	-0.013	-0.419***	-0.553***
Frequency	Annual	Annual	Annual	Annual	Annual	Annual	Annual	Annual	Annual	Annual
Years	1972-2011	1972-2011	1972-2011	1972-2011	1972-2011	1972-2011	1972-2011	1972-2011	1972-2011	1972-2011
Observations	39	39	39	39	39	39	39	39	39	39
Underlying sample	1,390,212	1,390,212	1,390,212	1,390,212	1,390,212	130,201	130,201	130,201	130,201	130,201

Notes: Each column reports a time-series OLS regression point estimate (and standard error below in parentheses) of a measure of uncertainty on a recession indicator. The recession indicator is the share of quarters in that year in a recession. Recessions are defined using the NBER data. In the bottom panel we report the mean of the dependent variable and its correlation with real GDP growth. In columns (1) to (5) the sample is the population of manufacturing establishments with 2 years or more of observations in the ASM or CM survey between 1972 and 2009, which contains data on 211,939 establishments across 39 years of data (one more year than the 38 years of regression data since we need lagged TFP to generate a TFP shock measure). In columns (6) to (10) the sample is the population of manufacturing establishments that appear in the 38 years (1972-2009) of the ASM or CM survey, which contains data on 3,449 establishments. In columns (1) and (6) the dependent variable is the cross-sectional standard deviation (S.D.) of the establishment-level ‘shock’ to Total Factor Productivity (TFP). This ‘shock’ is calculated as the residual from the regression of log(TFP) at year t+1 on its lagged value (year t), a full set of year dummies and establishment fixed effects. In columns (2) and (7) we use the cross-sectional skewness of the TFP ‘shock’, in columns (3) and (8) the cross-sectional kurtosis and in columns (4) and (9) the cross-sectional interquartile range of this TFP ‘shock’ as an outlier robust measure. In columns (5) and (10) the dependent variable is the interquartile range of plants’ sales growth. All regressions include a time trend and Census year dummies (for Census year and for 3 lags). Robust standard errors are applied in all columns to control for any potential serial correlation. *** denotes 1% significance, ** 5% significance and * 10% significance. Data available online at <http://www.stanford.edu/~nbloom/RUBC.zip>

Table A2: Estimates of AR coefficient for TFP Forecast Regressions for Calculation of M.E. Variance

	(1)	(2)	(3)
Dependent variable: $\log(\text{TFP}_{t+1})$	OLS	IV	IV
$\log(\text{TFP}_t)$	0.795*** (0.002)	0.906*** (0.002)	0.932*** (0.002)
Observations	413,401	413,401	413,401
Instrument	N/A	$\log(\text{TFP}_{t-1})$	$\log(\text{TFP}_{t-2})$

Notes: The dependent variable is $\log(\text{TFP})$ at the establishment level at year $t+1$. The right hand variable is $\log(\text{TFP})$ at t . The regression sample are manufacturing establishments with 25 years or more of observations in the ASM or CM survey between 1972 and 2009, which also have non-missing $\log(\text{TFP})$ value for $t-1$ and $t-2$. All columns have a full set of establishment and year fixed effects. In column (1) we use OLS regression. In column (2) we instrument for $\log(\text{TFP}_t)$ with $\log(\text{TFP}_{t-1})$, and in column (3) we instrument with $\log(\text{TFP}_{t-2})$. Standard errors are reported in brackets below every point estimate. *** denotes 1% significance, ** 5% significance and * 10% significance.

Table B1: Internal Accuracy Statistics for the Approximate Equilibrium Mappings in the Baseline Model

Aggregate State Position (A,S,S ₋₁)	Capital log(K _{t+1})		Price log(p _t)	
	RMSE (%)	R ²	RMSE (%)	R ²
(1,0,0)	0.38	0.99	0.63	0.95
(1,0,1)	0.20	1.00	0.09	1.00
(1,1,0)	0.73	0.98	0.29	0.99
(1,1,1)	0.49	0.99	0.29	0.98
(2,0,0)	0.51	0.99	0.58	0.96
(2,0,1)	0.42	0.99	0.08	1.00
(2,1,0)	0.34	1.00	0.23	0.99
(2,1,1)	0.41	0.99	0.30	0.98
(3,0,0)	0.42	0.99	0.50	0.97
(3,0,1)	0.43	0.99	0.08	1.00
(3,1,0)	0.52	0.99	0.23	0.99
(3,1,1)	0.55	0.98	0.32	0.97
(4,0,0)	0.40	0.99	0.45	0.98
(4,0,1)	0.56	0.94	0.08	1.00
(4,1,0)	0.36	0.99	0.22	0.99
(4,1,1)	0.39	0.98	0.31	0.96
(5,0,0)	0.44	0.99	0.51	0.97
(5,0,1)	0.26	1.00	0.07	1.00
(5,1,0)	0.44	0.99	0.09	1.00
(5,1,1)	0.32	0.99	0.27	0.97


