Income Dispersion and Counter-Cyclical Markups: Supplementary Material and Technical Appendix

Chris Edmond[*] Laura Veldkamp[†]

This draft: June 2009

This appendix provides the formal proof of Proposition 1, an analytic calculation showing that indivisible x-goods are not essential for our main results, and a numerical example showing that our results can also be obtained in a random coefficients discrete choice demand system like that of Berry, Levinsohn and Pakes (1995), commonly used in the industrial organization literature. The appendix also outlines how the model with GHH preferences is solved numerically, does sensitivity analysis, and describes the data and calibration procedures in more detail. Finally, the appendix sketches a version of the model with endogenous risk sharing.

A Proof of Proposition 1

Since marginal cost is constant and normalized to 1, the optimal markup $m(z,\sigma)$ is equal to the optimal price and so satisfies the first-order condition:

$$m(z,\sigma)-1=\frac{1}{h\left(\frac{\theta m(z,\sigma)/\nu-z}{\sigma}\right)}\frac{\nu\sigma}{\theta}, \tag{1}$$

where $h(\varepsilon)=f(\varepsilon)/(1-F(\varepsilon))>0$ is the hazard rate of the distribution of idiosyncratic effective labor productivity. By Assumption 1 the hazard is increasing $h'(\varepsilon)\geq 0$. Application of the implicit function theorem then gives:

$$\frac{\partial m(z,\sigma)}{\partial z}=\frac{h'\left(\frac{\theta m(z,\sigma)/\nu-z}{\sigma}\right)}{h\left(\frac{\theta m(z,\sigma)/\nu-z}{\sigma}\right)^2+h'\left(\frac{\theta m(z,\sigma)/\nu-z}{\sigma}\right)}\frac{\nu}{\theta}\geq 0, \tag{2}$$

[*]New York University, Stern School of Business and University of Melbourne.
[†]New York University, Stern School of Business, CEPR and NBER.

and similarly:

$$\frac{\partial m(z,\sigma)}{\partial \sigma} = \frac{h'\left(\frac{\theta m(z,\sigma)/\nu - z}{\sigma}\right)\left(\frac{\theta m(z,\sigma)/\nu - z}{\sigma}\right) + h\left(\frac{\theta m(z,\sigma)/\nu - z}{\sigma}\right)}{h\left(\frac{\theta m(z,\sigma)/\nu - z}{\sigma}\right)^2 + h'\left(\frac{\theta m(z,\sigma)/\nu - z}{\sigma}\right)} \frac{\nu}{\theta} > 0. \tag{3}$$

Notice that since $w_i > 0$ all i, the optimal markup is at least $m(z,\sigma) > \nu z/\theta$ so that both the numerator and denominator in this last expression are strictly positive. □

B Indivisibility of x goods not essential

Markups can be counter-cyclical because an increase in the dispersion of income will endogenously decrease the elasticity of demand, even without indivisible goods. To make this point, consider a variant of the illustrative model in Section 1. Households $i \in [0,1]$ have identical quasi-linear preferences $U_i = \log(c_i) + x_i$ over a competitive c-good and a monopolistically supplied x-good. The x-good is perfectly divisible in that a household can choose any $x_i \geq 0$. Households choose c_i and x_i to maximize utility subject to the budget constraint $c_i + px_i \leq y_i$ where p is the relative price of the x-good in units of the c-good and $y_i > 0$ is the household's total income. Total income $y_i > 0$ is IID uniform in the population on $[z-\sigma, z+\sigma]$ where $z > \sigma > 0$.

A household's optimal demand for the x-good is given by:

$$x_i = \max\left(0, \frac{y_i}{p} - 1\right). \tag{4}$$

As in the illustrative model, only households that have a high income relative to the price of the x-good will buy x-goods. The aggregate demand curve for the monopolist x-good producer is then:

$$x(p) := \int_0^1 x_i di = \int_{z-\sigma}^{z+\sigma} \max\left(0, \frac{y_i}{p} - 1\right) \frac{1}{2\sigma} dy_i. \tag{5}$$

Calculating the integral and simplifying gives:

$$x(p) = \int_p^{z+\sigma} \left(\frac{y_i}{p} - 1\right) \frac{1}{2\sigma} dy_i = \left[\frac{(z+\sigma)^2}{2p} + \frac{1}{2}p - (z+\sigma)\right] \frac{1}{2\sigma}. \tag{6}$$

As in the benchmark model, the monopolist x-good producer has constant marginal costs normalized to 1 and chooses price p to maximize profits,

$$\pi(p) := (p-1)x(p), \tag{7}$$

taking as given the demand curve (6).

To ensure that the x-good producer will operate:

Assumption 1. *Marginal costs are sufficiently low,* $1 \leq z + \sigma$.

Let $m(z, \sigma)$ denote the optimal markup chosen by the x-good producer. Since marginal costs are constant and normalized to one, this is the same as the optimal price. Qualitatively, the comparative static results with respect to z and σ are the same as in the main text:

Proposition 1. *In the model with divisible x-good, the optimal markup $m(z, \sigma)$ is increasing in mean income z and increasing in income dispersion σ.*

Hence the indivisibility of the x-good as used in the main text is not essential for the markup to be increasing in dispersion.

