Income Dispersion, Asymmetric Information and Fluctuations in Market Efficiency

Chris Edmond and Laura Veldkamp*
New York University
Stern School of Business

September 22, 2006

Abstract

The idea that information frictions amplify business cycles is hard to evaluate because information is not easily measured. We propose a quantifiable information friction that amplifies output fluctuations. In our simple model of decentralized trade, income dispersion measures uncertainty about buyer characteristics. Counter-cyclical income dispersion makes the asymmetric information friction stronger in recessions. Using income dispersion estimates to quantify the model's effect, we find that this simple friction more than doubles output persistence and increases output volatility 9-fold. The paper compares the model to aggregate data, tests its mechanism using state-level income dispersion, and examines its indirect predictions for markups and luxury goods.

Keywords: business cycles, counter-cyclical markups, asymmetric information.

JEL classifications: D82, E32.

*We thank George Alessandria, Mark Bils, Jeff Campbell, Aubhik Khan, John Leahy, Nick Souleles, Harald Uhlig, Lawrence Uren, Stijn Van Nieuwerburgh, Pierre-Olivier Weill, Michael Woodford and seminar participants at NYU, Philadelphia Fed, Rochester, Iowa, Oslo, Melbourne and the 2006 SED meetings for helpful comments and conversations. Laura Veldkamp also thanks Princeton University for their financial support through the Kenen fellowship.

What magnifies business cycle shocks? Recent work has proposed answers such as varying market competition, capital utilization and labor effort.[1] But a generation ago, researchers argued that information frictions contribute to output fluctuations.[2] Although no one could disprove these theories, modern business cycle research discarded information frictions because they are hard to quantify. Qualitative evidence links cycles to the uncertainty that information frictions create – the cyclical nature of business and consumer confidence, forecast errors, and forecast dispersion are examples.[3] But correlations alone cannot distinguish cause and effect. Is counter-cyclical uncertainty an innocuous by-product of the business cycle, or could it be quantitatively important in amplifying output fluctuations? To answer this question, we need a theory with measurable information frictions and testable implications.

This paper proposes a quantifiable information friction that amplifies output fluctuations. The environment is simple. When a buyer enters a store, she sees a posted price for a good and can purchase the good or not. If the buyer values the good more than the seller does, efficiency dictates that trade should occur. However, the seller does not know each buyer's income, and occasionally overestimates her willingness to pay. If the seller's posted price is too high, an efficient transaction will not take place. This asymmetric information friction is similar to that of Myerson and Satterthwaite (1983). Sellers know the distribution of buyers' incomes at each date, but see each buyer as an independent draw from that distribution. Thus, when income dispersion is high, sellers are more uncertain about each buyer's willingness to pay, and markets are less efficient. Lower income dispersion brings the economy closer to a full-information, efficient market.

Tying information asymmetry to income dispersion is useful because we know how income dispersion varies over the business cycle: it rises in recessions. We take counter-cyclical income dispersion as given, calibrate our model and ask how much inefficiency the mechanism produces. Surprisingly, this simple friction increases output volatility 9.2 times, and doubles output persistence. Our model depicts business cycles as small changes in productivity that affect uncertainty and result in large changes in market efficiency.

To illustrate the workings of the model's key mechanisms, sections 1 and 2 set up and

[1] Jaimovich (2006), Basu (1996), Burnside, Eichenbaum, and Rebelo (1993)

[2] Papers from this period include Grossman and Weiss (1982), Stiglitz and Weiss (1981), and Greenwald, Stiglitz, and Weiss (1984).

[3] For consumer confidence, see Stock and Watson (1999); for business confidence, see Roberts and Simon (2001); for forecast errors and forecast dispersion, see Van Nieuwerburgh and Veldkamp (2006).

analyze a static version of the model. The model's comparative statics describe the effect of changes in productivity and income dispersion. But the sign and magnitude of these effects depends on parameter values. To determine how much effect our information friction generates, section 3 calibrates and simulates a dynamic version of the model. Individuals' income processes are taken from the estimates of Storesletten, Telmer, and Yaron (2004). Aggregate productivity is taken from King and Rebelo (1999). Incomplete information about buyers' incomes has large quantitative effects on the cyclical variation in output.

