Aggregate Risk in Heterogeneous Agent Economies

ECO2101 Lecture 7 | Applications

Sebastian Dyrda, University of Toronto

25, March 2026

This lecture

Last time (Lecture 6): Krusell & Smith (1998) — the distribution as a state variable, bounded-rationality forecasting rules, approximate aggregation, computational algorithm

Today: three applications that test when and why heterogeneity matters

1. Khan & Thomas (*Econometrica*, 2008): Lumpy investment
   $\rightarrow$ Do micro nonconvexities matter for aggregate dynamics?

2. Khan & Thomas (*JPE*, 2013): Credit shocks
   $\rightarrow$ Financial frictions and the role of the distribution

3. Bloom, Floetotto, Jaimovich, Saporta-Eksten & Terry (*Econometrica*, 2018): Uncertainty shocks
   $\rightarrow$ Second-moment shocks and the "wait-and-see" effect

Khan & Thomas (*Econometrica*, 2008)

Idiosyncratic Shocks and the Role of Nonconvexities

in Plant and Aggregate Investment Dynamics

Motivation: plant-level investment is lumpy

- Micro data on establishment investment (LRD, ASM) reveals sharp departures from smooth adjustment:

- Spikes: ≈20% of plants have investment rates $> 20\%$ per year
- Inaction: ≈8–20% of plants have near-zero investment in any given year
- Asymmetry: large positive spikes far more common than large negative ones

- Natural interpretation: nonconvex (fixed) costs of capital adjustment generate (S, s)-style investment rules

- **Key question**: do these micro nonlinearities matter for *aggregate* investment dynamics?
  - $\rightarrow$ Caballero & Engel (1999): "yes" — synchronization of adjustments creates time-varying elasticity
  - $\rightarrow$ Khan & Thomas (2008): "no" — general equilibrium washes them out

The plant's problem: primitives

- **Plant state**: (k, ε, ξ) – capital, idiosyncratic TFP, adjustment cost draw
- **Aggregate state**: (z, μ) – aggregate TFP, distribution of plants over (ε, k)
- **Production**: $y = z\varepsilon F(k, n)$, with F increasing, concave (e.g. Cobb-Douglas)
- **Shocks**: z and ε follow independent Markov chains. Each period, every plant draws $\xi \sim G$ on $[0, \bar{\xi}]$ (iid across plants and time)
- **Capital**: law of motion $\gamma k' = (1-\delta)k + I$, where γ is the growth rate of labor-augmenting technology (all variables detrended)
- **Households**: representative, preferences $U(c, 1-n_h)$. In equilibrium, household optimality pins down the output price p and the real wage ω:

$$p(z,\mu) = D_1U(C, 1-N), \qquad \omega(z,\mu) = \frac{D_2U(C,1-N)}{D_1U(C,1-N)}$$

so ω equals the MRS between leisure and consumption

Investment technology: fixed costs and the constraint set $\Omega(k)$

Each period, every plant draws an adjustment cost $\xi \sim G$ on $[0, \bar{\xi}]$. Two options:

1. Pay fixed cost $\xi \cdot \omega$ (in labor units) $\Rightarrow$ choose *any* $k' \in \mathbb{R}_+$
2. Avoid the cost $\Rightarrow$ restricted to $k' \in \Omega(k)$, where

$$\Omega(k) \equiv \left[\frac{1-\delta+a}{\gamma}k, \ \frac{1-\delta+b}{\gamma}k\right], \qquad a < 0 < b$$

- Routine (small) adjustments are free; major reconfigurations require paying ξ
- Bounds scale with k (necessary for balanced growth; friction is size-neutral)
- a, b calibrated to match the fraction of small vs. large adjustments in the LRD
- Without $\Omega(k)$ the model predicts too much inaction — the problem in KT (2003)

The plant's problem: Bellman equation

Value of a plant with (ε, k) drawing adjustment cost ξ, given aggregate state (z, μ):

$$V^1(\varepsilon, k, \xi;\, z, \mu) = \max_{n \geq 0}\, p \cdot \left[z\varepsilon F(k, n) - \omega\, n + (1-\delta)k\right]$$
$$+ \max\left\{ \underbrace{-\xi\,\omega\, p + \max_{k' \geq 0}\, R(\varepsilon, k';\, z, \mu')}_{\text{pay } \xi\text{, adjust to any } k'},\quad \underbrace{\max_{k' \in \Omega(k)}\, R(\varepsilon, k';\, z, \mu')}_{\text{avoid cost, constrained } k'} \right\}$$

The **continuation value** of entering next period with capital k':

$$R(\varepsilon, k';\, z, \mu') = -\gamma k' p(z, \mu) + \beta\, \mathbb{E}\left[\frac{p(z', \mu')}{p(z, \mu)} \sum_{\varepsilon'} \Pi(\varepsilon'|\varepsilon)\, V^0(\varepsilon', k';\, z', \mu') \,\middle|\, z\right]$$

- First term: cost of purchasing k' units of capital at current output price p
- Second term: discounted expected value next period, using the household's SDF $\beta\, p'/p$

The plant's problem: Bellman equation

The **ex-ante value** (before the adjustment cost draw) integrates over ξ:

$$V^0(\varepsilon, k;\ z, \mu) = \int_0^{\bar{\xi}} V^1(\varepsilon, k, \xi;\ z, \mu)\, G(d\xi)$$

Note: n and k' choices are *separable*; $\mu' = \Gamma(z, \mu)$ is the equilibrium law of motion

Key objects: target capital and adjustment threshold

Define the **gross value of unconstrained** and **constrained adjustment**:

$$E(\varepsilon, z, \mu') \equiv \max_{k' \geq 0} \; R(\varepsilon, k'; \, z, \mu') \qquad \text{— choose any } k'$$

$$E^c(\varepsilon, k, z, \mu') \equiv \max_{k' \in \Omega(k)} \; R(\varepsilon, k'; \, z, \mu') \qquad \text{— constrained to } \Omega(k)$$

- The solution to the unconstrained problem: target capital $k^*(\varepsilon, z, \mu)$
  - $\rightarrow$ Independent of current k and ξ — depends on ε (through persistence of idiosyncratic TFP) and on aggregate conditions (z, μ)
  - $\rightarrow$ All adjusting plants with the same ε jump to the *same* k^*

Key objects: target capital and adjustment threshold

- **Threshold cost** $\xi^T(\varepsilon, k;\ z, \mu)$: the ξ that makes a (ε, k)-plant indifferent:

$$\xi^T(\varepsilon, k;\ z, \mu) = \frac{E(\varepsilon, z, \mu') - E^c(\varepsilon, k, z, \mu')}{p(z, \mu) \cdot \omega(z, \mu)}$$