Proof. First observe that, up to a positive scalar that is irrelevant for the optimal pricing decision, the x-good producer's profit function depends only on the sum $z + \sigma$ and not on z or σ separately. Given this, let $k := z + \sigma$. Now note the boundaries of the producer's problem: since profits are $\pi(p) = (p-1)x(p)$, setting $p = 1$ ensures zero profit as does setting the higher price $p = k$ (since $x(k) = 0$ from equation 6). Now let's turn to interior solutions.

The first order necessary condition characterizing the optimal price can be written:

$$\pi'(p) = (p-1)x'(p) + x(p) = 0. \tag{8}$$

Using the formula (6) for the demand curve gives:

$$\pi'(p) = p - k + \left[(k/p)^2 - 1\right]/2 = 0. \tag{9}$$

Similarly, the second order condition for a maximum is:

$$\pi''(p) = 1 - k^2/p^3 < 0. \tag{10}$$

There are two solutions in the interval $[1, k]$ to the first order condition, one interior and one on a boundary. From the second order condition, marginal profit $\pi'(p)$ is continuous and strictly decreasing in p for all $p \in (1, k^{2/3})$ and is then strictly increasing for all $p \in (k^{2/3}, k)$ and asymptotes to $\pi'(p) = 0$ as $p \to k$ from below. Moreover, $\pi'(1) > 0 > \pi'(k^{2/3})$. So from the intermediate value theorem the first order condition has a unique interior solution $p \in (1, k^{2/3})$ and has a second, higher, solution at the boundary $p = k$. This higher solution violates the second order condition and indeed, as noted above both this boundary solution $p = k$ and the other boundary solution $p = 1$ lead to zero profits. By contrast the interior solution, call it $p(k)$, leads to strictly positive profits — this is because $\pi(1) = 0$ and $\pi'(p) > 0$ for all $p < p(k)$. Therefore the unique interior local maximum at $p(k)$ is also the unique global maximum.

Applying the implicit function theorem shows that this optimal price $p(k)$ is increasing in k if and only if the marginal profit function is increasing in k, that is:

$$p'(k) \geq 0 \Leftrightarrow \frac{\partial}{\partial k}\{p - k + \left[(k/p)^2 - 1\right]/2\} \geq 0, \tag{11}$$

when the partial derivative is evaluated at the optimum, $p(k)$. Calculating the derivative and simplifying:

$$p'(k) \geq 0 \Leftrightarrow k^{1/2} \geq p(k). \tag{12}$$

Now let's check if this condition holds. Evaluating the marginal profit function $\pi'(p)$ at $p = k^{1/2}$ gives:

$$\pi'(k^{1/2}) = k^{1/2} - k/2 - 1/2 =: A(k). \tag{13}$$

Now since $A(1) = 0$ and $A'(k) < 0$ for all $k > 1$ it follows that $\pi'(k^{1/2}) \leq 0$ for all $k \geq 1$ with equality if and only if $k = 1$. But this implies the optimal price is indeed $p(k) \leq k^{1/2}$ (since $p(k)$ satisfies $\pi'(p(k)) = 0$ at a point where the second order condition holds). And so (12) implies $p'(k) \geq 0$ with equality if and only if $k = 1$. Since $k := z + \sigma$ and the optimal markup is equal to the optimal price, $m(z, \sigma) = p(z + \sigma)$. Therefore the optimal markup is increasing in mean income z and increasing in income dispersion σ. □

The intuition for this result is the same as in the benchmark model. An increase in the mean z shifts out and steepens the firm's marginal revenue curve, leading to higher sales of x and a higher markup. An increase in dispersion σ shifts in but also steepens the firm's marginal revenue curve so sales of x fall but the markup rises.

C Random coefficients discrete choice demand

This section outlines a version of our model with a discrete-choice random-coefficients demand system like that of Berry et al. (1995), commonly used in the industrial organization literature. This exercise also clarifies what model features are essential for more dispersion to raise markups.

In our benchmark model, the island structure prevents competition between producers. A consumer is randomly allocated to an island j and chooses to buy the good at the posted price or not. Given the indirect utility function, a consumer then chooses $x_{ij} = 1$ for all j such that $\theta p_j/\nu \leq w_i$. But in a BLP-type model, the choice among j differentiated goods is *mutually exclusive*. Each consumer faces a menu of differentiated goods with prices p_j and can choose one and only one of the j goods to consume.

Using the notation of our benchmark static model, the Lagrangian for consumer i is:

$$L = \log c_i - \theta n_i + \nu \int_0^1 x_{ij} dj + \lambda_i \left(w_i n_i + \pi - c_i - \int_0^1 p_j x_{ij} dj \right). \tag{14}$$

Collecting terms and noting that the first order conditions imply $\theta = \lambda_i w_i$ the indirect utility from consuming $x_{ij} = 1$ can be written as:

$$u_{ij} = \nu - \lambda_i p_j = \nu - \frac{\theta}{w_i} p_j. \tag{15}$$

This now takes the form of a random coefficients, discrete choice, demand system.