A recent literature has explored the role of information in business cycles (Beaudry and Portier (2004), Jaimovich and Rebelo (2006), and Lorenzoni (2006)). While that literature focuses on how information about productivity determines production decisions, we focus on how asymmetric information introduces frictions in bilateral market trade.

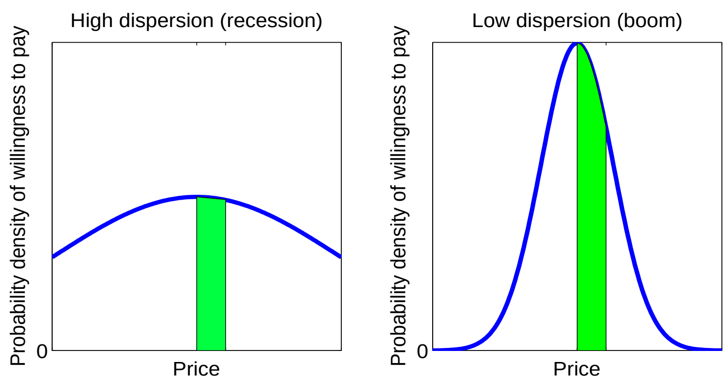

Figure 1: Lowering price is more beneficial when dispersion is low. The shaded area represents the increase in the probability of trade from lowering the price, by an amount equal to the width of the shaded area. This higher probability, times the expected gains from trade, is the marginal benefit to reducing the price. Willingness to pay is based on agents' income.

Our mechanism also resembles mechanisms that lower the price elasticity of demand in recessions, such as Gali (1994)'s changes in the composition of demand. The elasticity effect, which makes prices rise in recessions, comes from the model's counter-cyclical income dispersion. When incomes are more dispersed, buyers' willingness to pay is also more dispersed. If sellers were to reduce prices in recessions, they would attract few additional customers (as 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. In contrast, in booms when dispersion is low, a seller who lowers her price attracts many additional customers (as in the right panel of figure 1). Therefore in booms, sellers keep prices low. In our calibrated model, optimal prices are determined primarily by income dispersion. Since measured dispersion is counter-cyclical and smooth, prices in the model are counter-cyclical and smooth, consistent with GDP deflator data.

To lend support to our mechanism, we empirically verify its time-series and cross-sectional predictions in three ways. Section 5.1 shows that the model produces small counter-cyclical markups, whose fluctuations are similar in magnitude to those measured by Bils (1987). Section 5.2 uses state-level panel data to test the model's predicted relationships between income dispersion and prices. This same dispersion-price relationship is also confirmed in data on car prices across racial groups by Goldberg (1996). Section 5.3 extends the analysis to two goods. One good has the character of a necessity, the other a luxury. The traded quantities and prices of these goods have properties quite different from each other, but similar to luxuries and necessities data. Finally, to asses whether the friction is a realistic one, section 5.4 computes the welfare costs of incomplete information.

Our explanation raises an obvious question: Why does income dispersion rise in recession? One explanation is that job destruction in recessions is responsible (Caballero and Hammour 1994). Rampini (2004) argues that entrepreneurs' incentives change in recessions, making firm outcomes and owners' incomes 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 this model to produce a model whose only driving process is technology shocks.

1 Static model

Our end goal is to build a model and feed an exogenous income process into it, to quantify an information friction. Therefore, our model needs to take income as exogenous, an unconventional feature for a business cycle model. Our mechanism is compatible with a standard RBC model, but our empirical strategy is not. Therefore, we are not going to build a production economy. Instead, this section builds an endowment economy that can take a given amount of asymmetric information about income endowments and calculate its effect on economic fluctuations. To illustrate the working of this mechanism, we start by examining a static model of decentralized trade with one-sided asymmetric information about buyers' income.