- Plant pays cost and adjusts to k^* iff $\xi \leq \xi^T(\varepsilon, k; z, \mu)$.
- Plants *far* from target: large gap $E - E^c \Rightarrow$ high $\xi^T \Rightarrow$ likely to adjust
- Plants *near* target: small gap $\Rightarrow$ low $\xi^T \Rightarrow$ constrained adjustment (or target $k^* \in \Omega(k)$, so adjustment is free)

Generalized (S, s) investment rule

Combining the threshold and target, the **decision rule for capital** is:

$$k' = K(\varepsilon, k, \xi; z, \mu) = \begin{cases} k^*(\varepsilon, z, \mu) & \text{if } \xi \leq \xi^T(\varepsilon, k; z, \mu) \\ k^c(\varepsilon, k, z, \mu) & \text{if } \xi > \xi^T(\varepsilon, k; z, \mu) \end{cases}$$

where $k^c(\varepsilon, k, z, \mu) = \arg\max_{k' \in \Omega(k)} R(\varepsilon, k'; z, \mu')$ is the best constrained choice.

- This is a two-sided generalized (S, s) rule: the threshold ξ^T depends on the distance between current k and the target k^*, in *either* direction
- Within each (ε, k) group, fraction $G\big(\xi^T(\varepsilon, k; z, \mu)\big)$ of plants adjust
- **Two margins of aggregate investment response to shocks**:
  - $\rightarrow$ *Intensive margin*: how much do adjusting plants change k? (shift in k^*)
  - $\rightarrow$ *Extensive margin*: how many plants adjust? (shift in ξ^T and hence in $G(\xi^T)$)
- The central question: does the extensive margin generate aggregate nonlinearities absent from the representative-firm RBC model?

The (S, s) rule in action: adjustment hazard and distribution

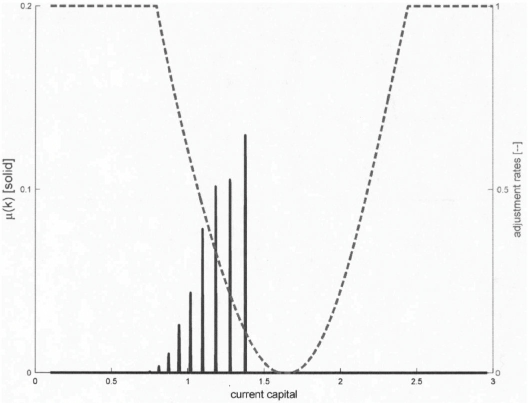

FIGURE 3.—Steady-state adjustment in the common productivity lumpy model.

Dashed line: adjustment rate by k (right axis). Bars: distribution of plants over k (left axis). Target $k^* \approx 1.38$. Hazard rises with distance from target — this *is* the (S, s) rule.

Aggregate state and connection to Krusell-Smith

- Aggregate state: (z, μ) where μ is the joint distribution of plants over (ε, k)
- Law of motion: $\mu' = \Gamma(z, \mu)$ – a mapping from distributions to distributions

- **KS approximation**: agents use the mean of capital $m = \int k \, d\mu$ as proxy for μ. Two forecasting rules, conditional on z:

$$\begin{aligned} \log m' &= \beta_0(z) + \beta_1(z) \log m \\ \log p &= \gamma_0(z) + \gamma_1(z) \log m \end{aligned}$$

- **Preference simplification**: Hansen-Rogerson indivisible labor $u(c, L) = \log c - \phi L \Rightarrow \omega = \phi/p$, so the wage is known once p is forecasted – no separate rule for ω needed

- Approximate aggregation holds: only m (the mean of k) is needed; higher moments add negligible forecasting power. $R^2 > 0.999$ for both rules

The mechanism: why GE kills the nonlinearities

Partial equilibrium vs. general equilibrium

- **Partial equilibrium** (prices fixed at steady-state values):
  - $\rightarrow$ Positive TFP shock $\Rightarrow$ target capital k^* jumps up for all ε types
  - $\rightarrow$ Many plants cross their adjustment threshold $\Rightarrow$ extensive margin surge
  - $\rightarrow$ Large, nonlinear response of aggregate investment (high skewness, kurtosis)
- **General equilibrium** (prices clear markets):
  - $\rightarrow$ Same TFP shock, but now wages ω rise and interest rates increase
  - $\rightarrow$ *Intensive margin dampened*: higher ω and higher interest rate reduce the present value of investing $\Rightarrow$ target k^* rises much *less* than in PE
  - $\rightarrow$ *Extensive margin dampened*: (i) smaller shift in k^* means fewer plants are far from target, and (ii) fixed cost $\xi \cdot \omega$ is more expensive at higher ω
  - $\rightarrow$ Both effects reduce ξ^T $\Rightarrow$ fewer plants cross the threshold
  - $\rightarrow$ Result: lumpy economy $\approx$ frictionless economy at the aggregate level
- GE dampens the microelasticity of plant investment to aggregate shocks by a factor of $\approx$**13**

How the extensive margin responds to shocks

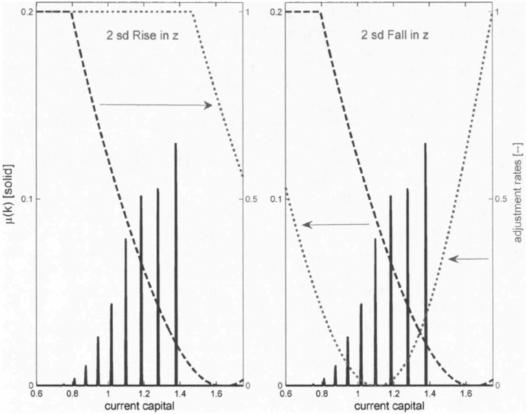

FIGURE 4.—Adjustment responses in the common productivity lumpy model.

Left: positive TFP shock — target shifts up, hazard shifts, more plants adjust.
Right: negative TFP shock — asymmetric response in adjustment rates.
In GE, these shifts are dramatically dampened by price movements.

