

NBER WORKING PAPER SERIES

INCOME DISPERSION AND COUNTER-CYCLICAL MARKUPS

Chris Edmond
Laura Veldkamp

Working Paper 14452
<http://www.nber.org/papers/w14452>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2008

We thank George Alessandria, David Backus, Mark Bilts, Mark Gertler, Carlos Garriga, Aubhik Khan, John Leahy, Emi Nakamura, Nick Souleles, Harald Uhlig, Lawrence Uren, Stijn Van Nieuwerburgh, Pierre-Olivier Weill, Michael Woodford, Randy Wright, and seminar participants at Columbia, Tokyo, Princeton, UCLA, NYU, Philadelphia FRB, Rochester, Iowa, Oslo, Melbourne and the 2008 Washington University in St Louis conference on money, credit and policy, the 2007 New York area monetary policy meetings, 2007 AEA and 2006 SED meetings for helpful comments and conversations. We especially thank Jeff Campbell whose detailed suggestions greatly improved the paper. Laura Veldkamp also thanks Princeton University for their hospitality and financial support through the Kenen fellowship. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Chris Edmond and Laura Veldkamp. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Income Dispersion and Counter-Cyclical Markups
Chris Edmond and Laura Veldkamp
NBER Working Paper No. 14452
October 2008
JEL No. E32

ABSTRACT

Recent advances in measuring cyclical changes in the income distribution raise new questions: How might these distributional changes affect the business cycle itself? We show how counter-cyclical income dispersion can generate counter-cyclical markups in the goods market, without any preference shocks or price-setting frictions. In recessions, heterogeneous labor productivity shocks raise income dispersion, lower the price elasticity of demand, and increase imperfectly competitive firms' optimal markups. The calibrated model explains not only many cyclical features of markups, but also cyclical, long-run and cross-state patterns of standard business cycle aggregates.

Chris Edmond
Department of Economics
Stern School of Business
New York University
44 West 4th Street
New York NY 10012
cedmond@stern.nyu.edu

Laura Veldkamp
Stern School of Business
New York University
44 W Fourth Street, Suite 7-77
New York, NY 10012
and NBER
lveldkam@stern.nyu.edu

A long line of empirical research suggests that prices vary less over the business cycle than marginal costs. In other words, markups are counter-cyclical. The question is why. We argue that the cross-sectional dispersion of earnings might play a role. In recessions, when earnings are more dispersed, buyers' willingness to pay is also more dispersed. If sellers reduce prices in recessions, they attract few additional buyers (the small shaded area in the left panel of figure 1). This low elasticity makes the marginal benefit of lowering prices smaller and induces firms to keep prices high. Therefore when dispersion is high, prices stay high but profits are low. In contrast, in booms when dispersion is low, sellers who reduce prices attract many additional buyers (the larger shaded area in the right panel of figure 1). Therefore in booms, sellers keep prices low but earn high profits.

While there have been many previous explanations for counter-cyclical markups, the mechanism we propose has two strengths: It is based on observables and can be embedded in a simple dynamic equilibrium model.¹ The observable variable is earnings dispersion. Embedding the earnings process estimated by Storesletten, Telmer, and Yaron (2004) in a production economy allows us to compare the model's predictions to business cycle aggregates. In particular, we deliver realistic pro-cyclical profit shares, a feature of the data that many models struggle with. In addition, the model can also explain long-run trends and cross-state variations in profit shares.

To illustrate our mechanism, section 1 analyzes a static version of the model. There is a competitive sector where price equals marginal cost and an imperfectly competitive sector where prices are marked up. In both sectors, the only input is effective labor. Households choose how much to work and how much of each good to buy. Income dispersion arises because some households are more productive. The main result is that more dispersed idiosyncratic productivity results in higher markups and higher prices.

Theory alone cannot tell us if the variation in earnings dispersion is a plausible source of counter-cyclical markups. The problem is that changes in aggregate productivity are a force

¹Seminal papers on counter-cyclical markups are Rotemberg and Saloner (1986) and Bilal (1989). For a review, see Rotemberg and Woodford (1999). Recent work on the related phenomenon of real price rigidity includes Nakamura and Steinsson (2006) and Menzies (2007).




Figure 1: LOWERING PRICE IS MORE BENEFICIAL WHEN DISPERSION IS LOW.

The shaded area represents the increase in the probability of trade from lowering the price, by an amount equal to the width of the shaded area. This higher probability, times the expected gains from trade, is the marginal benefit to reducing the price. Willingness to pay is based on households' earnings.

for pro-cyclical markups. Therefore section 2 calibrates and simulates a dynamic version of the model. For reasonable parameter values, we show that the earnings dispersion effect dominates the productivity effect. Since measured dispersion is counter-cyclical, markups are as well. Their correlation with GDP is almost as negative as in the data. The resulting prices look inflexible because they fluctuate less than marginal cost. Yet, there are no price-setting frictions.

One of the reasons economists pay attention to counter-cyclical markups is because they can amplify the effects of other business cycle shocks. In this model, when aggregate productivity is low, high markups keep prices from falling much. Higher prices mean fewer goods are sold, amplifying the effect of the productivity shock. In our quantitative results, the effect of the productivity shock is amplified seven-fold relative to a standard real business cycle model. Section 2.5 compares the model's predictions for GDP, employment, real wages, and profits to their empirical counterparts. Importantly, the model's ability to explain markups does not come at the cost of undermining its ability to match macroeconomic aggregates.

To keep heterogeneous earnings tractable, our model abstracts from important issues debated in the literature on income heterogeneity and welfare, such as risk sharing and capital accumulation (Krusell and Smith 1998, Rios-Rull 1996, Krueger and Perri 2005). In section 2.7 we show that our main results hold up if we re-calibrate idiosyncratic productivity

to the level of consumption dispersion documented in Krueger and Perri (2005). Omitting capital hurts the performance of the model by making aggregates too correlated with GDP.

A number of other mechanisms can generate counter-cyclical markups. One possibility is that sticky prices and pro-cyclical marginal costs make the difference between price and cost, the markup, counter-cyclical. The problem with this explanation is that, without additional labor market frictions, it implies counter-cyclical firm profits, strongly at odds with the data. Similarly, while firm entry and exit change the degree of market competition and thus the markup (Jaimovich 2006), free entry implies zero profits. Our model delivers the observed pro-cyclical profits. Booms are times when markups are low but volume is high enough to compensate. In Comin and Gertler (2006), the causality is reversed: They use shocks to markups as the source of business cycle fluctuations. Three models closely related to ours also produce a cyclical elasticity of demand due to changing production technology (Kimball 1995), changing demand composition (Gali 1994), or a change in product variety (Bilbiie, Ghironi, and Melitz 2006).

To argue that earnings dispersion is at least part of the reason for price variation, we look for other evidence that long-run changes and cross-sectional differences in earnings dispersion are correlated with differences in prices, profit shares, and output volatility, as predicted by the model. Section 3 shows that the observed increase in earnings dispersion is consistent with the observed slow-down in real wage growth and the accompanying increase in profit shares, and can generate a modest decline in business cycle volatility. Section 4.1 uses state-level panel data to test the model's predicted relationships between earnings dispersion and profit shares. Section 4.2 documents additional facts from the empirical pricing literature that when the customer base has more dispersed earnings, prices tend to be higher.

Our mechanism takes counter-cyclical earnings dispersion as a given.² But this raises an obvious question: Why does earnings dispersion rise in a recession? One explanation is that job destruction in recessions is responsible (Caballero and Hammour 1994). Rampini

²This is the general consensus in the literature, as exemplified by Storesletten, Telmer, and Yaron (2004) whose estimates we use. One exception is Barlevy and Tsiddon (2006). They come to the opposite conclusion because they consider the reduction in inequality in the years following the great depression.

(2004) argues that entrepreneurs' incentives change in recessions, making firm outcomes and owners' earnings more risky. Cooley, Marimon, and Quadrini (2004) and Lustig and Van Nieuwerburgh (2005) argue that low collateral values inhibit risk-sharing in recessions. Any one of these explanations could be merged with our mechanism to produce a model whose only driving process is aggregate productivity shocks.

1 An illustrative static model

Households. There is a continuum of households indexed by i with identical preferences over a numeraire consumption good c_i , labor supply n_i and a continuum of products x_{ij} indexed by $j \in [0, 1]$,

$$U_i = \log(c_i) - \theta n_i + \nu \int_0^1 x_{ij} dj, \quad \theta, \nu > 0. \quad (1)$$

Let w_i denote a household's idiosyncratic effective labor productivity and let p_j denote the price of good j , both in terms of the numeraire. The budget constraint for household i is:

$$c_i + \int_0^1 p_j x_{ij} dj \leq w_i n_i + \pi, \quad (2)$$

where π denotes lump-sum profits paid out by firms.

To simplify aggregation, we assume that each of the x -goods is indivisible. A household either buys an x -good or not, $x_{ij} \in \{0, 1\}$. Appendix A shows that our main results go through if households can demand any $x_{ij} \geq 0$.