Preferences and endowments There is a large but equal number of buyers and sellers. There are two real goods, one which we call c and the other, because it serves as a medium of exchange, we call 'money' m. Buyers and sellers have identical preferences over c and end-of-period money m':

$$U(c) + U(m') = \frac{(c+1)^{1-\alpha} - 1}{1-\alpha} + \frac{(m'+1)^{1-\alpha} - 1}{1-\alpha}, \quad \alpha > 0 \tag{1}$$

The additive form of constant relative risk aversion (CRRA) preferences over consumption and money $U(c) + U(m')$ makes sense if preferences over money are seen as representing preferences over future consumption. What is non-standard about these preferences is the $(+1)$ in each term. This represents a non-tradeable endowment that ensures gains from trade stay finite: acquiring none of a good yields zero utility ($U(0) = 0$), not negative infinite utility. If buyers had none of the sellers' good or sellers had none of the buyers' good, then CRRA preferences would make the marginal gains from trade infinite. Trading frictions only have an interesting role to play when gains from trade are positive, but finite.[4]

Both buyers' and sellers' endowments depend on aggregate productivity z. Buyers, indexed by i, have heterogeneous endowments m_i of money: $\log(m_i) = \log(z) + \varepsilon_i$ where $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$. We will refer to m_i as a buyer's income. Sellers are identical and are endowed with Az units of the consumption good, $A > 0$. Buyers and sellers both want balanced consumption bundles. An increase in z makes no-trade consumption bundles less balanced, increasing gains from trade.

Matching process Every buyer is matched with a seller and every seller is matched with a buyer. This assumption allows us to abstract from matching frictions and focus on the frictions that arise from asymmetric information. Matching is random: the probability of a seller meeting a buyer with a particular income depends only on the aggregate income distribution.

Bargaining and price-setting A seller posts a price p, which is a quantity of money she will accept for δ units of the consumption good. To keep choices simple, the size of trades δ is given. Since sellers have an endowment of Az, we naturally impose that sales must be smaller

[4] In a model with decentralized trade, Aruoba and Wright (2003) make a similar utility assumption to keep their search frictions well-behaved.

than the endowment, $\delta < Az$. The price is a take-it-or-leave-it offer. Given the quoted price p, a buyer decides whether to accept the offer, $a_i = 1$, or not, $a_i = 0$. If trade takes place, δ units of the consumption good changes hands. We use this form of price-setting because it's simple, realistic, and avoids the multiple equilibria problems of bargaining.

Information Everything is common knowledge, except buyers' income, m_i. The distribution of m_i is known, but the income of any particular buyer is not known. If sellers knew each buyer's income, then they would extract all the buyers' surplus, but every efficient trade would be executed. Because of the information asymmetry, some Pareto-improving trades do not occur. This is similar in spirit to the efficiency impossibility result of Myerson and Satterthwaite (1983). This is the source of the market inefficiency.

Equilibrium An equilibrium is an individual buyer acceptance rule $a(p, m_i)$, a probability of trade $\mathbb{E}\{a(p, m_i)\}$, and a price p^* such that:

1. Taking as given the price p, the acceptance rule $a(p, m_i)$ maximizes the utility (1) of buyer i subject to the constraints $m_i' = m_i - pa(p, m_i)$ and $c = a(p, m_i)\delta$.

2. Taking as given the probability of trade $\mathbb{E}\{a(p, m_i)\}$ and the distribution of buyers, sellers choose a price p^* to maximize expected utility (1), subject to the constraints: $m' = pa(p, m_i)$ and $c = Az - a(p, m_i)\delta$.