How the extensive margin responds to shocks

- The household's preference for consumption smoothing restrains equilibrium price movements
- Aggregate investment achieves through modest movements along *two* margins (intensive + extensive) what the frictionless model achieves through the intensive margin alone

Main results

1. **Micro fit**: model matches the cross-sectional distribution of plant investment rates (spikes, inaction, asymmetry) — first model to do so
2. **PE vs. GE**:
   → In PE: nonconvexities generate large aggregate nonlinearities (high skewness and kurtosis of aggregate investment)
   → In GE: aggregate dynamics are virtually identical to the frictionless representative-firm model
3. **Approximate aggregation**: $R^2 > 0.999$ for the log-linear forecasting rule, as in the household-side KS model
4. **Implication for estimation**: partial equilibrium estimation of adjustment costs overstates the upper support of the fixed cost distribution by a factor of ≈5

Results: investment rates

TABLE II

DISTRIBUTION OF PLANT INVESTMENT RATES

	Inaction	Positive Spike	Negative Spike	Positive Invest.	Negative Invest.
Establishment data[a]	*0.081*	*0.186*	*0.018*	*0.815*	*0.104*
Lumpy investment models[b]					
Traditional model	0.789	0.187	0.000	0.211	0.000
Extended model	0.073	0.185	0.010	0.752	0.175

[a] Data are from Cooper and Haltiwanger (2006). Inaction, $|i/k| < 0.01$; positive spike, $i/k > 0.20$; negative spike, $i/k < -0.20$; positive investment, $i/k \geq 0.01$; negative investment, $i/k \leq -0.01$.

[b] The traditional lumpy model has $\sigma_{\eta\varepsilon} = b = 0$ and upper support $= 0.014$ to fit positive spike observations. The extended model includes exempted range of investment rates and plant-specific shocks to productivity. Parameters are listed in Table I.

Results: aggregate investment rates in GE vs. PE

412 A. KHAN AND J. K. THOMAS

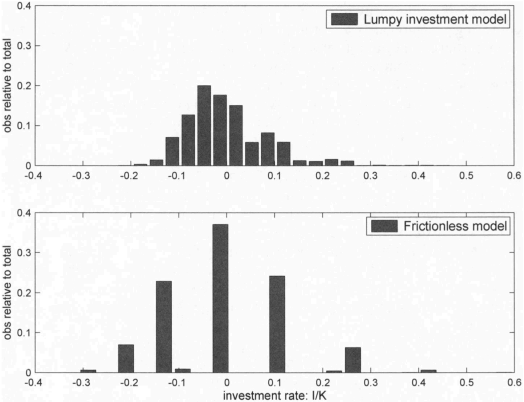

FIGURE 1.—Distribution of aggregate investment rates in partial equilibrium.

PE: wide, right-skewed distribution due to extensive margin nonlinearities.

Results: aggregate investment rates in GE vs. PE

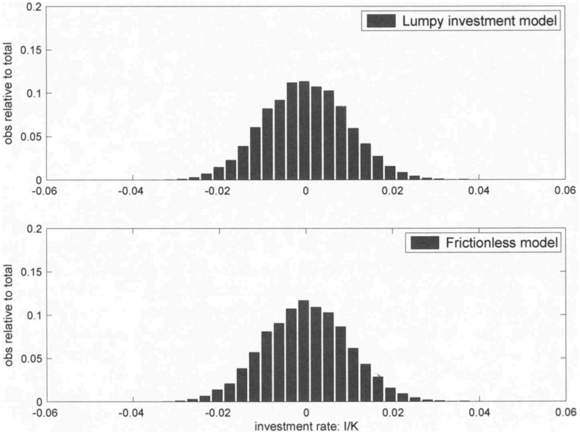

FIGURE 5.—Distribution of aggregate investment rates in general equilibrium.

GE: range compressed $\approx$ 10x, nearly identical to frictionless model.

Lessons from Khan & Thomas (2008)

Lesson 1: Micro nonconvexities can be *quantitatively irrelevant* for aggregate dynamics once general equilibrium effects are accounted for.

Lesson 2: Approximate aggregation holds — the Krusell-Smith approach works here. The distribution of firms across (k, ε) carries little information beyond aggregate K for forecasting prices.

Lesson 3: General equilibrium analysis is essential even for understanding *micro*-level investment dynamics. Partial equilibrium overstates both aggregate and plant-level responses.

But: this result depends on the *type of shock* (first-moment TFP) and the *type of friction* (real adjustment costs). What if we change either?
$\Rightarrow$ Applications II and III

Khan & Thomas (*JPE*, 2013)

Credit Shocks and Aggregate Fluctuations

in an Economy with Production Heterogeneity

Motivation: the 2007–2009 recession

- The Great Recession featured a sharp contraction in lending alongside the output collapse
- Can a financial shock alone — without any exogenous TFP decline — generate a large, persistent recession?
- Key features of the 2007–09 downturn that are hard to match with a standard TFP shock:
  - $\rightarrow$ GDP decline was *gradual*, not sudden; consumption initially rose
  - $\rightarrow$ Investment fell much more steeply than output
  - $\rightarrow$ Measured TFP declined *endogenously* — not driven by technology
  - $\rightarrow$ Employment losses concentrated among *small* firms
- KT13: a tightening of collateral constraints disrupts the *allocation* of capital across firms $\Rightarrow$ endogenous decline in aggregate TFP $\Rightarrow$ persistent recession

Motivation: the 2007–2009 recession (data)

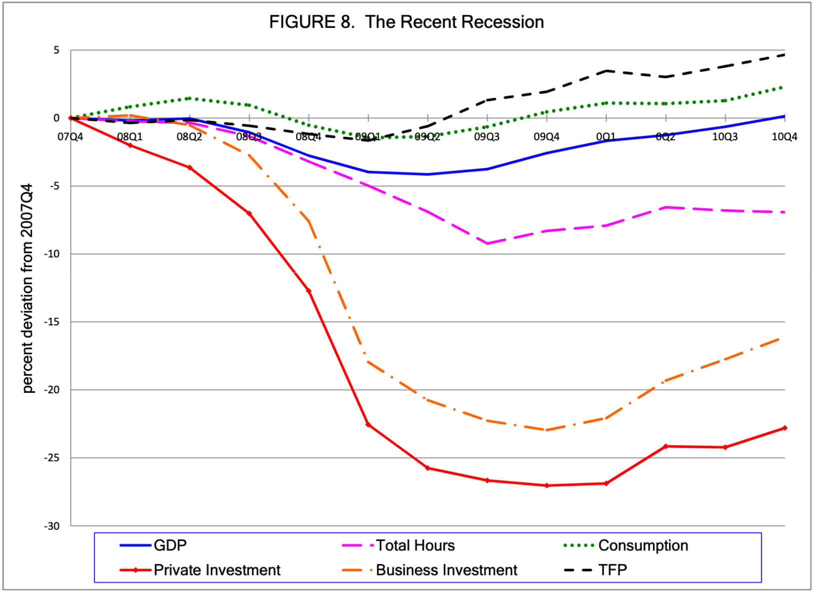

The firm's problem: primitives

- **Firm state**: (k, b, ε) – capital, one-period debt, idiosyncratic TFP
- **Aggregate state**: (s, μ) where $s = (z, \zeta)$ combines aggregate TFP z and a credit variable ζ; μ is the distribution of firms over (k, b, ε)
- **Production**: $y = z\varepsilon F(k, n)$, decreasing returns to scale
- **Two frictions**:
  - → Collateral constraint: $b' \leq \zeta\, \nu_k\, k$, where $\nu_k \in [0, 1]$ is the fraction of capital recoverable upon uninstalling, and ζ governs how much of that collateral firms can borrow against
  - → Partial irreversibility: selling capital recovers only $\nu_k < 1$ per unit $\Rightarrow$ generates (S, s) investment rules (as in KT08, but without iid fixed costs)
- **Entry/exit**: each period fraction p_d of firms exit (replaced by entrants with zero debt and low capital k_0). Entry of small, initially constrained firms ensures that borrowing limits remain relevant in the cross-section even though individual firms eventually outgrow them