Firms and market structure. The economy consists of a continuum of island locations. At each location is one firm that produces the x -good (a monopolist) and a large number of identical perfectly competitive firms that produce the numeraire c -good. Each island receives an IID random assignment of a unit mass of households drawn from the population. Each household supplies labor to a competitive labor market on its island. Producers of x and c goods hire that labor. Both types of goods are produced by a technology that transforms

effective labor 1-for-1 into final products. Since the aggregate supply of effective labor on an island is $\int_0^1 w_i n_i di$ the labor market clears on island j when

$$\int_0^1 c_i di + \int_0^1 x_{ij} di = \int_0^1 w_i n_i di. \quad (3)$$

The monopolist producer of an x -good on island j chooses its price to maximize profits π_j . Profits are price p_j times the quantity sold at that price $x(p_j)$ less cost:

$$\pi_j = (p_j - 1)x(p_j). \quad (4)$$

Since the competitive firms make zero profits, aggregate profits are $\pi := \int_0^1 \pi_j dj$. Each household gets an equal share of these aggregate profits.³

Notice that which island a household is assigned to is immaterial since the x -goods are perfect substitutes to households. Islands are identical; the only role they play in the analysis is in ensuring that each x -good producer is a (local) monopolist.

Earnings Dispersion. Heterogeneous effective labor productivity is the source of earnings dispersion. The distribution of productivity is summarized by its mean $z > 0$ and a measure of dispersion $\sigma > 0$. We write $w_i = z + \sigma \varepsilon_i$ where ε_i has mean zero and is IID in the population with probability density $f(\varepsilon)$ and cumulative distribution $F(\varepsilon)$ on support $[\underline{\varepsilon}, \bar{\varepsilon}]$. We suppose $\underline{\varepsilon} > -z/\sigma$ so that $w_i > 0$ for all i and that:

ASSUMPTION 1. The hazard $h(\varepsilon) := f(\varepsilon)/(1 - F(\varepsilon))$ is monotone increasing on $(\underline{\varepsilon}, \bar{\varepsilon})$.

As discussed by Bagnoli and Bergstrom (2005), this is equivalent to assuming that the survival function $1 - F(\varepsilon)$ is log-concave⁴ on $(\underline{\varepsilon}, \bar{\varepsilon})$.

³These profits are available to a household in time to be used as income for c and x purchases. One interpretation of this is that households are able to make some consumption purchases on credit backed by the (deterministic) profit income they will receive and then subsequently repay their creditors when the profit income is paid out by firms.

⁴That is, $\log(1 - F(\varepsilon))$ is concave. A sufficient condition for $1 - F(\varepsilon)$ to be log-concave is for the density $f(\varepsilon)$ to be log-concave. Examples of distributions with log-concave densities include the uniform, normal, exponential, logistic, extreme value and members of the gamma and beta families. Linear transformations of log-concave distributions are also log-concave.

Equilibrium. An equilibrium in this economy is: (i) a set of consumption choices c_i and x_{ij} and labor supply choices n_i for each household that maximize utility (1) subject to the budget constraint (2), (ii) a price p_j for the monopolist x -good producer on each island j that maximizes profit (4) taking as given the demand for the firm's product such that (iii) the markets for c -goods, x -goods, and labor (3) all clear on each island j .

Results. Optimal consumption of the x -goods follows a cutoff rule, household i buys the x -good on island j if the additional utility it provides exceeds the price p_j times the household's Lagrange multiplier on (2), i.e., if $\nu \geq p_j \lambda_i$. The first order condition for labor supply tells us that $\lambda_i = \theta/w_i$. Combining these two expressions yields the household's demand for the x -good:

$$x_{ij} = \begin{cases} 1 & \text{if } w_i \geq \frac{\theta}{\nu} p_j \\ 0 & \text{otherwise} \end{cases}. \quad (5)$$

The fraction of households who buy a differentiated product is just the probability that each household has a labor productivity higher than the cutoff value, so the demand curve facing an x -good producer is:

$$x(p_j) := \int_0^1 x_{ij} di = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \mathbb{1} \left\{ \varepsilon_i \geq \frac{\theta p_j / \nu - z}{\sigma} \right\} f(\varepsilon_i) d\varepsilon_i = 1 - F \left(\frac{\theta p_j / \nu - z}{\sigma} \right). \quad (6)$$

Differentiating the profit function (4) with respect to p_j yields the first order condition characterizing the profit-maximizing price:

$$p_j + \frac{x(p_j)}{x'(p_j)} = 1. \quad (7)$$

The left hand side is the firm's marginal revenue, the right hand side its constant marginal cost. Using the expression for the demand curve (6) and rearranging gives:

$$p_j - 1 = \frac{1 - F \left(\frac{\theta p_j / \nu - z}{\sigma} \right) \nu \sigma}{f \left(\frac{\theta p_j / \nu - z}{\sigma} \right) \theta} = \frac{1}{h \left(\frac{\theta p_j / \nu - z}{\sigma} \right)} \frac{\nu \sigma}{\theta}, \quad (8)$$

where $h(\varepsilon) := f(\varepsilon)/(1 - F(\varepsilon))$ is the hazard rate of the distribution of idiosyncratic labor productivity. By assumption, $h(\varepsilon)$ is monotone increasing. Therefore the right hand side of (8) is monotone decreasing in p_j while the left hand side is monotone increasing in p_j . The unique intersection of the two curves determines the optimal price set on island j . In a symmetric equilibrium this price is the same on every island, $p_j = p$ for all j .

Our interest here is in how the optimal markup varies with the parameters of the distribution of idiosyncratic labor productivity z, σ . Since marginal cost is constant and normalized to 1, the optimal markup is equal to the optimal price.

PROPOSITION 1. The optimal markup $m(z, \sigma)$ is increasing in aggregate productivity z and increasing in dispersion σ .

The formal proof is in appendix A. To get intuition for this result, it's instructive to use (6) to calculate the elasticity of demand:

$$\epsilon(p_j) := -\frac{x'(p_j)p_j}{x(p_j)} = 1 + h\left(\frac{\theta p_j/\nu - z}{\sigma}\right) \frac{\theta}{\nu\sigma}. \quad (9)$$

Both an increase in aggregate productivity z and an increase in dispersion σ reduce the elasticity of demand. When the elasticity of demand falls, lower prices generate few additional sales so the optimal markup and price rise.

Example. Proposition 1 holds for any distribution of idiosyncratic productivity $f(\varepsilon)$ with a (weakly) increasing hazard. To illustrate the economics a little further, we turn to a special case that can be worked out explicitly. Let ε_i be IID uniform on $[-1, +1]$ so that w_i is IID uniform on $[z - \sigma, z + \sigma]$ with constant density $1/2\sigma$. Then demand for x -goods on island j is linear in the price:

$$x(p_j) = \int_{\theta p_j/\nu}^{z+\sigma} \frac{1}{2\sigma} dw_i = \frac{z + \sigma}{2\sigma} - \frac{\theta/\nu}{2\sigma} p_j. \quad (10)$$

This demand curve implies the optimal markup as a function of the parameters of the distribution of idiosyncratic labor productivity:

$$m(z, \sigma) = 1 + \frac{1}{2} \left(\frac{z + \sigma}{\theta/\nu} - 1 \right). \quad (11)$$

Firms only produce if they earn non-negative profits, which is when $m(z, \sigma) \geq 1$. To ensure this we assume that marginal cost is sufficiently low: $1 \leq (z + \sigma)\nu/\theta$. If this assumption were violated, no firm would produce.

The elasticity of demand at the optimal price is:

$$\epsilon(z, \sigma) = \frac{\frac{z+\sigma}{\theta/\nu} + 1}{\frac{z+\sigma}{\theta/\nu} - 1}. \quad (12)$$

This elasticity is decreasing in both aggregate productivity and dispersion. An increase in aggregate productivity z shifts out and steepens the firm's marginal revenue curve, this leads to higher sales of x -goods and higher markups and prices as the firm uses its monopoly power to capture a share of the higher surplus generated by the additional demand. By contrast, an increase in dispersion σ shifts in but also steepens the firm's marginal revenue curve so that sales of x -goods fall but markups and prices rise. So as in Proposition 1, either higher z or σ increase markups and prices but higher z causes higher x -good sales while higher σ causes lower x -good sales.

Counter-cyclical markups? If business cycles involved only changes in productivity, then this mechanism predicts that markups would be pro-cyclical; an increase in z would increase markups. But in the data, earnings dispersion is counter-cyclical, suggesting σ falls when z rises. Can this off-setting force be strong enough to explain counter-cyclical markups? To answer this, section 2 builds a dynamic quantitative model.

2 A dynamic quantitative model

Our dynamic model departs from the static model in four ways. First, aggregate productivity z and dispersion σ fluctuate. Second, idiosyncratic labor productivity has a realistic lognormal distribution.⁵ Third, marginal cost is variable instead of constant, so that firm

⁵The lognormal distribution does not have an increasing hazard function and so is not covered by Assumption 1. The assumption of an increasing hazard function is sufficient but not necessary for an increase in dispersion to increase markups. More details available on request.

profit shares are realistic. Fourth, richer preferences deliver a more realistic wealth effect on the labor supply.

In the model, profits rise in booms. With a strong wealth effect on labor, cyclical profits can make labor counter-cyclical. Although other models encounter this problem, it is particularly acute here because imperfect competition in x goods makes profits larger and more volatile. We use “GHH” preferences (Greenwood, Hercowitz, and Huffman 1988) that eliminate the wealth effect on labor supply to deliver more realistic labor fluctuations.