To smooth out the problem facing firms, it's convenient to introduce an additional source of dispersion in valuations. In particular, suppose instead of a constant marginal utility ν for all i and j there are random ν_{ij} (i.e., taste shocks). Moreover suppose that ν_{ij} can be decomposed into a product-specific component $\bar{\nu}_j$ common across all consumers and an idiosyncratic component ε_{ij}, that is:

$$\nu_{ij} = \bar{\nu}_j + \varepsilon_{ij}, \tag{16}$$

so that indirect utility from consuming $x_{ij} = 1$ is:

$$u_{ij} = \nu_{ij} - \frac{\theta}{w_i} p_j = \bar{\nu}_j - \frac{\theta}{w_i} p_j + \varepsilon_{ij}. \tag{17}$$

To obtain the familiar logit formulas for demand, assume that ε_{ij} is IID extreme value (McFadden 1973). Then demand for a producer is given by the mixed logit formula:

$$x_j(p, \mathbf{p}) = \int_0^\infty \left(\frac{\exp(\bar{\nu}_j - \theta p / w_i)}{1 + \int_0^1 \exp(\bar{\nu}_k - \theta p_k / w_i) dk} \right) f(w_i | z, \sigma)\, dw_i, \tag{18}$$

See also Nevo (2000) or Train (2003) for further discussion and detailed derivations. Here $\mathbf{p} := (p_k)$ denotes the collection of prices charged by all firms and the demand curve is indexed by j because of the product-specific $\bar{\nu}_j$ terms. For each firm, the profit maximization problem can be written:

$$p_j(\mathbf{p}) \in \arg\max_p \left[(p-1) x_j(p, \mathbf{p}) \right]. \tag{19}$$

and we then need to solve for a fixed-point in prices. To keep the number of free parameters from proliferating, let $\bar{\nu}_j = \nu$ for all j. With this assumption the j index can be dropped from the demand curve, and, since firms are now identical, it's natural to look for a symmetric equilibrium where each firm charges the same price, $p_j = p$ for all j. This equilibrium price is a function of the parameters z, σ of the distribution of idiosyncratic productivities $f(w_i|z, \sigma)$.

The model is solved under two scenarios. In the first, as in our benchmark model, consumers can choose not to consume any of the x goods and instead simply consume more of the numeraire c good. In this *variable x-good demand* scenario, the total demand for x goods is not fixed. In our second scenario, there is *fixed x-good demand*, consumers necessarily choose one of the x-goods and it is only a question of which one (recall that the choice over x-goods is mutually exclusive).

Variable x-good demand. If, as in our benchmark model, consumers can take the 'outside option' of choosing no x good and instead simply consume more c good, the model cannot be solved analytically. We solved this version of the model numerically with the average utility weight on x-goods ν normalized to 1, the utility weight on leisure θ set to 3 so that labor supply is roughly 1/3rd, a lognormal productivity distribution with aggregate productivity $z = 0.7$ and we varied the degree of productivity dispersion. These parameters lead to markups for the x goods in the 20% – 40% range, as in Berry et al. (1995). As shown in figure C.1, with these parameter values the markup is increasing in earnings dispersion.

Fixed x-good demand. If, however, consumers do not have the outside option of simply consuming more c good, then this result is overturned. In this case, the demand curve facing a firm becomes:

$$x_j(p, \mathbf{p}) = \int_0^{\infty} \left(\frac{\exp(\bar{\nu}_j - \theta p/w_i)}{\int_0^1 \exp(\bar{\nu}_k - \theta p_k/w_i) dk} \right) f(w_i|z, \sigma)\, dw_i. \tag{20}$$

Again consider the symmetric case where $\bar{\nu}_j = \nu$ all j and compute equilibrium prices. It is straightforward to show analytically that in this case any mean-preserving increase in dispersion σ *reduces* the equilibrium price and markup (details available on request).

What this teaches us is that the cyclical variation in the amount of x-goods purchased is an essential component of our results. When, as in this second scenario, all consumers are required to buy one unit of x-goods, then an x-good firm that raises its price potentially loses customers to another x-good firm. But when — as in the first scenario and our benchmark model — consumers have the outside option, then a firm that raises its price loses customers who choose to buy the c good instead. The elasticity of demand in these two situations varies in opposite ways over the business cycle.

D Solving the model with GHH preferences

Households buy good x_j if and only if $\nu \geq \lambda_i p_j$, where λ_i is the Lagrange multiplier on the budget constraint and satisfies:

$$\lambda_i = \left(c_i - \theta \frac{n_i^{1+\gamma}}{1+\gamma} \right)^{-1}. \tag{21}$$

In equilibrium each seller sets the same price $p = p_j$ for all j and each household buys or not. Since x-goods are sold in discrete $\{0, 1\}$ amounts, each household's expenditure on x-goods is either p or 0.