The buyer's acceptance rule is a simple cutoff strategy. Buyer i maximizes utility by accepting all prices $p \leq \overline{p}(m_i)$ where the cutoff price $\overline{p}(m_i)$ leaves the buyer indifferent between accepting and not accepting the seller's offer: $U(m_i - \overline{p}) + U(\delta) = U(m_i) + U(0)$. Since $U(0) = 0$ the cutoff price is:

$$\overline{p}(m_i) := m_i - U^{-1}\left(U(m_i) - U(\delta)\right). \tag{2}$$

This maximum price a buyer is willing to pay is monotonically increasing in her income. Therefore, for each price p there exists a unique $\underline{m}(p)$ that represents the minimum income that a buyer must have to buy at price p.

Sellers choose price to maximize expected utility, taking buyers' optimal strategy as given:

$$p^* \in \arg\max_{p\geq 0} \left\{ \Pr\left(m_i \geq \underline{m}(p)\right)\left[U(Az-\delta)+U(p)\right] + \Pr\left(m_i < \underline{m}(p)\right)U(Az)\right\}. \tag{3}$$

Maximizing (3) is equivalent to maximizing the probability of trade times the utility gain from a trade, $\Pr(m_i \geq \underline{m}(p))\Delta U(p)$, where the utility gain is $\Delta U(p) := U(p)+U(Az-\delta)-U(Az)$. The probability of trade is:

$$\Pr\left(m_i \geq \underline{m}(p)\right) = \Phi\left(\frac{1}{\sigma}\log\left(\frac{z}{\underline{m}(p)}\right)\right), \tag{4}$$

where Φ denotes the CDF of the standard normal distribution. The utility-maximizing price p^* balances the gains from trade against probability of trade. More specifically, the first order condition tells us that the optimal price equates the price elasticity of utility gain, with the price elasticity of demand:

$$\frac{\partial\log(\Delta U(p))}{\partial\log(p)} = -\frac{\partial\log(\Pr\left(m_i \geq \underline{m}(p)\right))}{\partial\log(p)}. \tag{5}$$

2 Analytic results with log utility ($\alpha \to 1$)

We will explain our model's predictions using log utility, because this makes the mechanism more transparent. Later sections numerically illustrate the same effects with more general CRRA preferences.

Substituting log utility $U(c) = \log(c+1)$ in equation (2) reveals that buyers purchase the consumption good when:

$$p \leq \frac{\delta}{1+\delta}(m_i+1). \tag{6}$$

This decision rule implies that the probability of trade taking place at price p is:

$$\Pr\left(m_i \geq \frac{\delta+1}{\delta}p - 1\right) = \Phi\left(\frac{1}{\sigma}\log\left(\frac{z\delta}{(1+\delta)p-\delta}\right)\right). \tag{7}$$

Effect of dispersion changes If $\sigma = 0$, then there is no uncertainty about the buyer's willingness to pay, all buyers have $m_i = z$. In the absence of asymmetric information, the seller captures all gains from trade, by setting a take-it-or-leave-it price, equal to the buyer's valuation: $p^* = \delta(z+1)/(1+\delta)$. When uncertainty increases, sellers reduce their price to make sure that buyers will still trade with them. At this slightly reduced price, they still earn most of the gains from trade. With a large prize on the line, sellers are loathe to risk it by setting too high a price.

When uncertainty σ is high, however, the logic of a seller changes. The seller still sets the optimal price p^* by balancing the expected benefit of increased profits from raising the price against the expected lost probability of trade. But when the buyers' willingness to pay is very dispersed, increasing the price results in a smaller decrease in the probability of trade. This induces sellers to raise prices when σ becomes high. Figure 2 illustrates this non-monotonic effect of dispersion on price. Proposition 1 states the income dispersion-price relationship formally.