The firm's problem: Bellman equation

Ex-ante value (before learning survival):

$$V^0(k,b,\varepsilon;\, s,\mu) = p_d \max_n \; p\big[z\varepsilon F(k,n) - \omega n + \nu_k(1-\delta)k - b\big] + (1-p_d)\, V(k,b,\varepsilon;\, s,\mu)$$

Continuing firm chooses between **upward** and **downward** capital adjustment:

$$V(k,b,\varepsilon;\, s,\mu) = \max\left\{V^u(k,b,\varepsilon;\, s,\mu),\;\; V^d(k,b,\varepsilon;\, s,\mu)\right\}$$

Upward adjustment ($k' \geq (1-\delta)k$, investment price $= 1$):

$$\begin{aligned} V^u = \max_{n,k',b',D} \quad & p \cdot D + \beta \sum_{s'} \pi_{ss'} \sum_{\varepsilon'} \Pi(\varepsilon'|\varepsilon)\, V^0(k',b',\varepsilon';\, s',\mu') \\ \text{s.t.} \quad & D = z\varepsilon F(k,n) - \omega n + q\, b' - b - [k' - (1-\delta)k] \\ & k' \geq (1-\delta)k, \quad b' \leq \zeta\, \nu_k\, k, \quad D \geq 0 \end{aligned}$$

The firm's problem: Bellman equation

Downward adjustment ($k' \leq (1-\delta)k$, resale price $= \nu_k < 1$): same but $[k' - (1-\delta)k]$ is replaced by $\nu_k[k' - (1-\delta)k]$ in the budget constraint

Note: this is the reformulated problem where V is in utils (multiplied by p), so firms discount by β rather than by the SDF. The wage ω and bond price q are as in eqs. (7)–(8) of the paper.

Key differences from KT08

- **Three-dimensional heterogeneity**: firms differ in (k, b, ε), not just (k, ε)
  - $\rightarrow$ Debt b is a *genuine* state variable: it determines how constrained the firm is
  - $\rightarrow$ Cannot be summarized by "cash on hand" because capital and debt have different roles (capital is collateral, debt is a liability)
- **Partial irreversibility** replaces fixed adjustment costs
  - $\rightarrow$ Selling capital at $\nu_k < 1$ generates (S, s) bands endogenously
  - $\rightarrow$ Firms with too much k shed only to the upper edge of the inaction band; firms with too little k invest only to the lower edge
- **Collateral constraint** $b' \leq \zeta\, \nu_k\, k$: borrowing limited by *existing* capital
  - $\rightarrow$ Firms can relax the constraint by accumulating capital over time $\Rightarrow$ life-cycle dynamics: young firms are constrained, old firms outgrow the friction
  - $\rightarrow$ The credit shock is a drop in ζ: lenders accept less collateral
- **Dividends must be non-negative**: $D \geq 0$. Firms cannot raise equity, so internal funds (retained earnings) are the only buffer against tighter credit

Life-cycle dynamics: how firms outgrow the constraint

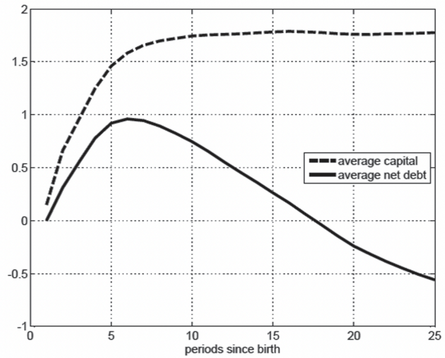

FIG. 2.—Cohort in steady state. Cohort average capital and net debt are taken from an unbalanced panel of firms born at date 1. Net debt is defined as debt less financial savings.

Young firms borrow and accumulate k; by period $\sim$7, debt peaks and firms start deleveraging; by period $\sim$16, the typical firm is a net saver and permanently unconstrained.

This *is* the policy function for (k, b) projected onto the age dimension.

Misallocation: expected return to capital by firm age

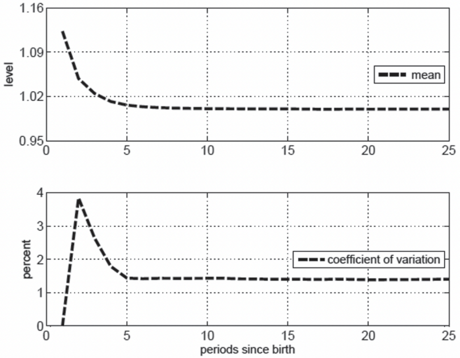

FIG. 3.—Cohort expected discounted return to capital in steady state. The top panel is period-by-period mean expected discounted marginal value of capital for an unbalanced panel of firms born in period 1. The bottom panel is the coefficient of variation in expected return from the same sample.

Top: mean return starts at 1.12 for young firms (constrained $\Rightarrow$ high marginal value of k), converges to 1.0 as firms age out of the friction.
Bottom: CV of returns falls with age — dispersion = misallocation.
Without frictions, both panels would be flat at 1.0 and 0.0.]