Now let $\hat{p}(w_i)$ denote the highest price that a household with idiosyncratic productivity w_i will pay for the x-good. Using the cutoff rule for x-good purchases, this price $\hat{p}(w_i)$ satisfies $\nu = \hat{p}(w_i)\lambda_i$. Substituting for λ_i from (21) gives:

$$\begin{aligned}
\nu &= \hat{p}(w_i) \left(c_i - \theta \frac{n_i^{1+\gamma}}{1+\gamma} \right)^{-1} \\
&= \hat{p}(w_i) \left(w_i \hat{n}(w_i) + \pi - \hat{p}(w_i) - \theta \frac{\hat{n}(w_i)^{1+\gamma}}{1+\gamma} \right)^{-1},
\end{aligned}$$

where the second line uses the budget constraint (equation 2 from the main text) to eliminate c_i and where $\hat{n}(w_i)$ is given by the labor supply curve (equation 13 from the main text). Solving for $\hat{p}(w_i)$ yields:

$$\hat{p}(w_i) = \frac{\nu}{1+\nu} \left(w_i \hat{n}(w_i) + \pi - \theta \frac{\hat{n}(w_i)^{1+\gamma}}{1+\gamma} \right). \tag{22}$$

This is a continuous, strictly increasing and strictly convex function of w_i.

Now let $\hat{w}(p)$ denote the inverse of the reservation price function $\hat{p}(w)$, i.e., $\hat{w}(\hat{p}) = 1$, so that $\hat{w}(p)$ represents the lowest idiosyncratic productivity draw that will lead a household to purchase at price p (so $\hat{w}(p)$ is strictly increasing and strictly concave). Then the total demand facing the x-firm on island j at price is:

$$x(p_j) = \Pr[w \geq \hat{w}(p_j)] = 1 - \Phi\left(\frac{1}{\sigma}\log\left(\frac{\hat{w}(p_j)}{z}\right)\right), \tag{23}$$

where Φ denotes the standard normal cumulative distribution function.

The resulting profit of firm j is revenue minus costs $\pi_j = p_j x(p_j) - x(p_j)^{1/\alpha}$. Aggregate profits are $\pi = \int_0^1 \pi_j dj$. Since in equilibrium $p_j = p$ for all j, $\pi_j = \pi$ for all j too. In equation (22), aggregate profits π enter the reservation price function $\hat{p}(w)$ and therefore enter the inverse $\hat{w}(p)$ too. To acknowledge this dependence, write $\hat{w}(p, \pi)$. To compute an equilibrium numerically, then, we have to solve a fixed point problem of the form $\pi = G(\pi)$ where:

$$G(\pi) := \max_{p\geq 0}\left\{p\left[1 - \Phi\left(\frac{1}{\sigma}\log\left(\frac{\hat{w}(p,\pi)}{z}\right)\right)\right] - \left[1 - \Phi\left(\frac{1}{\sigma}\log\left(\frac{\hat{w}(p,\pi)}{z}\right)\right)\right]^{1/\alpha}\right\}. \tag{24}$$

We solve this fixed point problem numerically, guessing an initial π_0, iterating on $\pi_{k+1} = G(\pi_k)$ for $k \geq 0$ and then iterating until $|\pi_{k+1} - \pi_k| < 10^{-6}$.

Steady state calibration targets. The six parameters $(\alpha, \gamma, \nu, \theta, \sigma, \overline{z})$ are chosen such that the steady state of our model delivers the following six properties:

$$\begin{aligned}
\text{elasticity of labor supply} &= \mathbb{E}[d\log(n_i)/d\log(w_i)] &= 1.67\\
\text{earnings dispersion} &= \overline{\sigma}_{\text{STY}} &= 0.29\\
\text{hours worked} &= \mathbb{E}[n_i] &= 0.33\\
\text{labor share} &= \mathbb{E}[w_i n_i]/\overline{y} &= 0.67\\
\text{aggregate markup} &= \overline{m} &= 1.10\\
x\text{-sector markup} &= \overline{m}_x &= 1.23
\end{aligned}$$

The following properties of the model are used repeatedly: individual labor supply is $n_i = (w_i/\theta)^{1/\gamma}$ (from equation 13 in the main text). Since log idiosyncratic productivity $\log(w_i)$ is normal with mean $\log(\overline{z})$ and standard deviation σ, log labor supply is:

$$\log(n_i) \sim \mathcal{N}[(1/\gamma)\log(\overline{z}/\theta), (\sigma/\gamma)^2],$$

and so:

$$\mathbb{E}[n_i] = (\overline{z}/\theta)^{1/\gamma}\exp[0.5(\sigma/\gamma)^2].$$

Similarly, log earnings is:

$$\log(w_i n_i) \sim \mathcal{N}[((1+\gamma)/\gamma)\log(\overline{z}) - (1/\gamma)\log(\theta), ((1+\gamma)\sigma/\gamma)^2],$$

and so:

$$\mathbb{E}[w_i n_i] = \overline{z}^{(1+\gamma)/\gamma}\theta^{-1/\gamma}\exp[0.5((1+\gamma)\sigma/\gamma)^2].$$

Given this, the average elasticity of labor supply is $\mathbb{E}[d\log(n_i)/d\log(w_i)] = 1/\gamma$ which equals 1.67 when $\gamma = 0.60$. The standard deviation of log earnings is $[(1+\gamma)/\gamma]\sigma$ which equals the Storesletten, Telmer and Yaron (2004) estimate of $\overline{\sigma}_{\text{STY}} = 0.29$ when $\sigma = (0.60/1.60)(0.29) = 0.11$.