Proposition 1. *For each z, there exists a unique cutoff level of income dispersion $\bar{\sigma}(z) > 0$ such that if $\sigma < \bar{\sigma}(z)$, then prices are decreasing in dispersion, $\partial p^*/\partial\sigma < 0$, and if $\sigma > \bar{\sigma}(z)$, prices are increasing in dispersion, $\partial p^*/\partial\sigma > 0$.*

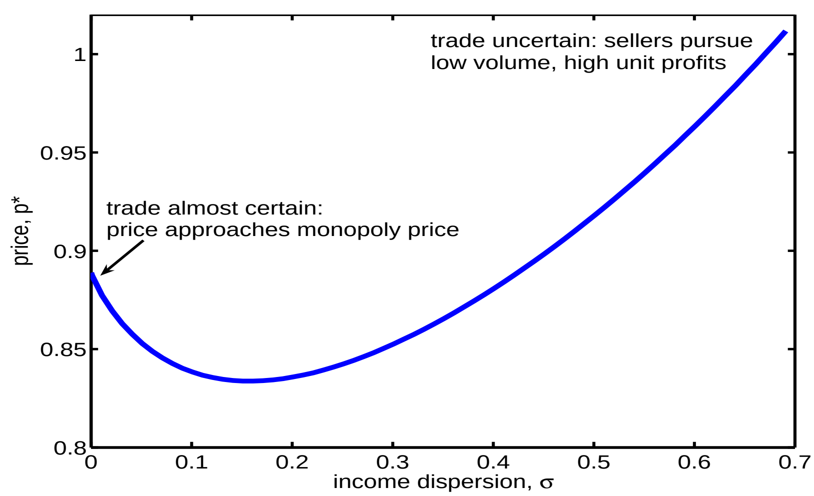

Figure 2: Optimal price p^* as function of dispersion σ.
Price is the seller's posted price for δ units of the good.

This result, proven in appendix A, comes from applying the implicit function theorem to (5). The left side, the elasticity of the seller's utility gain from trade, depends on price p but not on dispersion σ. The non-monotonicity arises because when σ is low, the demand elasticity (right side of (5)) is very price-sensitive. Sellers compensate for a moderate increase in σ by reducing price to ensure that most trade still occurs. When dispersion is high, demand is less price sensitive and the utility gain term dominates pricing decisions. Since an increase in dispersion makes demand less elastic, sellers raise prices.

Effect of productivity changes The effect of productivity changes is less obvious because productivity affects both the aggregate demand elasticity and the price elasticity of unit profits. Prices are non-monotonic in productivity; they can be pro- or counter-cyclical, depending on the mean level of productivity.

To simplify the discussion, we make a parametric assumption which will always be satisfied in our calibrated examples, which rules out a small number of 'perverse' outcomes, and which makes it easier to exposit the mechanics of our model.

Assumption 1. *The size of trades δ is sufficiently large, $\delta \geq \sqrt{1-\delta}$.*

Prices are non-monotonic in productivity. When productivity is very low, sellers do not have many goods, even for their own consumption, so the opportunity cost of selling one is high. Since consumption goods are scarce, their price is high. As productivity z rises, a seller gains more from trading and lowers price to make trade more likely. In this region, price is counter-cyclical. But when z is sufficiently high, the probability of trade and the elasticity of demand are not very sensitive to price. Therefore, sellers focus on capturing gains from trade. To do this, sellers set p^* high and approximately proportional to z, meaning price rises when the gains from trade rise. So, when the gains from trade and average probability of trade are sufficiently high, prices are pro-cyclical. Figure 3 illustrates these effects, which are summarized in the following proposition.