We omit capital and other assets to keep the model computationally tractable. Since households have no opportunity to share risk or smooth consumption, this assumption could distort our results. In section 2.7 we gauge the effect of this distortion by re-calibrating the model to match consumption data, which incorporates the effect of financial income, savings and transfers.

Our dynamic model is not intended to be a full model of business cycle fluctuations. Rather, it shows that with realistic parameters, our dispersion mechanism can generate counter-cyclical markups, the magnitude of the markup fluctuation is not trivial, and that including the mechanism does not undermine the model’s ability to match standard macroeconomic aggregates.

2.1 Model setup

Individuals have GHH preferences over the numeraire consumption good c_i and labor n_i and get additive utility from the x_{ij} goods:

$$U_i = \log \left(c_i - \theta \frac{n_i^{1+\gamma}}{1+\gamma} \right) + \nu \int_0^1 x_{ij} dj,$$

which they maximize subject to their budget constraint (2).

The log of aggregate productivity is an AR(1) process:

$$\log(z_t) = (1 - \rho) \log(\bar{z}) + \rho \log(z_{t-1}) + \varepsilon_{zt}, \quad \varepsilon_{zt} \sim \mathcal{N}(0, \sigma_z^2). \quad (13)$$

Idiosyncratic labor productivity is lognormal, $\log(w_{it}) = \log(z_t) + \varepsilon_{it}$ where $\varepsilon_{it} \sim \mathcal{N}(0, \sigma_t^2)$. Our model of idiosyncratic productivity follows Storesletten, Telmer, and Yaron (2004) who estimate an earnings process with persistent and transitory shocks. Let:

$$\begin{aligned}\varepsilon_{it} &= \xi_{it} + u_{it}, & u_{it} &\sim \mathcal{N}(0, \sigma_u^2) \\ \xi_{it} &= \rho_\xi \xi_{it-1} + \eta_{it}, & \eta_{it} &\sim \mathcal{N}(0, \sigma_{\eta t}^2).\end{aligned}\tag{14}$$

The key feature of the earnings process is that $\sigma_{\eta,t}^2$ increases when GDP is below its long-run mean, specifically $\sigma_{\eta,t}^2 = \sigma_H^2$ if $y_t \geq \bar{y}$ and $\sigma_{\eta,t}^2 = \sigma_L^2$ if $y_t < \bar{y}$, where $\sigma_H^2 < \sigma_L^2$, y_t is GDP, as defined in equation (17) below, and \bar{y} is its long-run mean.

Putting these elements together, our stochastic process for dispersion σ_t is given by:

$$\sigma_t^2 = \rho_\xi^2 \sigma_{t-1}^2 + (1 - \rho_\xi^2) \sigma_u^2 + \begin{cases} \sigma_H^2 & \text{if } y_t \geq \bar{y} \\ \sigma_L^2 & \text{if } y_t < \bar{y} \end{cases}.\tag{15}$$

Finally, we give x -good firms variable marginal costs. They transform effective labor into x -goods with a standard Cobb-Douglas technology, $x = n^\alpha$ with $0 < \alpha < 1$. Aggregate effective labor is $\int_0^1 w_i n_i di$ and the labor market clears on island j when $\int_0^1 c_i di + \int_0^1 x_{ij}^{1/\alpha} di = \int_0^1 w_i n_i di$. Profits for firm j are:

$$\pi_j = p_j x(p_j) - x(p_j)^{1/\alpha},\tag{16}$$

with variable marginal cost $x(p_j)^{(1-\alpha)/\alpha}/\alpha$ where $x(p_j)$ is the firm's aggregate demand curve.

Measuring GDP in the model. In order to calibrate and evaluate the model, we need a measure of total value-added to compare to GDP in the data:

$$y := \int_0^1 c_i di + \int_0^1 \int_0^1 p_j x_{ij} dj di.\tag{17}$$

GDP varies both because of changes in the production of each good and because of changes in the relative price of x -goods and c -goods.

2.2 Model solution

The first-order condition for labor choice tells us that labor depends only on the wage and on preference parameters:

$$n_i = \left(\frac{w_i}{\theta}\right)^{1/\gamma}. \quad (18)$$

This simple relationship, devoid of any wealth effect, is what GHH preferences are designed to deliver. These preferences also simplify our calibration procedure: they imply log earnings are proportional to log idiosyncratic productivity and so it is straightforward to match the empirical earnings distribution by a corresponding exogenous idiosyncratic productivity distribution. But GHH preferences complicate the model's solution because the cutoff rule for x -good demand is no longer linear in the wage. While household i still buys a unit of x_j if $\nu \geq p_j \lambda_i$, the Lagrange multiplier on their budget constraint is now $\lambda_i = \left(c_i - \theta \frac{n_i^{1+\gamma}}{1+\gamma}\right)^{-1}$. To derive the demand for x -goods, use (2) and (18) to substitute out c_i and n_i in the λ_i formula. Then, substitute λ_i into the cutoff rule at the indifference point ($p_j = \nu/\lambda_i$). This delivers a critical wage $\hat{w}(p_j)$ such that any household with wage higher than this threshold buys the good. Thus the aggregate demand curve facing the x -good producer on island j is $x(p_j) = \Pr[w_i \geq \hat{w}(p_j)]$.

Firms' prices are chosen to maximize profit (16) taking the aggregate demand curve as given. The first order condition for profit maximization equates marginal revenue and marginal cost,

$$p_j + \frac{x(p_j)}{x'(p_j)} = \frac{1}{\alpha} x(p_j)^{(1-\alpha)/\alpha}. \quad (19)$$

The set of equations that determine a solution to the model can no longer be solved in closed form. Appendix B details the fixed point problem solved in the following numerical analysis.

2.3 Calibration

Table 1 lists all parameters and their calibrated values. We choose the utility weight on leisure θ to match 33% of time spent working in steady state and the concavity of the x -sector technology α to match an aggregate labor share of 70%, both standard business cycle

Parameter	Calibration target			
utility weight on leisure	θ	15	steady state hours	0.33
concavity of production	α	0.24	steady state labor share	0.70
utility weight on x -goods	ν	100	steady state x -sector markup	30%
mean of productivity	\bar{z}	7.7	steady state aggregate markup	11%
inverse labor supply elasticity	γ	0.6	measured elasticity (GHH)	1.67
productivity innovation std dev	σ_z	0.0032	output std dev	0.017
productivity autocorrelation	ρ	0.80	output autocorrelation	0.80
transitory earnings std dev	σ_u	0.024	STY estimate (annual)	0.065
persistent earnings std dev $y > \bar{y}$	σ_H	0.012	STY estimate (annual)	0.032
persistent earnings std dev $y < \bar{y}$	σ_L	0.020	STY estimate (annual)	0.054
earnings autocorrelation	ρ_ξ	0.988	STY estimate (annual)	0.952

Table 1: PARAMETERS AND THE MOMENT OF THE DATA EACH PARAMETER MATCHES. Appendix B derives the steady state moments of the model. Appendix D details our transformation of STY moments from annual to quarterly.

calibration targets. The calibrated α differs from the typical value of 0.70 because in this two-sector model, the degree of diminishing returns to labor in one of the sectors is not equivalent to the aggregate labor share. Appendix C explores model results with higher and lower α 's. The second moments of the productivity process match the persistence and standard deviation of output as reported in Stock and Watson (1999).

Markups in the x -sector are defined as price divided by marginal cost. Estimates of markups vary widely, depending on the sector of the economy being measured. At the high end, Berry, Levinsohn, and Pakes (1995) and Nevo (2001) document markups of 27-45% for automobiles and branded cereals. For the macroeconomy as a whole, Chari, Kehoe, and McGrattan (2000) argue for a markup of 11%. Since the competitive c -sector has zero markup, the markup for the economy as a whole is the x -sector markup times the x -sector expenditure share. Our calibration uses the mean of productivity and the utility weight on x -goods (which determines the x -good expenditure share) to match both the x -sector and aggregate markup facts. In particular, we choose to match an x -sector markup of 30% and an aggregate markup of 11%.

The relationship between earnings dispersion and output is not something we can manipulate directly because both earnings and GDP are endogenous variables. Therefore, we

choose idiosyncratic labor productivity to fit the earnings data. Because log labor supply is proportional to log productivity (equation 18), productivity dispersion and hourly wage dispersion both have the same correlation with log output. Because an individual’s labor supply is positively correlated with their productivity, total earnings $w_i n_i$ have higher dispersion than productivity, by a factor of $(1 + \gamma)/\gamma = 1.60/0.60 = 2.67$. Therefore, our idiosyncratic productivity parameters are the STY estimates, transformed from yearly to quarterly, divided by 2.67. Appendix D gives further details.