The remaining four parameters $(\alpha, \nu, \theta, \overline{z})$ have to be solved for simultaneously. From our previous calculations, one condition is immediate:

$$\mathbb{E}[n_i] = (\overline{z}/\theta)^{1/0.60}\exp[0.5(0.11/0.60)^2] = 0.33. \tag{25}$$

Labor's share is $\mathbb{E}[w_i n_i]/\overline{y} = 0.67$ and since $\overline{y} = \mathbb{E}[w_i n_i] + \overline{\pi}$, average earnings $\mathbb{E}[w_i n_i] = (0.67/0.33)\overline{\pi}$. To calculate $\overline{\pi}$ the fixed point problem $\overline{\pi} = G(\overline{\pi})$ outlined above has to be solved. The solution of the fixed point problem depends on all four parameters and to acknowledge this write $\overline{\pi}(\alpha, \nu, \theta, \overline{z})$. Using the formula for average earnings derived above gives a second equation in the four unknowns:

$$\mathbb{E}[w_i n_i] = \overline{z}^{(1.60)/0.60}\theta^{-1/0.60}\exp[0.5((1.60)(0.11)/0.60)^2] = \frac{0.67}{0.33}\overline{\pi}(\alpha, \nu, \theta, \overline{z}). \tag{26}$$

The solution to the fixed point problem also gives us an optimal price $\overline{p}(\alpha, \nu, \theta, \overline{z})$ and associated percentage markup $\overline{m}_x(\alpha, \nu, \theta, \overline{z})$ for the x-sector. Our third equation in the four unknowns is therefore:

$$\overline{m}_x(\alpha, \nu, \theta, \overline{z}) = 1.23. \tag{27}$$

Let $\overline{x}$ denote the amount of the x-good sold by the firm at the price $\overline{p}$ and let $\overline{e}_x := (\overline{px})/\overline{y}$ denote the expenditure share on the x-sector. Define the aggregate markup as the expenditure share weighted average of the x-sector and c-sector markups, $\overline{m} := \overline{e}_x\overline{m}_x + (1 - \overline{e}_x)$, since the c-sector percentage markup is zero by definition. Rearranging gives our fourth equation:

$$\overline{e}_x(\alpha, \nu, \theta, \overline{z}) = \frac{\overline{m} - 1}{\overline{m}_x - 1} = \frac{0.10}{0.23}. \tag{28}$$

In short, the four parameters $(\alpha, \nu, \theta, \overline{z})$ are found by solving the four equations (25)-(28) simultaneously. This gives us $\alpha = 0.24, \nu = 100, \theta = 15$ and $\overline{z} = 7.68$.

E Sensitivity analysis

Diminishing returns parameter α. The top panel of table E.1 shows that halving or doubling α leaves most of our main results intact. The most notable exception is that when α is very low, real wages become counter-cyclical. This happens because firms' marginal costs are very volatile. Since those costs are pro-cyclical, it makes prices strongly pro-cyclical and real wages counter-cyclical.

Utility weight on x-good. The middle panel of table E.1 shows that halving or doubling ν has almost no effect on either our model calibration targets or the model's implications for other macro aggregates. Our calibration procedure does not precisely pin down a value for ν, but the model's ability to reproduce business cycle facts does not depend crucially on our benchmark ν value.

Level of aggregate productivity $\bar{z}$. The bottom panel of table E.1 shows that our model is much more sensitive to the level of the aggregate productivity. This is because of our non-homothetic preferences. Calibration targets, such as hours and profit shares, change rapidly as $\bar{z}$ increases or decreases. The advantage of this sensitivity is that it means that the data provide very precise information about what the level of $\bar{z}$ should be. The disadvantage, of course, it that the model's ability to reproduce business cycle facts deteriorates when $\bar{z}$ is changed. A 10% rise in the level of aggregate productivity (from the benchmark $\bar{z} = 7.7$ to 8.5) leads the model to predict counter-cyclical real wages. But if the model is re-calibrated by making offsetting changes in other parameters — e.g., if the weight θ on leisure is increased so as to pull labor supply back down — the model regains its ability to explain business cycle facts.

Results broken down by sector. While the paper reports aggregate statistics for the economy as a whole, it is also useful to see what happens within each sector. Table E.2 reports markups, profit shares and expenditure shares for the x and c sectors separately for our benchmark parameters.

F Decline in business cycle volatility.

Higher dispersion can lower business cycle volatility because high earnings dispersion reduces aggregate demand elasticity. Therefore, shocks to labor productivity have less effect on who buys what products. Since producers are producing in anticipation of changes in aggregate demand, when aggregate demand becomes less volatile, GDP volatility falls as well. While our model does not explain the bulk of the fall in business cycle volatility, as shown in table F.1, it can — with the balanced growth correction — generate a modest decline.

In each decade, the model's earnings dispersion by decade is chosen so that its log change from the previous decade matches Heathcote, Storesletten and Violante (2006). When only that change is made to the model, business cycle volatility increases because of an unintended side-effect: When dispersion increases, the average productivity level does as well because idiosyncratic productivity is lognormal. Higher productivity raises hours worked and shifts the expenditure from c-good to consumption to x-good consumption. x-good consumption is more volatile because of its linear utility.