Proposition 2. *For each $\sigma > 0$, there exists a unique cutoff level of productivity $\overline{z}(\sigma) \geq \delta/A$ such that if $z < \overline{z}(\sigma)$, then prices are decreasing in productivity, $\partial p^*/\partial z < 0$, but if productivity is sufficiently high, $z > \overline{z}(\sigma)$, then prices are increasing in productivity $\partial p^*/\partial z > 0$.*

The effect of productivity on prices dictates the cyclical behavior of the probability of trade. Equation (7) shows that changes in productivity affect the probability of trade in two ways. First, an increase in productivity increases each buyer's income m_i and hence directly increases the probability of trade, $\Pr(m_i \geq \underline{m}(p))$. Second, an increase in productivity changes the price p^* and therefore the probability of trade. For low levels of productivity, the

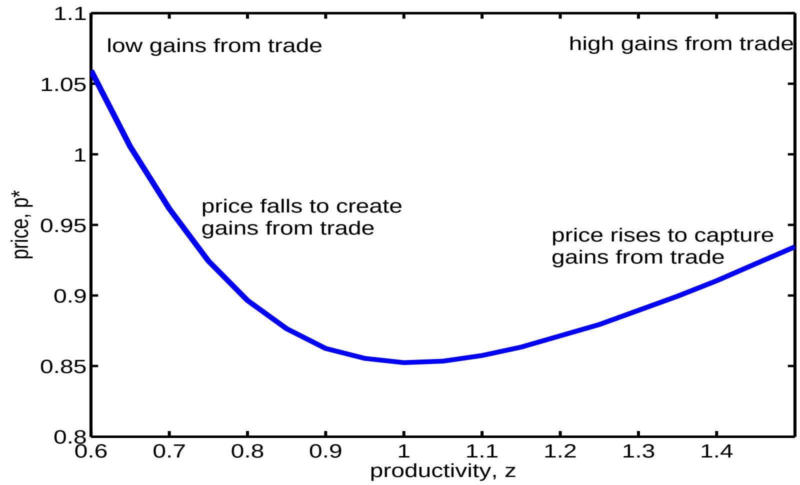

Figure 3: Optimal price p^* as function of productivity z.
Prices fall, then rise, as productivity increases.

price p^* is decreasing. So in low-productivity times, the probability of trade is clearly pro-cyclical. But when productivity is higher, p^* is increasing in productivity; this indirect effect offsets the direct effect of higher average income. In the calibrated examples that follow, the price effect is always dominated by the direct effect of higher average income, making the probability of trade always pro-cyclical.

The static model delivers three insights that explain why the dynamic calibrated model replicates features of the data. (1) The probability of trade is pro-cyclical. This makes exchange less efficient in a recession and amplifies the effect of productivity shocks on traded output. (2) Prices are increasing in productivity z if the gains from trade are high, are insensitive for medium levels of the gains from trade, and are decreasing in productivity z for low levels of gains from trade. (3) More dispersion in buyer incomes causes prices to fall if dispersion is low, but causes prices to rise if dispersion is high. The second and third insights cannot be tested without determining the relevant range of gains from trade and income dispersion. Also, since changes in gains from trade and income dispersion are correlated, their combined effect depends on their relative volatilities. To resolve these questions and compare the model's indirect implications to data, we turn to a dynamic calibrated model.

3 Dynamic model

The dynamic model is simply a repeated static model where productivity z and income dispersion σ fluctuate. Utility is constant relative risk aversion, as in (1). Time is discrete and infinite $t = 0, 1, \ldots$. Productivity follows an AR(1) process:

$$\log(z_t) = (1-\rho)\log(\overline{z}) + \rho \log(z_{t-1}) + \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2). \tag{8}$$

As before, an individual's income is correlated with productivity, but has an idiosyncratic component as well: $\log(m_{it}) = \log(z_t) + \varepsilon_{it}$.

3.1 Calibration

The static model tells us that the cyclical properties of income dispersion are crucial for the behavior of prices and quantities. Storesletten, Telmer, and Yaron (2004) estimate that income dispersion is counter-cyclical. Their income process has idiosyncratic income shocks with persistent and transitory components and the following parameters:

$$\varepsilon_{it} = \xi_{it} + u_{it}, \qquad u_{it} \sim \mathcal{N}(0, 0.065^2) \tag{9}$$

$$\xi_{it} = 0.988\xi_{it-1} + \eta_{it}, \qquad \eta_{it} \sim \mathcal{N}(0, \sigma_{\eta t}^2), \tag{10}$$

and dispersion that increases when productivity is below average

$$\sigma_{\eta t} = \begin{cases} 0.032 & \text{if} \quad z_t \geq \overline{z}, \\ 0.054 & \text{if} \quad z_t < \overline{z}. \end{cases} \tag{11}$$

Appendix B details all our data and its transformations.