To determine whether the economy is in the high or low dispersion state (σ_H or σ_L), we first simulate the model in one state and then check whether GDP is higher or lower than its steady-state level. If realized GDP is inconsistent with the dispersion state, we re-simulate with the correct dispersion parameter. The resulting correlations of dispersion and GDP are quite accurate: -0.29 in the model and -0.30 in the data.⁶

Issues in measuring dispersion. Storesletten, Telmer, and Yaron (2004)’s estimates have been controversial, because of the difficulty identifying transitory and permanent shocks. Guvenen (2005) and others argue that, because of unmeasured permanent differences in earnings profiles, the persistence of earnings shocks is overestimated. While this distinction is crucial in a consumption-savings problem, it is not relevant for our model, which has no savings. Whether earnings dispersion is persistent because each person gets persistent shocks or because new workers with more dispersed characteristics enter the sample — this does not matter to our seller who sets the price facing a distribution of willingness to pay. Thus both sides in this debate hold views consistent with our model’s predictions.

2.4 Results: counter-cyclical markups

Recessions are times when firms pursue low-volume, high-margin sales strategies. The correlation of markups and log GDP in the simulated model is -0.21 , compared to -0.27 in

⁶A simpler alternative procedure is to link σ_H or σ_L to productivity. Because productivity is exogenous, this would eliminate the need for an iterative procedure. However, this both increases the distance between the model and STY’s estimates and results in less counter-cyclical dispersion, which hurts the model’s performance. This distinction becomes an issue because GDP and productivity are not so tightly linked in this model as they are in more standard business cycle models.

the data (Rotemberg and Woodford 1999). The standard deviation of markups is 0.85 times the standard deviation of log output, compared to 0.36 in the data. Thus, markups are counter-cyclical and smoother than GDP, but more volatile than in the data. In contrast, in a perfectly competitive market, the markup would always be zero. Figure 2 illustrates a simulated time-series of markups.




Figure 2: SIMULATED MARKUPS, EARNINGS DISPERSION AND PRODUCTIVITY.

In the data, counter-cyclical markups have been documented by Rotemberg and Woodford (1999) using three different methods, by Murphy, Shleifer, and Vishny (1989) using input and output prices, by Chevalier, Kashyap, and Rossi (2003) with supermarket data, by Portier (1995) with French data and by Bils (1987) inferring firms' marginal costs. Besides their negative correlation with output, the other salient cyclical feature of markups is that they lag output. Figure 3 shows that the model's markup is negatively correlated as a lagging variable, but turns to positively correlated when it leads, just as in the data. The difference is that the model's markup must lead by 5 quarters, rather than 2 quarters, to achieve a positive correlation.

The reason that the model's markups are lagging is that the earnings dispersion process is highly persistent. In low-productivity periods, it is the shocks to the persistent component of earnings that become more volatile (equation 14). As these high-volatility shocks continue to arrive, the earnings distribution fans out. When productivity picks up and shocks become less volatile, there is enormous dispersion in the persistent component of earnings. It takes

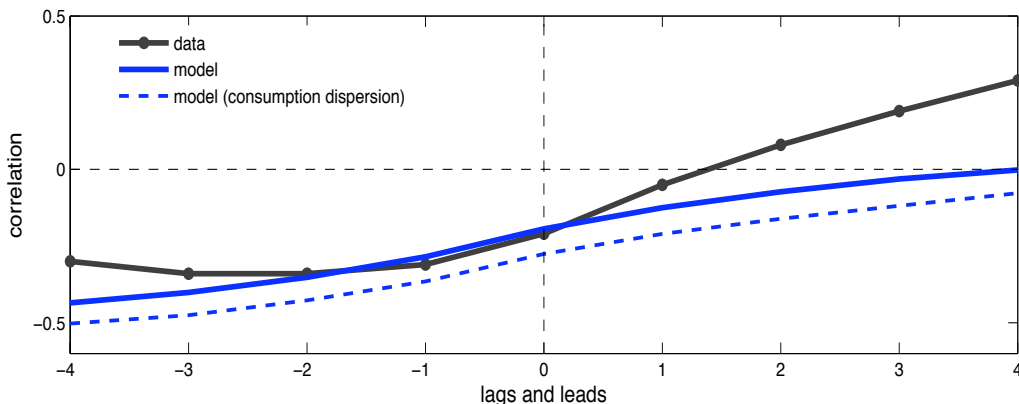


Figure 3: LEADS AND LAGS OF MARKUP-GDP CORRELATIONS.

Entries are $\text{corr}(\log(\text{markup}_t), \log(y_{t+k}))$. Positive numbers indicate leads and negative numbers indicate lags. Data from Rotemberg and Woodford (1999) (table 2, column 2). Markup is estimated using the labor share in the non-financial corporate business sector and an elasticity of non-overhead labor of -0.4 .

many periods of low-volatility shocks for the earnings distribution to become less dispersed. Since markups are driven by earnings dispersion, which is a lagging variable, markups are lagging as well. This feature of the model is similar to Bilbiie, Ghironi, and Melitz (2006). In their setting, a large fixed cost causes firms to delay entry. Since markups depend on how many firms enter, markups lag the cycle.

2.5 Does the model match standard business cycle moments?

Our explanation for counter-cyclical markups is not useful if it implies counter-factual fluctuations in macroeconomic aggregates. Of course, there are some facts that our model cannot speak to because of its simplicity. For example, we cannot compare consumption to output because without savings, the two are identical. But we do compare GDP to labor and to profits. Likewise, we cannot evaluate the properties of the model's prices with a measure like the CPI, which is the rate of exchange between goods and money, because the model has no money. But we do report a relative price, the real wage, which is the price of labor relative to the expenditure-weighted price index of x and c goods.

Table 2 compares the model aggregates to data. An important result is that the profit share is pro-cyclical (although too pro-cyclical and a little too volatile). Pro-cyclical profits

Model variable	relative std dev	corr with GDP
profit share	0.94	0.69
labor	0.49	0.96
real wages	0.28	0.22
Data variable	relative std dev	corr with GDP
profit share	0.80	0.22 (0.37)
labor (employment)	0.84 (0.82)	0.81 (0.89)
labor (hours)	0.97 (0.98)	0.88 (0.92)
real wages	0.39 (0.36)	0.16 (0.25)*
King and Rebelo (1999)	relative std dev	corr with GDP
profit share	0.00	0.00
labor	0.48	0.97
real wages	0.54	0.98

Table 2: SECOND MOMENTS OF AGGREGATE VARIABLES IN THE MODEL AND DATA. Standard deviations are divided by the standard deviation of GDP. Most statistics are from Stock and Watson (1999). Labor and wage numbers in parentheses are from Cooley and Prescott (1995). Number with an asterisk is from Rotemberg and Woodford (1999) All capital share statistics come from the labor share statistics reported in Gomme and Greenwood (1995). The second correlation, in parentheses, comes from NIPA data. But the NIPA-based measure counts all proprietors' earnings as profits, although it is part profit and part labor compensation. The first correlation corrects for this by removing proprietor's earnings.

distinguish this model from sticky price theories, models with free-entry or standard business cycle models such as King and Rebelo (1999). Labor and real wages do slightly less well, but not worse than the standard model. In the model, the real wage is the relative price of labor to the expenditure-weighted price index of x and c goods. Without a capital stock in the model, wages, labor and output are more driven by changes in productivity. This makes their correlations with output too high.

Amplification of business cycle fluctuations. Understanding counter-cyclical markups is important for business cycle research because the resulting prices are less flexible (less volatile); price rigidity amplifies the effects of productivity shocks on output. In our model, prices are only 2/3rds as volatile as they would be in a standard competitive economy where price equals marginal cost. If our prices were more flexible, they would fall further in recessions so that more x -goods would be sold. From table 2, it appears as though our model explains no more of macro volatility than the standard model. But the similarity is misleading. Recall that we calibrated our aggregate productivity process to match the

volatility of GDP. Our calibrated shocks are 1/7th as volatile as those in King and Rebelo (1999).⁷ Because our model’s recessions are deeper, using the King and Rebelo (1999) productivity process would make our business cycles many times more volatile.




Figure 4: REAL WAGES, EARNINGS DISPERSION AND GDP IN THE SIMULATED MODEL.

Figure 4 illustrates the behavior of real wages and GDP. It has two features that look familiar. First, real wages look like the mirror image of markups. The intuition for this is that wages are the main component of marginal costs and so wages relative to the x -good price behave like the reciprocal of the markup. Furthermore, both real wages and markups are closely correlated with dispersion. Rotemberg and Woodford (1999) argue that many empirical measures of markups are functions of the inverse real wage. The fact that simulated real wages are pro-cyclical means that this alternative measure of the model’s markups is counter-cyclical as well. Second, the measure of GDP looks quite similar productivity, plotted in figure 2. This tells us that, although dispersion has an effect, GDP is still primarily driven by aggregate productivity shocks.

⁷In our calibration, aggregate log productivity has a quarterly persistence of 0.80 and an innovation standard deviation of 0.032 which implies an unconditional quarterly standard deviation of log productivity that is a function of the persistence and volatility of the innovations: $0.0032/\sqrt{1-0.80^2} = 0.005$. In King and Rebelo (1999), the quarterly standard deviation of log productivity is $0.0072/\sqrt{1-0.979^2} = 0.035$.