With the balanced growth correction, results improve. Our model predicts essentially flat business cycle volatility. To achieve a decrease in business cycle volatility, the dispersion process needs to be smaller and less rapidly growing, like that for consumption dispersion. Storesletten et al. (2004) report that food consumption has about 2/3rds the dispersion of earnings. We match the level of dispersion in the 70's and its cyclical properties to their estimate. For the long-run increase in consumption dispersion, we use the 5% per-decade increase in non-durable consumption dispersion reported by Krueger and Perri (2005) for 1970-2000. The 10% rise in consumption dispersion from the 1970's-90's results in a 24% drop in the standard deviation of log real GDP, but it does not reproduce the halving of business cycle volatility in the data.

G Data and simulation details

Making annual dispersion quarterly. The quarterly persistence and standard deviation of income are derived from the annual estimates of Storesletten et al. (2004) as follows: $\rho_\xi = 0.952^{1/4}$, the standard deviation to the persistent component is $0.125Q$ when productivity is above average and is $0.211Q$ when productivity is below average while the standard deviation of the transitory component is $0.255Q$ where the adjustment factor is $Q := 1/(1 + \rho_\xi + \rho_\xi^2 + \rho_\xi^3) = 0.2546$.

Storesletten et al. (2004) also report consumption dispersion estimates, using food consumption data from the PSID. The same procedure as for earnings was used to transform annual to quarterly estimates. Since the persistence of consumption is slightly lower than earnings, the annual to quarterly conversion factor is different. The quarterly AR1 coefficient in persistent piece of individual earnings is $\rho_\xi = 0.862^{1/4} = 0.964$. This delivers a factor for converting annual to quarterly standard deviations $Q = (1 + \rho_\xi + \rho_\xi^2 + \rho_\xi^3)^{-1} = 0.264$. This conversion yields the parameters of the idiosyncratic earnings process. To convert these to parameters of the idiosyncratic productivity process as fed into the model, multiply each by $\gamma/(1 + \gamma)$. Thus, $\sigma_B = 0.172Q\gamma/(1 + \gamma) = 0.017$, $\sigma_R = 0.222Q\gamma/(1 + \gamma) = 0.02$, and $\sigma_u Q\gamma/(1 + \gamma) = 0.283Q = 0.028$. The resulting steady state dispersion of consumption is 0.21, about 2/3rds of the steady state earnings dispersion (0.29) from the benchmark model. None of the other parameters are changed.

Simulations. All simulations in this paper begin by sampling the exogenous state variables for a 'burn-in' of 1000 quarters. This eliminates any dependence on arbitrary initial conditions. A cross-section of 2500 individuals is tracked for 200 quarters, corresponding to the dimensions in Storesletten et al. (2004). Realizations of endogenous variables are then computed. The moments discussed in the text are averages over the results from 100 runs of the simulation (that is, averages over $200 \times 100 = 20000$ observations).

Aggregate data. All data is quarterly 1947:1-2006:4 and seasonally adjusted. Real GDP is from the Bureau of Economic Analysis (BEA). This is nominal GDP deflated deflated by the BEA's chain-type price index with a base year of 2000. Aggregate labor productivity and real wages are measured as real output per hour and real compensation per hour in the non-farm business sector, both from the Bureau of Labor Statistics (BLS) Current Employment Survey (CES) program. Nominal output and compensation are deflated by the BLS's business sector implicit price deflator with a base year of 1992.

H A model with endogenous risk-sharing

Because risk-sharing determines how much heterogeneity in demand for goods arises from heterogeneity in earnings, an important extension of this model would be to allow households to share risk. Perfect risk-sharing is both unrealistic and problematic: By ensuring that all households have the same consumption, it would collapse heterogeneity entirely, rendering our mechanism irrelevant. Rather, one would want to incorporate some limited risk-sharing. A common approach is to allow households to trade non-state-contingent bonds. They can then ensure their consumption stream by borrowing and lending. We first sketch a setup of our model with borrowing and lending and then discuss some of the technical challenges that make solving such a model another research project, in itself.

Model setup. Suppose that consumers can trade non-state-contingent bonds that are in zero net supply. The individual state for a consumer is their current idiosyncratic productivity w and asset level a, call this $s = (a, w)$. The aggregate state of the economy is the joint distribution $\mu(s)$ over individual states (knowing μ implies a current aggregate productivity level z and dispersion σ).