The estimates of Storesletten, Telmer, and Yaron (2004) have been controversial, because of the difficulty identifying transitory and permanent shocks. Guvenen (2005) and others argue that, because of unmeasured permanent differences in income profiles, the persistence of income shocks is overestimated. While this distinction is crucial in a consumption-savings problem, it is not relevant for our aggregate model. Whether income 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 and determines the probability of trade. Thus both sides in this debate hold views consistent with our model's predictions.

We calibrate the productivity process z using the estimates of King and Rebelo (1999). The quarterly persistence of the Solow residual is $\rho = 0.98$, and the standard deviation of its innovation is $\sigma_\epsilon = 0.0072$. The coefficient of relative risk aversion is set to $\alpha = 1.5$, a commonly used, conservative value. The size of a trade $\delta = 0.8$ is approximately the same size trade as in the competitive economy in section 3.2.

The most difficult quantities to tie to data are average productivity $\overline{z}$ and the scale factor A. If endowments are too small, most trading pairs have no gains from trade. If they are too large, trading frictions disappear. Since we think of the US economy as a place where people normally produce specialized goods and therefore have moderate gains from trade, we calibrate the endowments to make the unconditional average probability of trade about 0.66. Jointly with other parameters, $A = \bar{z} = 1$ achieves this balance. Model outcomes are sensitive to these parameters. But as long as markets are neither almost efficient nor nearly inoperative, it doesn't matter much what combinations of A and $\bar{z}$ we use to achieve that middle ground. After presenting our main results, we examine the behavior of the economy in a time when gains from trade are very low. We use those results to characterize a depression.

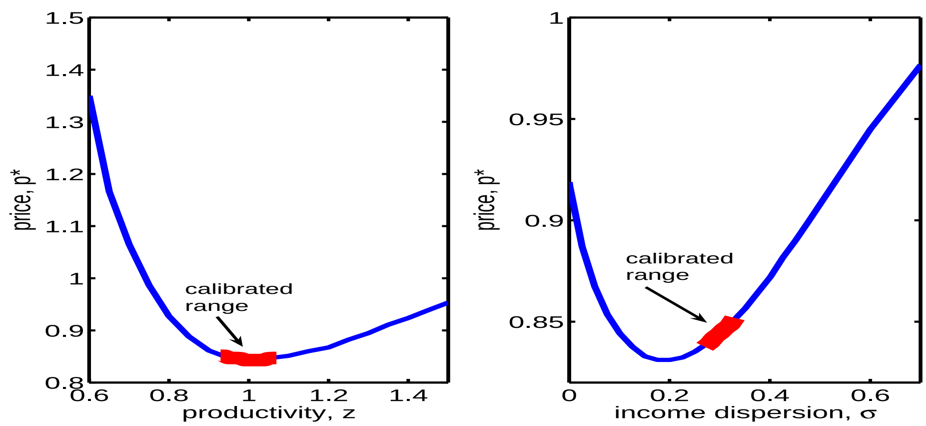

Figure 4: Calibrated range of parameters.
The direct effect of z on p^* is ambiguous, but the effect of σ on p^* is positive. In the left panel, the shaded part of the line is the 2 standard deviation range of productivity, based on estimates from King and Rebelo (1999). In the right panel, the 2 standard deviation range of income dispersion is indicated, based on the estimates of Storesletten, Telmer and Yaron (2004). Preferences are CRRA with $\alpha = 1.5$.