Aggregate state and connection to Krusell-Smith

- Aggregate state: (s, μ) where $s = (z, \zeta)$ and μ is the distribution over (k, b, ε)
- Two exogenous aggregate variables (z, ζ) following a joint Markov chain

- **KS approximation**: two forecasting rules, one for capital, one for the output price:

$$\begin{aligned}\log m' &= \beta_0(s) + \beta_1(s)\log m + \beta_2(s)\,c_1 + \beta_3(s)\,c_2 \\ \log p &= \gamma_0(s) + \gamma_1(s)\log m + \gamma_2(s)\,c_1 + \gamma_3(s)\,c_2\end{aligned}$$

  where $m = \int k\, d\mu$ and c_1, c_2 are crisis indicator variables ($c_j = \mathbf{1}(\zeta_{t-j} = \zeta^L)$).
  Coefficients are state-contingent: 6 regressions each ($N_s = 6$ aggregate states)

- More regressors than in KT08: the history of credit conditions (c_1, c_2) is needed because past ζ shapes the current distribution of net worth
- Forecasting accuracy remains high (Den Haan max errors $\approx$ KS 1998 benchmarks)

Limits of the Krusell-Smith approach

- KS works here — but at a cost:
  - $\rightarrow$ Which moments of the (k, b, ε) distribution should enter the forecasting rule?
  - $\rightarrow$ Choice is *ad hoc*, found by trial and error, model-specific
  - $\rightarrow$ No guarantee that moments that work in one calibration transfer to another
- As models grow in complexity (more state variables, more shocks), the number of candidate moments grows combinatorially
- This is a fundamental limitation of the KS algorithm
- $\Rightarrow$ Motivates the sequence-space Jacobian methods (Lectures 8–9), which solve for equilibrium without ever choosing which moments to track

The mechanism: how credit shocks cause recessions

1. ζ drops $\Rightarrow$ collateral constraint $b' \leq \zeta \nu_k k$ binds for more firms
2. Constrained (young/small) firms cannot reach their efficient k; unconstrained (old/large) firms expand, exploiting the lower real interest rate
3. Capital misallocation worsens $\Rightarrow$ endogenous decline in aggregate TFP — same production functions, same z, but k and ε become less correlated
4. Nonnegative-dividend constraint ($D \geq 0$) and tight collateral slow rebuilding of net worth $\Rightarrow$ **persistent** recession even after ζ returns to normal
5. Contrast with KT08: the constraint $b' \leq \zeta \nu_k k$ is a *quantity limit* — GE price adjustment (lower r) cannot relax it

Capital misallocation: how the distribution shifts after a credit shock

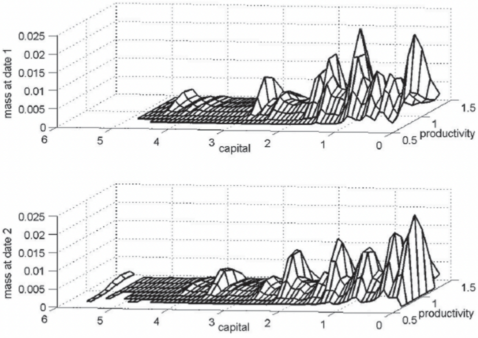

FIG. 7.—Persistent financial crisis: capital-productivity distribution. Response to a drop in the credit variable from ζ_o to ζ_l. Expectations are consistent with the calibrated shock process. The top panel is the predetermined distribution of firms over capital and productivity at the date of the shock; the bottom panel is the same distribution at the start of the next period. The *y*-axes measure the mass at firms at each (k, ε) combination.

Top: predetermined distribution at date of shock.

Bottom: distribution one period later — more small firms, fewer medium firms, largest (unconstrained) firms grow even larger exploiting low interest rates.

The gap between constrained and unconstrained firms *widens* $\Rightarrow$ misallocation rises.

Results: Credit shock

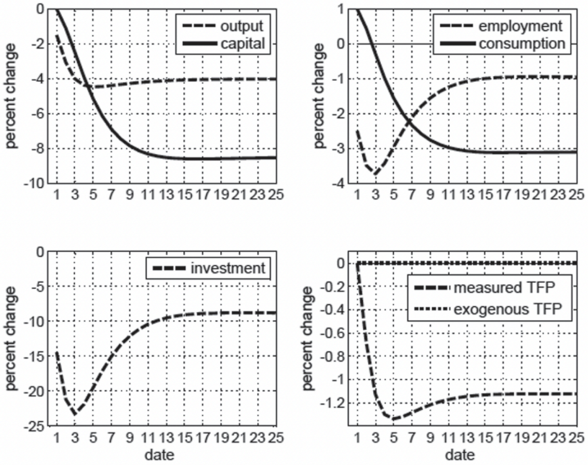

FIG. 6.—Persistent financial crisis. Response to a drop in the credit variable from ζ_o to ζ_l. Expectations are consistent with the calibrated shock process; however, the credit variable is maintained at its crisis level through period 25. The y-axes measure percentage deviations from simulation means.

Results: Credit shock

TABLE 5
PEAK-TO-TROUGH DECLINES: US 2007 RECESSION AND MODEL

	Trough	GDP	I	N	C	TFP	Debt
Data	2009Q2	5.59	18.98	6.03	4.08	2.18	25.94
Four-period crisis	4	4.38	21.88	3.42	.99	1.30	25.96
z-shock (ρ_z)	1	3.16	11.63	1.66	1.52	2.18	.24
ζ-shock (ρ_z)	4	3.80	18.40	2.85	.97	1.09	22.14

NOTE.—Debt data entry is the fall in real stock of commercial and industrial loans at commercial banks between 2008Q4 and 2011Q4; all other series are measured at the GDP trough date listed in col. 2. Excluding debt, data in row 1 are described in fig. 5. Row 2 is the four-period credit crisis described in the text; results here are the same whether the financial recovery starting in date 5 is immediate or gradual. Row 3 is a negative productivity shock with persistence ρ_z and date 1 value chosen to match the empirical drop in measured TFP. Row 4 is a credit shock with persistence ρ_z and date 1 value chosen to match the empirical drop in debt.

Main results

1. Credit shock alone $\Rightarrow$ GDP falls $\approx$4.5%, investment collapses, consumption decline is delayed — matching key 2007–09 patterns
2. Misallocation explains $>$50% of the observed TFP drop — no exogenous technology shock needed
3. Employment losses concentrated among small firms, consistent with Census evidence
4. A standard TFP shock, by contrast, generates weaker declines and *equal* losses across firm sizes

Lessons from Khan & Thomas (2013)

Lesson 1: The distribution of firms is economically central — misallocation of k across ε types drives endogenous TFP declines. KS forecasting still works, but requires more moments, chosen ad hoc. This does not scale $\Rightarrow$ motivates SSJ methods (Lectures 8–9).

Lesson 2: Credit shocks cause “supply-side” recessions without technology regress — a channel absent from standard RBC models.

Lesson 3: Whether heterogeneity matters depends on the *type of friction*:

- Real adjustment costs (KT08): GE prices neutralize $\Rightarrow$ distribution irrelevant
- Financial frictions (KT13): quantity constraint $\Rightarrow$ distribution matters

$\Rightarrow$ Application III: what if we change the *type of shock* instead?