Consumer's problem: Consumers take as given indivisible goods prices p, interest rate r, and lump-sum profits π and chooses consumption c, labor supply n, indivisible goods x_j for $j \in [0, 1]$ and next period's asset position a' to solve the Bellman equation:

$$v(s, \mu, p, r, \pi) = \max_{c,n,x} \{U(c, n, x) + \beta\mathbb{E}[v(s', \mu', p', r', \pi')|s, \mu]\} \tag{29}$$

subject to the budget constraint:

$$c + \int_0^1 p_j x_j dj + a' \leq w(s, \mu)n + (1 + r)a + \pi$$

with $c, a' \geq 0$, $n \in [0, 1]$, and $x_j \in \{0, 1\}$. Let the endogenous law of motion for the aggregate state be $\mu' = \Psi(\mu)$ and let the distribution of idiosyncratic productivities be $w(s, \mu)$. When forming their conditional expectation in (29), consumers know the endogenous functions that map the aggregate state μ into p, r, π, in equilibrium. Write the policy functions of an individual consumer as $c(s, \mu, p, r, \pi), n(s, \mu, p, r, \pi), x(s, \mu, p, r, \pi)$ and $a' = g(s, \mu, p, r, \pi)$.

Producer's problem: Symmetric producers of x-goods take as given aggregate demand $\bar{x}(\mu, p, r, \pi) := \int x(s, \mu, p, r, \pi)d\mu(s)$ and solve a static profit maximization problem:

$$\pi(\mu) = \max_p \left\{p\bar{x}[\mu, p, r(\mu), \pi(\mu)] - \bar{x}[\mu, p, r(\mu), \pi(\mu)]^{1/\alpha}\right\}.$$

Let $p(\mu)$ denote the profit-maximizing optimal price.

Market clearing: Labor markets clear if

$$\bar{x}[\mu, p(\mu), r(\mu), \pi(\mu)]^{1/\alpha} = \int \{w(s, \mu)n[s, \mu, p(\mu), r(\mu), \pi(\mu)] - c[s, \mu, p(\mu), r(\mu), \pi(\mu)]\}d\mu(s).$$

Asset markets clear if

$$\int g[s, \mu, p(\mu), r(\mu), \pi(\mu)] d\mu(s) = 0.$$

A *recursive equilibrium* is a law of motion Ψ, individual functions v, c, x, n, g, pricing functions p, r and profit function π such that (i) v, c, x, n, g solve the consumer's problem, (ii) p and π solve the producer's problem, (iii) r clears the competitive asset market, and (iv) Ψ is generated by g and the exogenous distribution of idiosyncratic productivities.

Solving the model. The key computational difficulty in solving a model of this kind is the presence of the distribution μ, a high-dimensional object, as a state variable. Krusell and Smith (1998) propose an approach to solving such models where the distribution μ is summarized by its mean. Implementing their algorithm to solve our model is not straightforward.

An obvious issue is that the mean of the productivity distribution is no longer a nearly-sufficient statistic for the distribution itself. All the novel effects of the model come from the second moment of the productivity distribution. At the very least, the variance of the distribution would be a second state variable. Although that would make the problem more unwieldy, it is not the biggest hurdle.

The key technical difference is that in Krusell and Smith (1998), the interest rate r is a *known* function of the mean of the capital stock distribution (from a competitive marginal product condition). Because of their Cobb-Douglas technology, the interest rate does not depend on any other properties of the distribution μ. These assumptions ensure that the distribution μ is relevant only because households need to know the future distribution μ' to forecast future interest rates. It is not needed to determine current period economic conditions.

In our model with endogenous risk sharing, not only is the distribution of μ' is relevant for forecasting future interest rates (and profits), it is also directly needed to solve the x good producer's current problem and hence to determine current consumption of both goods and labor supply. The interest rate r and lump-sum profit π functions must be simultaneously determined with all the other equilibrium objects as part of a larger fixed point problem. Thus, even the within-period optimization problem is a difficult fixed-point problem.

Finally, because the distribution itself plays a more central role in our problem than in Krusell and Smith (1998), it is unlikely that their computational approach – summarizing the distribution with a small number of moments – will deliver a close approximation to the model's solution.

References

Berry, S., Levinsohn, J., Pakes, A. 1995. Automobile prices in market equilibrium. Econometrica 63(4), 841–890.

Heathcote, J., Storesletten, K., Violante, G. 2006. The macroeconomic implications of rising wage inequality in the United States. Unpublished Manuscript, NYU.

Krueger, D., Perri, F. 2005. Does income inequality lead to consumption inequality? Evidence and theory. Review of Economic Studies 73, 163–193.

Krusell, P., Smith, A.A. 1998. Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy 106(5), 867–896.

McFadden, D. 1973. Conditional logit analysis of qualitative choice behavior. In: Zarembka, P. (Ed.), Frontiers of Econometrics. Academic Press, New York.

Nevo, A. 2000. A practitioner's guide to estimation of random-coefficients logit models of demand. Journal of Economics and Management Strategy 9(4), 513–548.

Storesletten, K., Telmer C., Yaron, A. 2004. Cyclical dynamics in idiosyncratic labor market risk. Journal of Political Economy 112(3), 695–717.

Train, K.E. 2003. Discrete choice methods with simulation. Cambridge University Press, Cambridge.

Tables

Table E.1
Sensitivity of the benchmark model to alternative parameters.