2.6 Benchmark economies

Our results are driven by the earnings dispersion mechanism, not the preferences or the two-sector structure. Both the presence of earnings dispersion and its time-variation are essential for counter-cyclical markups. To show this, we compare our model to two benchmarks. The first is an economy where earnings dispersion is constant, always equal to its steady-state value. The second benchmark is an economy where there is no earnings dispersion, only a representative consumer. In both benchmarks, the correlation of markups with GDP is positive. This recalls our qualitative result from the static illustrative model, where holding dispersion fixed, an increase in aggregate productivity causes markups to increase.

When dispersion is constant and equal to its unconditional mean, many of our calibration targets look similar. The x -good markup (33%), the aggregate markup (12%), the average labor supply (0.33), average profit share (0.30), and the standard deviation and autocorrelation of output (0.017, 0.78) are all essentially unchanged. The key difference is the correlation of markups and GDP (0.23). Markups switched from being counter-cyclical to pro-cyclical. Alternatively, when earnings dispersion is zero, aggregates are either insufficiently volatile or almost perfectly cyclical.

2.7 Earnings vs. consumption dispersion

An important question is whether earnings dispersion is a meaningful measure of inequality to feed into the model. The reason we calibrate to earnings dispersion because we have reliable estimates of its cyclical properties. To measure dispersion in a number of business cycles requires a panel data set with a long time-series dimension. Either income, including capital income and transfers, or consumption are arguably more appropriate measures. But replacing earnings with income is unlikely to change our results: while labor earnings are only 63% of income for the average household, earnings and income dispersion have remarkably similar levels and cross-sectional variation (Diaz-Gimenez, Quadrini, and Rios-Rull 1997).

A more serious challenge is to replace earnings dispersion with consumption dispersion. Storesletten, Telmer, and Yaron (2004) claim that consumption dispersion is only 72% of

earnings dispersion. They create a long panel data set by using age characteristics of the PSID respondents to construct synthetic food consumption data back to 1930, just like they did with earnings. The long sample allows them to estimate a consumption process with counter-cyclical dispersion. Since, in our model, household earnings $w_i n_i$ and total consumption expenditure $c_i + \int p_j x_{ij} dj$ are identical (up to the common π), we reset our idiosyncratic productivity parameters (equation 14) to match our earnings process to STY's consumption expenditure process.⁸ We do not re-calibrate our other parameters. Table 3 shows that using consumption dispersion instead of income dispersion strengthens our main result: Markups become more counter-cyclical (-0.36 vs. -0.21) and much smoother (0.42 vs. 0.85). This is accomplished without sacrificing much fit with the other business cycle moments. The average markup falls from 30% to 24% on x -goods and from 11% to 10% for the average good. Profit shares and labor hours are unchanged.

Of course, food consumption is not an ideal measure of overall consumption. Krueger and Perri (2005) compare the dispersion of after-tax labor earnings, plus transfers, to the dispersion of consumption, which includes expenditures on non-durables, services, small durables, plus imputed services from housing and vehicles. They find that consumption dispersion was about 74% of earnings dispersion in 1980 and was about 67% of earnings dispersion in 2003.⁹ An alternative approach is to ask what a model with realistic risk-sharing predicts about the relative size of consumption and income dispersion. Aiyagari (1994)'s model has a coefficient of variation for consumption that is 50-70% that of income.

Every one of these findings suggests that consumption dispersion is between 50-75% of earnings dispersion. The second and third columns of table 3 show that when the model's earnings dispersion is scaled down to the level of consumption dispersion, the model's main results survive. Only when consumption dispersion is 25% of earnings dispersion, well below any estimates, do counter-cyclical markups and pro-cyclical real wages disappear.

⁸We use the persistence of the consumption dispersion process estimated by Storesletten, Telmer, and Yaron (2004) to set the persistence parameter ρ_ξ and we use their estimates of the variances of the persistent and transitory components of household consumption to set $\sigma_u^2, \sigma_H^2, \sigma_L^2$. See appendix D for details.

⁹Krueger and Perri (2005) report the following cross-sectional variances for consumption and earnings: 0.19/0.35 in 1980 and 0.25/0.55 in 2003. Because our dispersion is a standard deviation, we use the square root of these ratios.

		STY consumption	Amount of dispersion			Data
			75%	50%	25%	
markups	std dev	0.42	0.66	0.53	0.53	0.36
	corr GDP	-0.36	-0.16	-0.11	0.02	-0.27
profit share	std dev	0.78	0.83	0.73	0.67	0.80
	corr GDP	0.70	0.80	0.94	0.88	0.22
labor	std dev	0.44	0.46	0.42	0.38	0.90
	corr GDP	0.98	0.96	0.97	0.98	0.8-0.9
real wages	std dev	0.18	0.26	0.21	0.16	0.39
	corr GDP	0.32	0.33	0.12	-0.37	0.47

Table 3: LOW-DISPERSION MODEL RESULTS.

Numbered columns contain results from a model where $\sigma_H, \sigma_L, \sigma_u$ are re-calibrated so that steady state earnings is the listed percentage of its benchmark level. The “STY consumption” column uses the stochastic process for food consumption estimated by Storesletten, Telmer, and Yaron (2004). This process has both lower innovation variances and lower persistence. See appendix D for details. Standard deviations are relative to GDP.

The difficulty with endogenous risk sharing. Ideally, our model should include a consumption-savings choice, as in Aiyagari (1994) and Krusell and Smith (1998). Then, calibrated counter-cyclical earnings dispersion could endogenously generate the counter-cyclical consumption dispersion that moves markups. Appendix E sketches such a model and shows that the wealth distribution would be a state variable. Krusell and Smith (1998) approximate such a large-dimensional state variable with a small number of moments. However, their approach is unlikely to deliver a close approximation to our model’s true solution.

Krusell and Smith’s wealth distribution only matters because it forecasts this period’s interest rate and next period’s capital rental rate. These are *known* functions of the mean of the wealth distribution. Since this period’s mean wealth is a good forecaster of next period’s mean wealth, and thus future rental rates, keeping track only of only the mean of the wealth distribution results in a close approximation to the true solution.

In our model, not only does the wealth distribution forecast future interest rates, it also determines x -good producer’s prices, current consumption of both goods, labor supply, and profits. These are not known functions of the wealth distribution’s moments. Rather, they are determined by a fixed-point problem that uses all the information in the distribution. When higher moments become important to the solution of the model, the Krusell and Smith

(1998) technique can produce misleading results (Carroll 2000).

3 Evaluating long-run predictions

While our model is built to explain fluctuations at business cycle frequencies, there has been a long-run increase in the level of earnings dispersion that should cause low-frequency changes too. Earnings dispersion increased by 20% from 1967-1996, an average annual rate of 0.66% (Heathcote, Storesletten, and Violante 2006). To determine if our model's predictions are consistent with the long-run facts, we simulate six model economies, with different levels of earnings dispersion and productivity. Each economy represents a decade from the 1950's to the 2000's. We choose the 1970's to be the same as our benchmark calibration.

In doing this exercise, we run into a well-known problem. GHH preferences are inconsistent with balanced growth. The standard solution to this problem is to scale up the utility weight on leisure as productivity increases so as to keep hours flat. The formula for individual labor supply is $n_i = (w_i/\theta)^{1/\gamma}$. If we proportionately scale up θ with w_i , average hours will not change. In our model with linear preferences over the x -goods, we also need to scale up the utility weight ν . Changing these two parameters at the rate of productivity growth keeps both average hours and expenditure shares constant. Below we refer to results as having 'no correction' when we do not shift preference parameters and to 'balanced growth' results when we do.

3.1 Long-run slowdown in real wage growth

A long-run change that has been of particular concern to policy-makers is the slowdown in the growth of real wages. The left panel of figure 5 illustrates how real wages were keeping pace with labor productivity until the 1970's, when real wage growth slowed down. In the figure, real wages are measured as BLS real compensation per hour, including benefits, in the non-farm business sector, while labor productivity is BLS real output per hour. Both series come from the payroll survey.

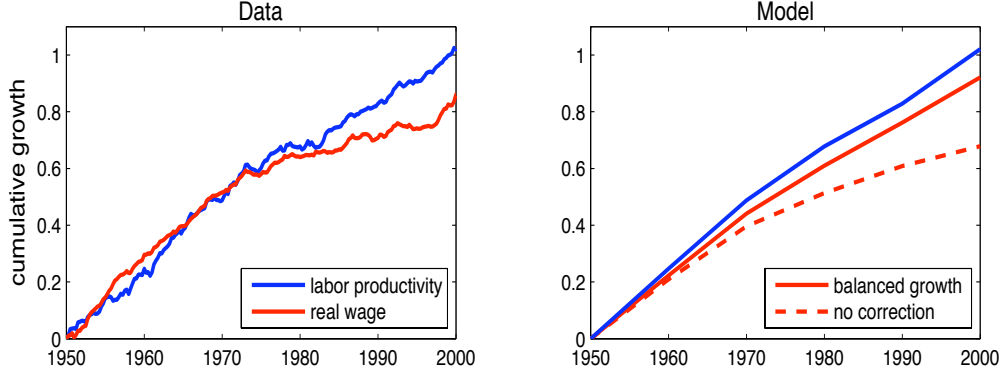


Figure 5: PRODUCTIVITY AND WAGE GROWTH IN THE DATA AND THE SIMULATED MODEL. The trend added to the model is the decade-by-decade increase in labor productivity and in earnings dispersion estimated by Heathcote, Storesletten, and Violante (2006). All other parameter values are listed in table 1. Data: real wages measured as BLS real compensation per hour in the non-farm business sector, labor productivity is BLS real output per hour. See appendix D for further details.