Notes: Approximate equilibrium forecast mappings for market-clearing prices are $\log(p_t) = \alpha_p(A_t, S_t, S_{t-1}) + \beta_p(A_t, S_t, S_{t-1})\log(K_t)$ and for next period's capital are $\log(K_{t+1}) = \alpha_K(A_t, S_t, S_{t-1}) + \beta_K(A_t, S_t, S_{t-1})\log(K_t)$. The accuracy statistics above are computed from an unconditional simulation of 5000 quarters of the baseline model, discarding an initial 500 quarters. Each row in the table above displays the performance of the equilibrium mapping conditional upon a subsample of the data characterized by a given triplet of discretized grid points for aggregate productivity A_t , uncertainty S_t , and lagged uncertainty S_{t-1} . RMSE represents the root mean squared error of the indicated rule's static or one-period ahead forecasts, and the R^2 measure is the standard R^2 measure computed from the log-linear regression on the appropriate subsample of simulated data.

Table B2: Marginal Utility Price Forecasting Accuracy, Alternative Computational Strategies

Den Haan Statistics (%)	Stat.\Horizon	3 Years	4 Years	5 Years	12 Years
Baseline: Lagged Unc., Convexified Clearing	Mean	0.63	0.75	0.84	1.25
	Max	1.29	1.54	1.76	2.80
Extra Lags of Uncertainty Convexified Clearing	Mean	0.64	0.75	0.84	1.16
	Max	1.27	1.51	1.72	2.56
Extra Forecast Moment Convexified Clearing	Mean	0.61	0.73	0.83	1.23
	Max	1.53	1.87	2.15	3.27
One Lag of Uncertainty Nonconvexified Clearing	Mean	0.43	0.49	0.53	0.68
	Max	0.97	1.13	1.25	1.79


Notes: The table above reports the Den Haan (2010) accuracy statistics “Den Haan Statistics” for the forecasting system for market-clearing marginal utility prices p . Den Haan statistics are based on forward iteration of the forecasting system for marginal utility out to a pre-specified horizon conditioning only on exogenous processes, substituting s -period ahead forecasts as inputs for $s+1$ -period ahead forecasts and so on. “Mean” and “Max” refer to the mean and maximum of the approximate percentage errors $100 | \log(p_t) - \log(p_t^{DH}) |$ resulting from these iterative forecasts. The table reports averages of these error statistics using 2,000 forecast start dates in an unconditional simulation of the model. The statistics are calculated for each of four alternative model solution strategies, denoted in the first column.

Figure A1: 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: Macro volatility calculated from a GARCH(1,1) model estimated from aggregate TFP growth




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

Figure B1: Convexified excess demand leads to market clearing with essentially arbitrary accuracy



Notes: The figure above plots two series from an unconditional simulation of the model for 5000 quarters, after discarding an initial 500 quarters. The top panel plots the market clearing level of consumption C_t relying on convexification of the excess demand function to clear goods markets. The bottom panel plots the percentage clearing error, i.e. the percentage difference between implied consumption C_t and the reciprocal of the marginal utility price p_t governing firm investment choices, along with the numerical clearing bounds in dotted lines. Over the unconditional simulation of the model, the maximum clearing error is 0.0099%, and the mean error is 0.00032%.


Figure B2: The size of uncertainty's effect varies little with alternative forecasting and market clearing algorithms




Notes: Based on independent simulations of 2500 economies of 100-quarter length. We impose an uncertainty shock in the quarter labelled 1. The baseline economy (x symbols) uses a forecast rule for market clearing prices p , i.e. marginal utility, including lagged uncertainty as a state variable, and market clearing consumption is obtained by clearing a convexified excess demand function. The extra lags, convexified economy (* symbols) includes three lags of uncertainty as state variables and clears markets with a convexified excess demand function. The extra moment, convexified economy (diamonds) adds the cross-sectional covariance of capital and productivity to the forecast rule and clears markets with a convexified excess demand function. The one lag, nonconvexified economy (o symbols) includes lagged uncertainty as a state variable and clears markets with a nonconvexified excess demand function. From the left, we plot the percent deviations of cross-economy average output, consumption, and the price forecasting error (marginal utility price p – forecast price p) from quarter 0.