Bloom, Floetotto, Jaimovich, Saporta-Eksten & Terry (*Econometrica*, 2018)

Really Uncertain Business Cycles

Motivation: uncertainty is countercyclical

- Census microdata (1972–2011, $\approx$50,000 manufacturing plants):
- The cross-sectional dispersion of plant-level TFP shocks rises sharply in recessions
  - $\rightarrow$ Great Recession: variance of TFP shocks rose 76%, variance of sales growth rose 152%
  - $\rightarrow$ Pattern is robust across samples, measures, and levels of aggregation
- Same pattern at firm level (Compustat) and industry level (SIC 4-digit)
- Skewness and kurtosis of TFP shocks are *not* significantly cyclical — recessions are a negative first-moment + positive second-moment shock
- **Key question**: is uncertainty a *cause* of recessions, or just a byproduct?

Dispersion of plant-level TFP shocks over the business cycle

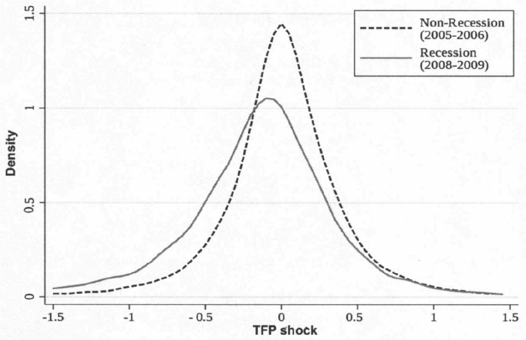

FIGURE 1.—The variance of establishment-level TFP shocks increased by 76% in the Great Recession. Notes: Constructed from the Census of Manufactures and the Annual Survey of Manufactures using a balanced panel of 15,752 establishments active in 2005–2006 and 2008–2009. TFP Shocks are defined as residuals from a plant-level log(TFP) AR(1) regression that also includes plant and year fixed effects. Moments of the distribution for non-recession (recession) years are: mean 0 (−0.166), variance 0.198 (0.349), coefficient of skewness −1.060 (−1.340) and kurtosis 15.01 (11.96). The year 2007 is omitted because according to the NBER the recession began in 12/2007, so 2007 is not a clean "before" or "during" recession year.

Dispersion of plant-level TFP shocks over the business cycle

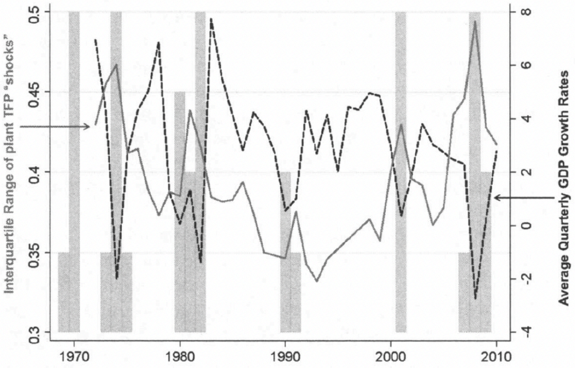

FIGURE 3.—TFP "shocks" are more dispersed in recessions. Notes: Constructed from the Census of Manufactures and the Annual Survey of Manufactures establishments, using establishments with 25+ years to address sample selection. Grey shaded columns are the share of quarters in recession within a year.

The firm's problem: primitives

- **Production**: $y_{jt} = A_t\, z_{jt}\, k_{jt}^{\alpha}\, n_{jt}^{\nu}$, $\quad \alpha + \nu < 1$ (DRS)
- **Productivity**: product of aggregate A_t and idiosyncratic z_{jt}, each AR(1):

$$\log A_t = \rho_A \log A_{t-1} + \sigma_{A,t}\, e_t$$
$$\log z_{jt} = \rho_z \log z_{j,t-1} + \sigma_{z,t}\, e_{jt}$$

- Time-varying volatility: $\sigma_{A,t} \in \{\sigma_A^L, \sigma_A^H\}$ and $\sigma_{z,t} \in \{\sigma_z^L, \sigma_z^H\}$ are governed by a common two-state uncertainty regime $S_t \in \{L, H\}$: when macro uncertainty is high, micro uncertainty is also high. Firms observe the current regime before shocks are drawn
- **Adjustment costs**:
  - → Capital: fixed disruption cost F_K (fraction of output for any $|i| > 0$) $+$ partial irreversibility (resale loss of 34%)
  - → Labor: fixed disruption cost F_L $+$ linear hiring/firing cost $\propto$ wage
- **Firm state**: (k, n_{-1}, z). **Aggregate state**: $(A, \sigma_A, \sigma_z, \mu)$

The firm's problem: Bellman equation

$$V(k, n_{-1}, z;\, A, \sigma_A, \sigma_z, \mu) = \max_{i,n}\; p \cdot \left[y - \omega\, n - i - AC^k - AC^n\right] + \beta\, \mathbb{E}\left[V(k', n, z';\, A', \sigma'_A, \sigma'_z, \mu')\right]$$

where:

$$\begin{aligned} k' &= (1-\delta_k)k + i \\ AC^k &= \mathbf{1}(|i| > 0) \cdot y \cdot F_K + \mathbf{1}(i < 0) \cdot (1-\nu_k)|i| \\ AC^n &= \mathbf{1}(|n - (1-\delta_n)n_{-1}| > 0) \cdot y \cdot F_L + |n - (1-\delta_n)n_{-1}| \cdot H \cdot \omega \end{aligned}$$

- Expectation over $(A', z', \sigma'_A, \sigma'_z)$ — four sources of uncertainty
- Nonconvex costs (F_K, F_L): create inaction regions for both k and n
- Partial irreversibility ($\nu_k < 1$): disinvestment incurs an additional linear cost — selling capital recovers only ν_k per unit
- Labor is a *state variable* (due to n_{-1} in the hiring/firing cost), unlike KT08/KT13

Connection to KT08 and Krusell-Smith

- Builds directly on the KT08 framework: heterogeneous firms, nonconvex adjustment costs, KS algorithm for the aggregate state
- **Key departure**: the *variance* of productivity shocks is itself stochastic
  - $\rightarrow$ KT08: only first-moment shocks (level of z) — nonconvexities washed out in GE
  - $\rightarrow$ Bloom et al.: second-moment shocks (variance of z and A) — nonconvexities are the *entire transmission mechanism*
- Without nonconvex adjustment costs, there are no inaction regions and firms adjust smoothly. A mean-preserving spread in shocks has only minor effects on aggregates (through Jensen's inequality / Oi-Hartman-Abel, which actually *raises* long-run output)
- With nonconvexities: higher variance $\Rightarrow$ wider (S, s) inaction bands $\Rightarrow$ more firms pause $\Rightarrow$ labor and capital depreciate without replacement $\Rightarrow$ output falls

KS implementation in Bloom et al.