To ask the if the model produces the same effect, we need to calibrate not just the changes in earnings dispersion, but also the changes in aggregate productivity. To do this, we use annual estimates of labor productivity from the BLS, averaged by decade. The right panel of figure 5 shows a pattern in the model similar to that in the data. While our model with balanced growth preferences predicts only half the relative decline in real wages, the baseline model without correction produces an effect twice as strong as that in the data. In contrast, a standard business cycle model would predict that wages and productivity grow in tandem.

The flip side of the relative decline in real wages is an increase in firms' profit shares. The balanced growth model's share of output paid as profits rises steadily from 25% in 1950's to 35% (our calibrated value) in the 1970's to 38% in the 2000's. What happens is that higher dispersion reduces demand elasticity, prompting higher markups, and, in conjunction with higher productivity, this delivers higher profit shares. In the no correction model, rising productivity has stronger effects, making the rise in profit share more extreme (18% in 1950 to 73% in 2000). As higher productivity increases earnings, the composition of demand changes. Consuming more x -goods means consuming a broader variety of goods and is therefore not subject to the same diminishing marginal returns that set in when c -good consumption increases. Therefore, when earnings increase, x -good consumption rises more than c -good consumption. This effect is big. The expenditure share for x -goods is 22%

in 1950 and 75% in 2000. Higher demand for x -goods prompts firms to increase markups, raising profits.

In the data, the evidence on the size of the increase in profit shares is mixed. The share of output not paid out as labor earnings – a very broad definition of profits – rose only by about 5% from 1970-96 (NIPA data). Meanwhile, Lustig and Van Nieuwerburgh (2007) document that net payouts to security holders as a fraction of each firm’s value added – a much more narrow definition of profits – rose from 1.4% to 9% (based on flow of funds data) or 2.3% to 7.5% (NIPA data). While the broad measure suggests that our model over-predicts the rise in profits, the 3- to 6-fold rises in profits reflected in the latter measure suggest that the dramatic profit increases predicted by the model may not be so far from the truth.

3.2 The long-run decline in business cycle volatility

One of the most discussed low-frequency changes in the U.S. economy has been the decline of business cycle volatility. One might think that the increase in earnings dispersion over the last 50 years would make the model’s cycles more volatile because the individual earnings processes have become more volatile. This concern is not founded. 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, with the balance growth correction, it can generate a modest decline (see appendix F for details).

4 Evaluating cross-sectional predictions

The model’s key mechanism is that higher earnings dispersion increases markups but decreases the volume of goods traded. The change in volume is the key feature of the model that produces sizable business cycles out of small productivity shocks. This section looks for

cross-sectional evidence of that relationship and surveys related evidence in the industrial organization literature.

4.1 Testing the model with state-level panel data

Although we would like to test the hypothesis that more dispersion increases markups directly, we cannot because no state-level markup data is available. But we can measure the combined effect of higher markups and lower trade volumes by using the share of state GDP earned as firm profits, or its opposite, the labor share. If our mechanism is operative, then in years and U.S. states with similar productivity, the higher-earnings-dispersion state should have a lower profit share and a higher labor share.

The data is a panel of annual observations on 49 U.S. states from 1969-2004. For each state s and year t , our panel contains average labor productivity $z_{s,t}$, average labor share, and a measure of cross-county earnings dispersion $\sigma_{s,t}$. As a proxy for state labor productivity, we use real state GDP per employed worker. To measure a state's earnings dispersion, we take the log average salaries in each of the state's counties, weight them by the number of jobs in the county, and take their cross-sectional standard deviation. Of course, the cross-county dispersion is lower than cross-individual dispersion would be. The state labor share is total state nominal wages, salaries and supplements divided by state nominal GDP. We take logs of all variables and run the regressions in differences to remove any state-specific fixed effects. Appendix D details the data sources and transformations.

To determine if in years and U.S. states with similar productivity, the higher-earnings-dispersion state has a higher labor share, we estimate their relationship using a linear and a quadratic specification. To compare these results to our model, we simulate a panel of 5,000 economies followed for 200 quarters. We sample simulated data every four quarters and estimate the same annual relationships in the model.

In table 4, the null hypothesis that dispersion is uncorrelated with the labor share is rejected at the 99% confidence level. This is true after controlling for productivity, in both the linear and quadratic specifications. All of the signs of the coefficients are the same in

Dependent variable: labor share	Linear		Quadratic	
	data	model	data	model
dispersion	0.03 (5.59)	0.10	0.03 (5.91)	0.10
productivity	-0.67 (58.1)	-0.82	-0.66 (54.6)	-0.82
dispersion ²			-0.02 (3.06)	-0.05
productivity ²			-0.79 (5.90)	1.03
dispersion \times productivity			0.07 (0.07)	0.88

Table 4: EFFECT OF DISPERSION AND PRODUCTIVITY ON LABOR’S SHARE OF GDP. Both specifications are estimated with OLS on pooled data. The linear specification estimates: labor share = $\alpha + \beta_1$ dispersion_{s,t} + β_2 productivity_{s,t} + $\epsilon_{s,t}$, for state s in year t . The quadratic specification estimates: labor share = $\alpha + \beta_1$ dispersion_{s,t} + β_2 productivity_{s,t} + β_3 dispersion²_{s,t} + β_4 productivity²_{s,t} + β_5 dispersion_{s,t} \times productivity_{s,t} + $\epsilon_{s,t}$. See text and appendix D for further details. The t-statistic for the two-sided test of the null hypothesis of no significance in parentheses.

the model and the data, except for the productivity-squared term and all are statistically significant, except for the productivity-dispersion interaction term. The magnitudes of the coefficients are systematically larger in the model than in the data. In particular, a 1% increase in dispersion corresponds to a 0.03% increase in the labor’s share of GDP in the data and a 0.10% increase in the model.¹⁰ This is to be expected if the model has only dispersion and productivity as driving forces and the data has other sources of shocks as well as measurement error. In short, our prediction that higher dispersion is correlated with a fall in a state’s profit share holds up in the state panel data.

4.2 Evidence from empirical pricing studies

Our results are also qualitatively consistent with the findings of Chevalier, Kashyap, and Rossi (2003). Periods of good-specific high demand (e.g., beer on the fourth of July) are

¹⁰Comparing first and second derivatives of the labor share with respect to dispersion and productivity in the model and data corrects for the fact that the levels of the variables are not identical and that this can influence the second-order terms. This estimation yielded nearly identical results to the reported quadratic coefficients. Likewise, using the real wage instead of the profit share as a dependent variable yielded an equally good match between model and data. Finally, adding state-level fixed effects and running the regressions in levels made a negligible difference.

times when consumers' values for the goods are more similar. While one might expect that high demand would increase prices, the authors find that prices and markups fall. The same outcome would arise in our calibrated model if productivity dispersion σ falls.

Our model would also predict a higher markup when a good is sold to customers with more dispersed valuations. Studies of the effect willingness-to-pay dispersion has on car dealers' markups and sales supports this prediction. Using CES data, Goldberg (1996) estimates that blacks' valuations for new cars are more dispersed than whites' and that females' valuations are more dispersed than males'. She finds that, compared to the price paid by white males for the same car, black females (who have the most dispersed valuations) pay \$430 more, black males pay \$270 more and white females pay \$130 more, on average. While the standard errors on Goldberg's estimates are large, methodologically distinct studies, such as Ayers (1991), obtain almost identical estimates.¹¹

5 Conclusion

Our production economy is set up to capture the intuition that when earnings dispersion is higher, the price elasticity of demand is lower, so sellers optimally raise markups. However, without quantifying the model, the cyclical behavior of prices and markups is ambiguous because the productivity and earnings dispersion effects work in opposite directions. Using estimates of the time-series variation in the earnings distribution, we calibrate the model. Although the model is a simple one, it does a reasonable job of matching the business cycle features of markups and some traditional macro aggregates.

Our model provides a theory of real price rigidity, meaning prices that fluctuate less than marginal cost. By themselves, rigid real prices make business cycles more costly. When interacted with a form of nominal rigidity, real rigidities also amplify the real effects of

¹¹Goldberg (1996)'s price differentials are not statistically significant, which leads her to conclude that there is no evidence of discrimination. But the high standard errors come from dropping more than half the sample due a missing or inconsistent response. Ayers (1991) finds black males pay \$280 more than white males while Goldberg (1996) estimates that differential to be \$270. For the most imprecisely estimated case, white women, Ayers (1991) finds a differential of \$190 while Goldberg (1996) estimates \$130. See Harless and Hoffer (2002) and Ross (2003) for further discussion of these and related results.

nominal shocks (Ball and Romer 1990, Kimball 1995). Future work could merge this theory with a nominal price-setting friction. If such a model delivered enough amplification of small nominal shocks, it could further our understanding of monetary policy's role in the business cycle.

Appendix

A Analytic results

This appendix provides two analytic results, the formal proof of Proposition 1 and a calculation showing that indivisible x -goods are not essential for our main results.