Figure B3: Alternative measures of misallocation comove


Corr (Productivity, Labor)




Corr (Productivity, Capital)




Marginal Product Dispersion, Labor




Marginal Product Dispersion, Capital



Cov(Productivity, Labor)




Cov (Productivity, Capital)



Quarters (uncertainty shock in quarter 1)


Notes: Based on independent simulations of 2500 economies of 100-quarter length. We impose an uncertainty shock in the quarter labelled 1, allowing normal evolution of the economy afterwards. In each panel, we plot the average deviation of an alternative proxy for allocative efficiency from its pre-shock value, in the raw units associated with that series in quarter 0. The top row plots the correlation of productivity and inputs: $\text{corr}(\log k, \log z)$ and $\text{corr}(\log n, \log z)$. The middle row plots the dispersion of the marginal product of inputs: the standard deviations of $(\log y / k)$ and $(\log y / n)$. The bottom row plots the covariance of productivity and inputs: $\text{cov}(\log k, \log z)$ and $\text{cov}(\log n, \log z)$.

Figure B4: The impact of uncertainty is robust to alternative methods of computing impulse responses




Notes: For the baseline economy (x symbols), we simulate 2500 independent economies of 100-period length. We impose an uncertainty shock in about quarter 50, labeled 1 above. Clockwise from the top left, we plot the percent deviations of cross-economy average output, labor, consumption, and investment from their values in quarter 0. For the simulation differencing case (► symbols), we simulate 2500 pairs of independent economies. The exogenous shocks to all series are identical within a pair. However, we impose an uncertainty shock in one of the two economies within a pair, allowing all other series and outcomes to evolve normally afterwards. In each panel, the figure plots the average across economy pairs of the percent deviation of the indicated series in the shocked economy versus the no shock economy. The parameters are identical across the two impulse responses plotted above, equal to the baseline estimates.

Figure B5: The impact of uncertainty is robust to using a denser grid for productivity shocks




Notes: Based on independent simulations 100-quarter length. For each path, we impose an uncertainty shock in quarter 1. For the baseline economy (x symbols), we use a grid with 5 points for macro productivity A and micro productivity z. For the denser grid economy (▶ symbols), we use a grid with 9 points for macro productivity A and micro productivity z. The parameterizations are identical in each case and reflect the baseline estimates. Clockwise from the top left, we plot the percent deviations of cross-economy average output, labor, consumption, and investment from their values in quarter 0.

Figure C1: Input observation timing mismatch & measurement error explain lower variation in the model uncertainty proxy



Notes: The IQR series in the left hand panel were drawn from a representative 120-quarter period in the unconditional simulation of the model. “Quarterly” is the underlying quarterly cross-sectional IQR of micro productivity shocks. “Annual” is the IQR of shocks with variance equal to the variance of the sum of the quarterly shocks throughout a given year. “Mismatch” is the IQR of simulated TFP shocks in that year, with TFP computed as in the data using a Solow Residual approach, first-quarter labor, fourth-quarter capital, and the sum of yearly output. “Measurement Error” is the baseline IQR of simulated TFP shocks accounting for measurement error. The bars in the right hand panel are equal to the coefficient of variation of each indicated IQR series from the left hand side, computed over the entire unconditional simulation of the model, along with the coefficient of variation of the annual uncertainty IQR proxy computed from our Census of Manufactures and Annual Survey of Manufactures data from 1972-2010.

Figure D1: The impact of a capital destruction shock in the Brock-Mirman model



Notes: Based on independent simulations of 75000 economies of 100-year length in the simplified and illustrative representative agent Brock-Mirman model as described in Appendix E. In the year labelled 1 in each economy, we impose a capital destruction shock leading to about a 10% loss of capital on average, allowing the economy to evolve normally thereafter. Clockwise from the top left, we plot the percent deviations of cross-economy average output, capital, consumption, and investment from their values in year 0.