- Distribution μ is over *three* firm-level states: (k, n_{-1}, z) — richer than KT08's (k, ε)
- Approximate aggregate state: (A, S, S_{-1}, K) where $S \in \{L, H\}$ is the joint uncertainty regime and S_{-1} is its lag. Conditioning on (A, S, S_{-1}) with A on a 5-point grid gives $5 \times 2 \times 2 = 20$ states

- Two log-linear forecasting rules, each conditioned on (A, S, S_{-1}):

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

- Preference assumption $\chi = 1 \Rightarrow \omega = \theta/p$, so no separate wage rule needed
- Total: $2 \times 20 = 40$ regressions. $R^2 \geq 0.94$ for all (Table B1, Online Appendix)

What needs to be forecasted: comparison across papers

	KT08	KT13	Bloom
Distribution μ over	(ε, k)	(k, b, ε)	(k, n_{-1}, z)
Conditioning states	z (11)	$s = (z, \zeta)$ (6)	(A, S, S_{-1}) (20)
Forecast $\log m'$ / $\log K'$	✓	✓	✓
Forecast $\log p$	✓	✓	✓
Forecast ω	× (ϕ/p)	× (D_2U/p)	× (θ/p)
Regressor	$\log m$	$\log m, c_1, c_2$	$\log K$
# regressions	$2 \times 11 = 22$	$2 \times 6 = 12$	$2 \times 20 = 40$
R^2 (min)	> 0.999	> 0.999	≥ 0.94

- All three papers exploit preference assumptions to avoid forecasting ω
- As complexity grows (more states, more conditioning), the number of regressions proliferates — from 22 to 40 — and the choice of what to include remains ad hoc

The mechanism: "wait and see"

1. Uncertainty rises ($\sigma_{z,t}$, $\sigma_{A,t}$ switch to high regime)
2. The option value of waiting increases: with nonconvex costs, it pays to delay investment/hiring until the shock is realized
3. Firms pause: inaction bands for both k and n widen $\Rightarrow$ hiring and investment freeze
4. **Misallocation rises**: productive firms stop expanding, unproductive firms stop contracting $\Rightarrow$ reallocation freezes $\Rightarrow$ measured TFP falls
5. When uncertainty subsides: pent-up demand $\Rightarrow$ overshoot in investment and hiring, but recovery is slowed by persistent misallocation and a depleted capital stock

The mechanism: “wait and see”

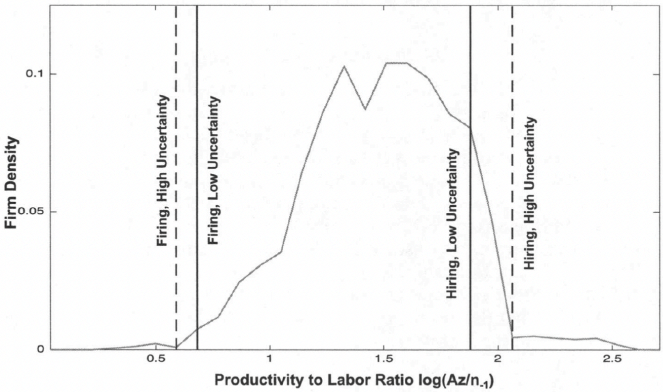

FIGURE 5.—The impact of an increase in uncertainty on the hiring and firing thresholds. Notes: The figure plots the simulated cross-sectional marginal distribution of micro-level labor inputs after productivity shock realizations and before labor adjustment. The distribution plots a representative period with average aggregate productivity and low uncertainty levels. The vertical hiring and firing thresholds are computed based on firm policy functions with average micro-level productivity realizations, taking as given the aggregate state of the economy with low uncertainty (solid lines) and a high uncertainty counterfactual (dotted lines).

Impulse response to an uncertainty shock

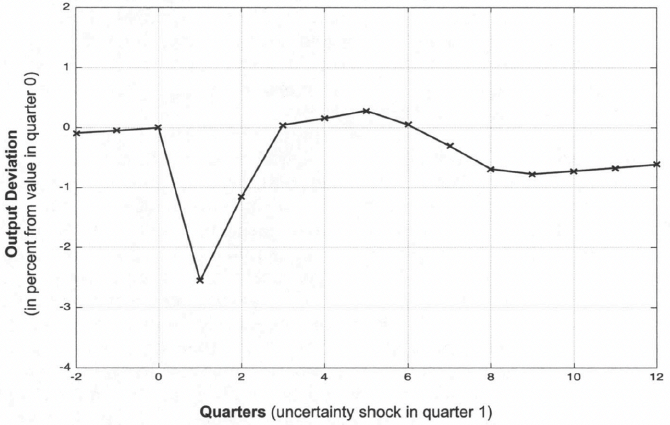

FIGURE 6.—The effects of an uncertainty shock. Notes: Based on independent simulations of 2500 economies of 100-quarter length. We impose an uncertainty shock in the quarter labelled 1, allowing normal evolution of the economy afterwards. We plot the percent deviation of cross-economy average output from its value in quarter 0.

Impulse response to an uncertainty shock

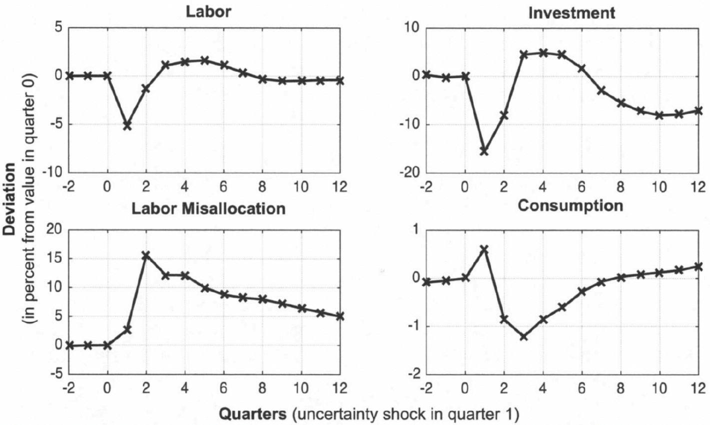

FIGURE 7.—Labor and investment drop and rebound, misallocation rises, and consumption overshoots then falls. Notes: Based on independent simulations of 2500 economies of 100-quarter length. We impose an uncertainty shock in the quarter labelled 1, allowing normal evolution of the economy afterwards. Clockwise from the top left, we plot the percent deviations of cross-economy average labor, investment, consumption, and the dispersion of the marginal product of labor from their values in quarter 0.