A.1 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:

$$m(z, \sigma) - 1 = \frac{1}{h\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)} \frac{\nu \sigma}{\theta}, \quad (20)$$

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:

$$\frac{\partial m(z, \sigma)}{\partial z} = \frac{h'\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)}{h\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)^2 + h'\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)} \frac{\nu}{\theta} \geq 0, \quad (21)$$

and similarly:

$$\frac{\partial m(z, \sigma)}{\partial \sigma} = \frac{h'\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right) \left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right) + h\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)}{h\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)^2 + h'\left(\frac{\theta m(z, \sigma)/\nu - z}{\sigma}\right)} \frac{\nu}{\theta} > 0. \quad (22)$$

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. \square

A.2 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, we use 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:

$$x_i = \max\left(0, \frac{y_i}{p} - 1\right). \quad (23)$$

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:

$$x(p) := \int_0^1 x_i di = \int_{z-\sigma}^{z+\sigma} \max\left(0, \frac{y_i}{p} - 1\right) \frac{1}{2\sigma} dy_i. \quad (24)$$

Calculating the integral and simplifying gives:

$$x(p) = \int_p^{z+\sigma} \left(\frac{y_i}{p} - 1\right) \frac{1}{2\sigma} dy_i = \left[\frac{(z + \sigma)^2}{2p} + \frac{1}{2}p - (z + \sigma)\right] \frac{1}{2\sigma}. \quad (25)$$

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

$$\pi(p) := (p - 1)x(p), \quad (26)$$

taking as given the demand curve (25).

To ensure that the x -good producer will operate, we need:

ASSUMPTION 2. 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 2. 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 25). Now let's turn to interior solutions.

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

$$\pi'(p) = (p - 1)x'(p) + x(p) = 0. \quad (27)$$

Using the formula (25) for the demand curve, we get:

$$\pi'(p) = p - k + [(k/p)^2 - 1] / 2 = 0. \quad (28)$$

Similarly, the second order condition for a maximum is:

$$\pi''(p) = 1 - k^2/p^3 < 0. \quad (29)$$

We now show that 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 we see that the 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 $p \rightarrow k$ 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:

$$p'(k) \geq 0 \Leftrightarrow \frac{\partial}{\partial k} \{p - k + [(k/p)^2 - 1] / 2\} \geq 0, \quad (30)$$

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

$$p'(k) \geq 0 \Leftrightarrow k^{1/2} \geq p(k). \quad (31)$$

We now need to check if this condition holds. Evaluating the marginal profit function $\pi'(p)$ at $p = k^{1/2}$ we have:

$$\pi'(k^{1/2}) = k^{1/2} - k/2 - 1/2 =: A(k). \quad (32)$$

Now since $A(1) = 0$ and $A'(k) < 0$ for all $k > 1$ we have $\pi'(k^{1/2}) \leq 0$ 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 from (31) we conclude $p'(k) \geq 0$ with equality if and only if $k = 1$. Since $k := z + \sigma$ and the optimal markup is equal to the optimal price, we have $m(z, \sigma) = p(z + \sigma)$. Therefore the optimal markup is increasing in mean income z and increasing in income dispersion σ . \square

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.

B 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:

$$\lambda_i = \left(c_i - \theta \frac{n_i^{1+\gamma}}{1+\gamma} \right)^{-1}. \quad (33)$$

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 (33) 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 (2) to eliminate c_i and $\hat{n}(w_i)$ is given in (18). Solving for $\hat{p}(w_i)$ yields:

$$\hat{p}(w_i) = \frac{\nu}{1+\nu} \left[w_i \hat{n}(w_i) + \pi - \theta \frac{\hat{n}(w_i)^{1+\gamma}}{1+\gamma} \right]. \quad (34)$$

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:

$$x(p_j) = \Pr[w \geq \hat{w}(p_j)] = 1 - \Phi \left(\frac{1}{\sigma} \log \left(\frac{\hat{w}(p_j)}{z} \right) \right), \quad (35)$$

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 (34), 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 need to solve a fixed point problem of the form $\pi = G(\pi)$ where:

$$G(\pi) := \max_{p \geq 0} \left\{ p \left[1 - \Phi \left(\frac{1}{\sigma} \log \left(\frac{\hat{w}(p, \pi)}{z} \right) \right) \right] - \left[1 - \Phi \left(\frac{1}{\sigma} \log \left(\frac{\hat{w}(p, \pi)}{z} \right) \right) \right]^{1/\alpha} \right\}. \quad (36)$$

We do this 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. We solve for six parameters $(\alpha, \gamma, \nu, \theta, \sigma, \bar{z})$ such that the steady state of our model delivers the following six properties:

$$\begin{aligned} \text{elasticity of labor supply} &= \mathbb{E}[d \log(n_i)/d \log(w_i)] &= 1.67 \\ \text{earnings dispersion} &= \bar{\sigma}_{\text{STY}} &= 0.29 \\ \text{hours worked} &= \mathbb{E}[n_i] &= 0.33 \\ \text{labor share} &= \mathbb{E}[w_i n_i]/\bar{y} &= 0.70 \\ \text{aggregate markup} &= \bar{m} &= 1.11 \\ \text{x-sector markup} &= \bar{m}_x &= 1.30 \end{aligned}$$

We use the following properties of the model repeatedly: individual labor supply is, from (18), $n_i = (w_i/\theta)^{1/\gamma}$. Since log idiosyncratic productivity $\log(w_i)$ is normal with mean $\log(\bar{z})$ and standard deviation σ , log labor supply is:

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

and so:

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

Similarly, log earnings is:

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

and so:

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

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 $\bar{\sigma}_{\text{STY}} = 0.29$ when $\sigma = (0.60/1.60)(0.29) = 0.11$.

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

$$\mathbb{E}[n_i] = (\bar{z}/\theta)^{1/0.60} \exp[0.5(0.11/0.60)^2] = 0.33. \quad (37)$$

We also use labor's share $\mathbb{E}[w_i n_i]/\bar{y} = 0.70$ and since $\bar{y} = \mathbb{E}[w_i n_i] + \bar{\pi}$, we need to have $\mathbb{E}[w_i n_i] = (0.70/0.30)\bar{\pi}$. To calculate $\bar{\pi}$ we need to solve the fixed point problem $\bar{\pi} = G(\bar{\pi})$ outlined above. The solution of the fixed point problem depends on all four parameters and to acknowledge this write $\bar{\pi}(\alpha, \nu, \theta, \bar{z})$. Using the formula for average earnings derived above, we then have a second equation in the four unknowns:

$$\mathbb{E}[w_i n_i] = \bar{z}^{(1.60)/0.60} \theta^{-1/0.60} \exp[0.5((1.60)(0.11)/0.60)^2] = \frac{0.70}{0.30} \bar{\pi}(\alpha, \nu, \theta, \bar{z}). \quad (38)$$

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

$$\bar{m}_x(\alpha, \nu, \theta, \bar{z}) = 1.30. \quad (39)$$

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

$$\bar{e}_x(\alpha, \nu, \theta, \bar{z}) = \frac{\bar{m} - 1}{\bar{m}_x - 1} = \frac{0.11}{0.30}. \quad (40)$$

In short, we solve for the four parameters $(\alpha, \nu, \theta, \bar{z})$ by solving the four equations (37)-(40) simultaneously. This gives us $\alpha = 0.24, \nu = 100, \theta = 15$ and $\bar{z} = 7.68$.

C Sensitivity analysis

Diminishing returns parameter α . The top panel of table 5 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 5 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 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

Diminishing returns	low α model ($\alpha = 0.12$)			high α model ($\alpha = 0.48$)		
	level	std dev	corr	level	std dev	corr
x -good markup	28%	0.95	-0.38	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	4.4	0.64	-0.26	6.8	0.24	0.89
Weight on x-goods	low ν model ($\nu = 50$)			high ν model ($\nu = 200$)		
	level	std dev	corr	level	std dev	corr
x -good markup	23%	0.66	-0.16	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	5.9	0.24	0.40	5.7	0.28	0.18
Productivity \bar{z}	low \bar{z} model ($\bar{z} = 5.80$)			high \bar{z} model ($\bar{z} = 8.50$)		
	level	std dev	corr	level	std dev	corr
x -good markup	19%	1.38	-0.08	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	6.3	0.37	0.99	4.9	0.58	-0.48

Table 5: SENSITIVITY ANALYSIS.

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 divided by the standard deviation of GDP, correlations with GDP.

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 8.50 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. See the discussion of the balanced growth correction in section 3 for more details.

D Data and simulation details

Making annual dispersion quarterly. The quarterly persistence and standard deviation of income are derived from the annual estimates of Storesletten, Telmer, and Yaron (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, Telmer, and Yaron (2004) also report consumption dispersion estimates, using food consumption data from the PSID. We use the same procedure as for earnings

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 that we feed into the model, multiply each by $\gamma/(1 + \gamma)$. Thus, $\sigma_H = 0.172Q\gamma/(1 + \gamma) = 0.017$, $\sigma_L = 0.222Q\gamma/(1 + \gamma) = 0.02$, 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, Telmer, and Yaron (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 by the BEA’s chain-type price index with a base year of 2000. We measure aggregate labor productivity by real output per hour and real wages by 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.