Diminishing returns	low α model ($\alpha = 0.12$)			high α model ($\alpha = 0.48$)		
	level	relative std dev	correlation	level	relative std dev	correlation
x-good markup	1.28	0.95	−0.38	1.35	0.75	−0.09
profit share	0.46	0.96	0.66	0.21	0.93	0.80
labor	0.33	0.44	0.67	0.33	0.52	0.99
real wages	4.40	0.64	−0.26	6.80	0.24	0.89

Weight on x-goods	low ν model ($\nu = 50$)			high ν model ($\nu = 200$)		
	level	relative std dev	correlation	level	relative std dev	correlation
x-good markup	1.23	0.66	−0.16	1.23	0.66	−0.17
profit share	0.32	0.78	0.82	0.33	0.84	0.79
labor	0.33	0.46	0.97	0.33	0.44	0.95
real wages	5.90	0.24	0.40	5.70	0.28	0.18

Productivity	low $\bar{z}$ model ($\bar{z} = 5.80$)			high $\bar{z}$ model ($\bar{z} = 8.50$)		
	level	relative std dev	correlation	level	relative std dev	correlation
x-good markup	1.19	1.38	−0.08	1.28	0.49	−0.05
profit share	0.22	0.50	0.88	0.40	1.00	0.82
labor	0.21	0.55	0.99	0.39	0.36	0.81
real wages	6.30	0.37	0.99	4.90	0.58	−0.48

Alternative levels of diminishing returns in producing x-goods α, utility weight on x-goods ν, and level of aggregate productivity $\bar{z}$. Standard deviations are the standard deviation of the variable in logs divided by the standard deviation of log GDP. Similarly correlations are of the variable in logs with log GDP.

Table E.2

Key moments of the model, disaggregated by sector.

Markups	x-sector	c-sector	aggregate
level	1.23	1.00	1.10
relative std dev	0.63	0.00	0.29
correlation	−0.19	0.00	−0.19
Profit Share			
level	0.81	0.00	0.33
std dev	2.71	0.00	0.82
corr	0.92	0.00	0.81
Expenditure Share			
level	0.40	0.60	1.00
relative std dev	0.93	0.64	0.00
correlation	0.76	−0.76	0.00

Relative standard deviations are the standard deviation of the variable in logs divided by the standard deviation of log GDP. Similarly correlations are of the variable in logs with log GDP. All moments are for the benchmark model using the parameters given in table 1 of the main text.

Table F.1
Business cycle volatility as dispersion increases.

Decade	model (benchmark σ)	model (low σ)	data
50s	1.50		2.1
60s	1.49		1.3
70s	1.55	2.10	2.2
80s	1.55	1.68	1.7
90s	1.53	1.70	0.9
00s	1.50		1.0

Volatility measured as standard deviation of $\log(y_t)$ in percent. The trend added to the model is the decade-by-decade increase in earnings dispersion estimated by Heathcote et al. (2006). 'Low σ' refers to the model calibrated to match the rise in consumption dispersion as estimated by Krueger and Perri (2005). Both cases use the balanced growth correction described in the text. All other parameter values are listed in table 1 of the main text. Data are standard deviations of quarterly log real GDP, by decade, in percent.

Figure

Figure C.1
Markup increasing in dispersion, random coefficients demand.

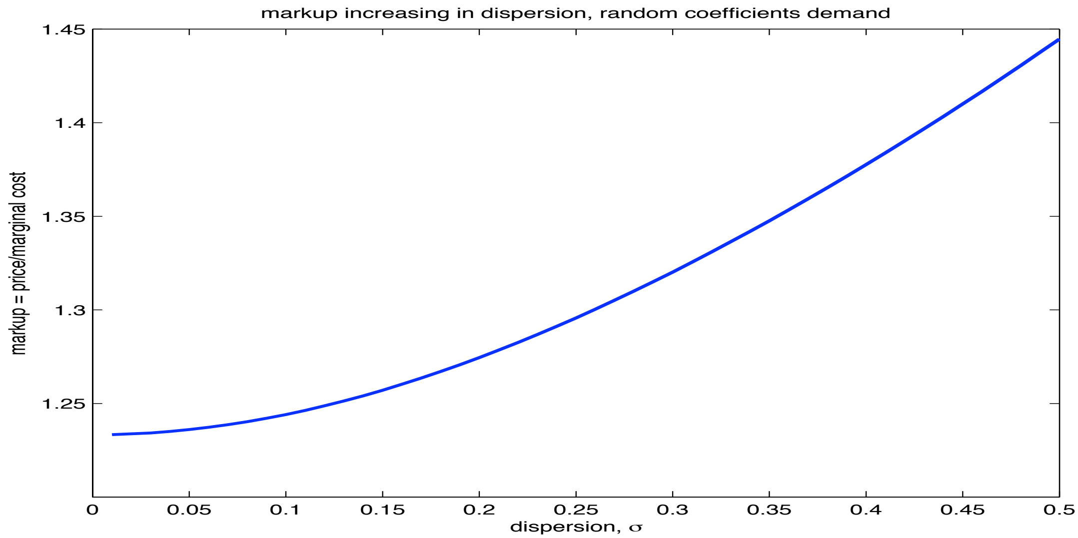

The average markup (price/marginal cost) for x goods increases in earnings dispersion σ, in a BLP-type random coefficients model where there is variable x-good demand in the sense that consumers can choose not to consume any of the x goods and instead simply consume more of the numeraire c good (so there may be expenditure switching).