Main results

1. Uncertainty shocks generate GDP drops of ≈2.5% through the wait-and-see channel
2. Misallocation (dispersion of marginal products) worsens endogenously — echoes the KT13 TFP mechanism but through a different channel
3. Uncertainty shocks alone predict a *rise* in consumption on impact (investment freeze frees up resources). Need first-moment shocks to fit consumption
4. Recessions are best modeled as *joint* negative first-moment + positive second-moment shocks
5. Uncertainty dampens the effectiveness of first-moment policies (e.g., wage subsidies) — firms are more cautious in responding to price changes

Adding first-moment shocks

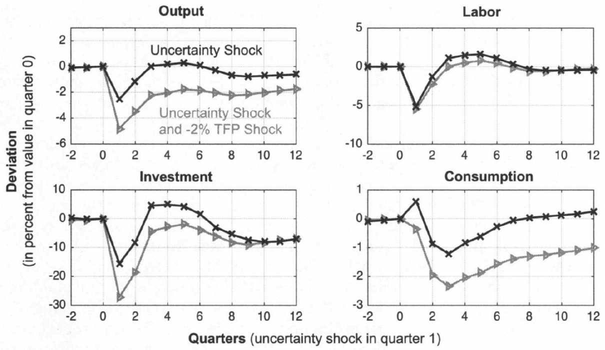

FIGURE 8.—Adding a −2% first-moment shock increases the output fall and eliminates a consumption overshoot. Notes: Based on independent simulations of 2500 economies of 100-quarter length. For the baseline (× symbols) we impose an uncertainty shock in the quarter labelled 1. For the uncertainty and TFP shock (► symbols), we also impose an aggregate productivity shock with average equal to −2%, allowing normal evolution of the economy afterwards. 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.

Adding a -2% TFP shock eliminates the consumption overshoot and generates a more realistic recession profile.

Role of GE

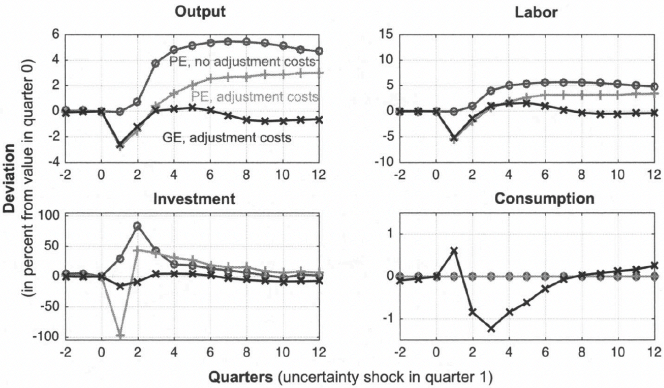

FIGURE 10.—The impact of an uncertainty shock combines Oi–Hartman–Abel, real options and consumption smoothing effects. Notes: Based on independent simulations of 2500 economies of 100-quarter length. For all simulations we impose an uncertainty shock in the quarter labelled 1, allowing normal evolution of the economy afterwards. GE, adjustment costs (× symbols) is the baseline, and the PE responses are partial equilibrium paths with adjustment costs (+ symbols) and without adjustment costs (∘ symbols). 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. Note that PE economies have no consumption concept, with deviations therefore set to 0.

Lessons from Bloom et al. (2018)

Lesson 1: Nonconvexities do matter for aggregates when the shock is to the *second moment*. This completes the triptych:

- KT08: first-moment TFP shock $+$ nonconvex costs $\Rightarrow$ GE washes out
- KT13: credit shock $+$ collateral constraint $\Rightarrow$ distribution matters via misallocation
- Bloom: second-moment shock $+$ nonconvex costs $\Rightarrow$ nonconvexities are the mechanism

Lesson 2: Recessions involve *both* first- and second-moment shocks. Neither alone matches the data.

Lesson 3: Uncertainty shocks reduce policy effectiveness — a cautionary note for stabilization policy design.

Taking Stock

Taking stock: what we have learned

1. **Krusell & Smith**: when agents face aggregate shocks, the distribution becomes a state variable. The KS algorithm approximates it with a few moments — and often this works ("approximate aggregation")
2. **Khan & Thomas (2008)**: micro nonconvexities in investment vanish at the aggregate level once GE price adjustment is accounted for. First-moment TFP shocks $+$ real frictions $\Rightarrow$ distribution doesn't matter
3. **Khan & Thomas (2013)**: financial frictions change the picture. Credit shocks cause misallocation and endogenous TFP declines. The distribution is economically central, and the KS algorithm requires more (ad hoc) moments
4. **Bloom et al. (2018)**: second-moment shocks interact with nonconvexities to cause recessions through the wait-and-see channel. The type of shock, not just the type of friction, determines whether heterogeneity matters

Where do we go from here?

The Krusell-Smith algorithm has served us well, but it has fundamental limitations:

- **Moment selection**: which moments of μ to track? Ad hoc, model-specific, doesn't scale
- **Simulation-based**: requires simulating large panels of agents — slow, noisy
- **No clean characterization**: doesn't tell us *how* the distribution matters for GE, only *whether* approximate aggregation holds after the fact

Next lectures: a different approach that resolves all three issues:

- MIT shocks as numerical derivatives: the impulse response to a small one-time shock *is* the Jacobian of the equilibrium mapping — the first paper to make this conceptual connection (Boppart, Krusell & Mitman, 2018).
- Sequence-space Jacobians: compute these Jacobians *analytically* via the fake news algorithm, then combine them across blocks using the chain rule. Once computed, impulse responses to *any* shock are immediate.(Auclert, Bardóczy, Rognlie & Straub, 2021)

Readings

This lecture

1. Krusell, P. and A. Smith (1998). "Income and Wealth Heterogeneity in the Macroeconomy," *JPE*, 106(5), 867–896.

   `doi:10.1086/250034`

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

   `doi:10.1111/j.1468-0262.2008.00837.x`

3. Khan, A. and J. Thomas (2013). "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity," *JPE*, 121(6), 1055–1107.

   `doi:10.1086/674142`

4. Bloom, N. et al. (2018). "Really Uncertain Business Cycles," *Econometrica*, 86(3), 1031–1065.

   `doi:10.3982/ECTA10927`

Readings (continued)

Next lectures — MIT shocks and sequence-space Jacobians

5. Boppart, T., P. Krusell, and K. Mitman (2018). "Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative," *Journal of Economic Dynamics and Control*, 89, 68–92.

   `doi:10.1016/j.jedc.2018.01.002`

6. Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89(5), 2375–2408.

   `doi:10.3982/ECTA17434`

Supplementary

7. Algan, Y., O. Allais, W. Den Haan, and P. Rendahl (2014). "Solving and Simulating Models with Heterogeneous Agents and Aggregate Uncertainty," *Handbook of Computational Economics*, Vol. 3, Ch. 6.

   `doi:10.1016/B978-0-444-52980-0.00006-2`