State-level data. First, we describe the data sources. The state employment, earnings and wage data come from three sources. The first is the County Wage and Salary Summary (CA34), produced by the BEA’s Regional Economic Accounts (REA). The data are reported annually from 1963-2004. GDP by state, also from the REA, is aggregated based on weights from the Standard Industrial Classification (SIC) codes from 1963-97, and based on the revised North American Industry Classification System (NAICS) from 1997-2005. Due to missing data for Alaska, we use 49 states. The District of Columbia is excluded because computing dispersion is impossible with only one county. We use a second source for state consumer price indices, which we use to construct real wages. The BEA reports state price indices, but only from 1990-2005. Del Negro (1998) estimated the indices for 1969-1995. Both measures are annual. The result is a balanced panel of 49 states and 36 years for a total of 1764 observations.

Next, we construct productivity z and earnings dispersion σ . Productivity is measured by output per worker (labor productivity). It is state GDP, divided by the state CPI to get real state GDP, then divided again by total state employment. Earnings dispersion uses county-level data on the average labor earnings per capita. In each period, state earnings

dispersion is the standard deviation of log earnings, across all the counties in the state. The state real wage is the average nominal wage, divided by the state CPI.

To make our data stationary we remove trends from all variables. While we could remove state-specific deterministic trends, we instead remove a single national trend. This helps preserve as much cross-sectional information as possible. For example, according to our measure a state with persistently lower inequality relative to the national average always has below-trend dispersion; if we had instead removed a state-specific trend then this state would sometime have below trend dispersion and sometime above trend dispersion. We calculate national trends for each variable v as the average across states with each state weighted by its total number of employed workers: $v_{national,t} = \sum_{s=1}^{49} v_{s,t} l_{s,t} / (\sum_{s=1}^{49} l_{s,t})$, where l is the number of employed workers in state s and year t . The de-trending is done by computing the log deviation of the state series from the national average: $\log(v_{s,t}) - \log(v_{national,t})$.

E 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. Then we 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:

$$v(s, \mu, p, r, \pi) = \max_{c, n, x} \{U(c, n, x) + \beta \mathbb{E}[v(s', \mu', p', r', \pi') | s, \mu]\} \quad (41)$$

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 (41), 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:

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

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, we would need to keep track of this variance as 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.

F Computing the decline in business cycle volatility

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, we need a smaller and less rapidly growing dispersion process, like that for consumption dispersion. Storesletten, Telmer, and Yaron (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. (See section 2.7 for further details.) 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.

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

Table 6: BUSINESS CYCLE VOLATILITY AS DISPERSION INCREASES.

Volatility measured as standard deviation of $\log(y)$ in percent. The trend added to the model is the decade-by-decade increase in earnings dispersion estimated by Heathcote, Storesletten, and Violante (2006). ‘Low σ ’ refers to the model calibrated to match the rise in consumption dispersion as estimated by Krueger and Perri (2005). In both cases we use the balanced growth correction described in the text. All other parameter values are listed in table 1. Data: standard deviations of quarterly GDP, by decade.

References

- AIYAGARI, R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109(3), 659–684.
- AYERS, I. (1991): “Fair Driving: Gender and Race Discrimination in Retail Car Negotiations,” *Harvard Law Review*, 104(4), 817–279.
- BAGNOLI, M., AND T. BERGSTROM (2005): “Log-Concave Probability and its Applications,” *Economic Theory*, 26, 445–469.
- BALL, L., AND D. ROMER (1990): “Real Rigidities and the Non-Neutrality of Money,” *Review of Economic Studies*, 57(2), 183–203.
- BARLEVY, G., AND D. TSIDDON (2006): “Earnings Inequality and the Business Cycle,” *European Economic Review*, 50(1), 55–89.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BILBIE, F., F. GHIRONI, AND M. J. MELITZ (2006): “Endogenous Entry, Product Variety, and Business Cycles,” Working Paper.
- BILS, M. (1987): “The Cyclical Behavior of Marginal Cost and Price,” *American Economic Review*, 77(5), 838–855.
- (1989): “Pricing in a Customer Market,” *Quarterly Journal of Economics*, 104(4), 699–718.
- CABALLERO, R., AND M. HAMMOUR (1994): “The Cleansing Effect of Recessions,” *American Economic Review*, 84, 1350–68.
- CARROLL, C. (2000): “Requiem for the Representative Consumer? Aggregate Implications of Microeconomic Consumption Behavior,” *American Economic Review*, 90(2), 110–115.
- CHARI, V., P. KEHOE, AND E. MCGRATTAN (2000): “Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?,” *Econometrica*, 68(5), 1151–1179.
- CHEVALIER, J., A. KASHYAP, AND P. ROSSI (2003): “Why Don’t Prices Rise During Periods of Peak Demand? Evidence from Scanner Data,” *American Economic Review*, 93, 15–37.
- COMIN, D., AND M. GERTLER (2006): “Medium Term Business Cycles,” *American Economic Review*, 96(3), 523–551.
- COOLEY, T., R. MARIMON, AND V. QUADRINI (2004): “Aggregate Consequences of Limited Contract Enforceability,” *Journal of Political Economy*, 112, 817–847.

- COOLEY, T., AND E. PRESCOTT (1995): “Economic Growth and Business Cycles,” in *Frontiers in Business Cycle Research*, ed. by T. Cooley. Princeton University Press.
- DEL NEGRO, M. (1998): “Aggregate Risk Sharing Across US States and Across European Countries,” FRB Atlanta Working Paper.
- DIAZ-GIMENEZ, J., V. QUADRINI, AND J. RIOS-RULL (1997): “Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income and Wealth,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 21(2), 3–21.
- GALI, J. (1994): “Monopolistic Competition, Business Cycles, and the Composition of Aggregate Demand,” *Journal of Economic Theory*, 63, 73–96.
- GOLDBERG, P. (1996): “Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey,” *Journal of Political Economy*, 104(3), 622–654.
- GOMME, P., AND J. GREENWOOD (1995): “On the Cyclical Allocation of Risk,” *Journal of Economic Dynamics and Control*, 19, 19–124.
- GREENWOOD, J., Z. HERCOWITZ, AND G. HUFFMAN (1988): “Investment, Capacity Utilization, and the Real Business Cycle,” *American Economic Review*, 78(3), 402–417.
- GUVENEN, F. (2005): “An Empirical Investigation of Labor Income Processes,” UT-Austin Working Paper.
- HARLESS, D. W., AND G. E. HOFFER (2002): “Do Women Pay More for New Vehicles? Evidence from Transaction Price Data,” *American Economic Review*, 92(1), 270–279.
- HEATHCOTE, J., K. STORESLETTEN, AND G. VIOLANTE (2006): “The Macroeconomic Implications of Rising Wage Inequality in the United States,” NYU Working Paper.
- JAIMOVICH, N. (2006): “Firm Dynamics, Markup Variations, and the Business Cycle,” Stanford University Working Paper.
- KIMBALL, M. (1995): “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money Credit and Banking*, 27, 1241–1277.
- KING, R., AND S. T. REBELO (1999): “Resuscitating Real Business Cycles,” in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford. Elsevier.
- KRUEGER, D., AND F. PERRI (2005): “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *Review of Economic Studies*, 73, 163–193.
- KRUSELL, P., AND A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5).

- LUSTIG, H., AND S. VAN NIEUWERBURGH (2005): “Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective,” *Journal of Finance*, 60(3), 1167–1219.
- (2007): “Why Have Payouts to U.S. Corporations Increased?,” Working Paper.
- MENZIO, G. (2007): “A Search Theory of Rigid Prices,” University of Pennsylvania Working Paper.
- MURPHY, K., A. SHLEIFER, AND R. VISHNY (1989): “Building Blocks of Market-Clearing Models,” *NBER Macroeconomics Annual*, pp. 247–286.
- NAKAMURA, E., AND J. STEINSSON (2006): “Price Setting in Forward-Looking Customer Markets,” Harvard University Working Paper.
- NEVO, A. (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69(2), 307–342.
- PORTIER, F. (1995): “Business Formation and Cyclical Markups in the French Business Cycle,” *Annales D’Economie et de Statistique*, 37-38, 411–440.
- RAMPINI, A. (2004): “Entrepreneurial Activity, Risk, and the Business Cycle,” *Journal of Monetary Economics*, 51, 555–573.
- RIOS-RULL, J. (1996): “Life-Cycle Economies and Aggregate Fluctuations,” *Review of Economic Studies*, 63, 465–490.
- ROSS, S. L. (2003): “What Is Known about Testing for Discrimination: Lessons Learned by Comparing across Different Markets,” University of Connecticut Working Paper.
- ROTEMBERG, J., AND G. SALONER (1986): “A Supergame-Theoretic Model of Price Wars during Booms,” *American Economic Review*, 76(3), 390–407.
- ROTEMBERG, J., AND M. WOODFORD (1999): “The Cyclical Behavior of Prices and Costs,” in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford. Elsevier.
- STOCK, J., AND M. WATSON (1999): “Business Cycle Fluctuations in US Macroeconomic Time Series,” in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford. Elsevier.
- STORESLETTEN, K., C. TELMER, AND A. YARON (2004): “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 112(3), 695–717.