One of the reasons that it is important to calibrate the model is because the cyclical behavior of prices cannot be inferred from theory alone. Only when the relevant region of the parameter space is identified can the model's predictions be compared to data. The left hand panel of figure 4 is the same theory-based mapping between productivity and price as figure 3, but with CRRA instead of log preferences. The calibrated 2 standard deviation interval of productivity around its mean is highlighted. In this interval, prices react very little to changes in productivity but react strongly to changes in dispersion (right panel). Therefore, prices are determined primarily by income dispersion — counter-cyclical income dispersion generates counter-cyclical prices.

With our calibrated parameters, the probability of trade is given by a cumulative distribution function which is steep for moderate levels of productivity. Small changes in productivity are associated with relatively large changes in the probability of trade.

3.2 A competitive benchmark

A natural benchmark for comparisons is an economy with competitive markets. Suppose buyers and sellers face competitive prices. The only difference between this setting and our model is the agents' budget constraints. Now each seller has the budget constraint $pc+m' \leq pAz$ while buyer i has budget constraint $pc+m' \leq m_i$. The first order condition for all agents can be written $[(m'+1)/(c+1)]^\alpha = p$. After substituting p into constraints, each seller demands $c_s = (pAz+1-p^{1/\alpha})/(p+p^{1/\alpha})$ and buyer i demands $c_i = (m_i+1-p^{1/\alpha})/(p+p^{1/\alpha})$ units of consumption. Market clearing requires $c_s + \mathbb{E}\{c_i\} = Az$. This determines the competitive price:

$$p_c = \left(\frac{\mathbb{E}\{m_i\}+2}{Az+2}\right)^\alpha, \tag{12}$$

where $\mathbb{E}\{m_i\} = z\exp(\frac{1}{2}\sigma^2)$.

3.3 Defining GDP

In the competitive benchmark, GDP and output are equivalent because everything produced is sold. In a model with decentralized trade, some output is not sold. Since GDP data only includes final goods traded on a market, GDP in the model also includes only traded goods.

In the model, per-capita real GDP is $y := Az\Pr(m_i \geq \underline{m}(p))$. In a frictionless model the probability of trade would equal one and GDP would be equal to the seller's endowment. That endowment Az is like a potential GDP. With market frictions, some goods are not traded and are consumed by their producers. These goods are home production and are not counted in GDP. Alternatively, it could be that Az is potential production of services. But the services are not actually rendered unless they are sold. Our results are nearly identical if GDP is the average amount of consumption goods that a seller and buyer exchange in the market, $\Pr(m_i \geq \underline{m}(p))\delta$.

4 Main results

In recessions, sellers pursue a low volume, high margin strategy. In booms, they earn lower profit per unit, but make higher total profits by selling more. This section quantifies these effects and compares their magnitudes to evidence from macroeconomic aggregates.

4.1 Amplification of GDP and re-calibration

We start with a standard calibration of a productivity process from King and Rebelo (1999) to compare our model to others that use the same technology shocks. The innovations to that shock process are chosen to let the real business cycle model match the volatility of output. But there is no reason to think that this driving process is appropriate for our model. In fact, it produces too much volatility in output.

Result 1. *Productivity shocks are amplified.*

The relationship between output, productivity and income dispersion is illustrated in figure 5. While a standard real business cycle model's output volatility is 15% higher than productivity volatility, this model's output volatility is 368% higher.

Since we want the model outcomes to resemble the data, we recalibrate the productivity process to match the persistence and volatility of real GDP. The required autocorrelation ρ falls from 0.98 to 0.40. Less persistence of productivity is needed because income dispersion is a highly persistent process. It is the source of long-run swings in output. The standard deviation of innovations to productivity falls by 80% from 0.0072, to 0.0014. Less volatile productivity is required because improvements in productivity are amplified by increases in the probability of trade. The model now produces output volatility 9.2 times higher than productivity volatility. All the qualitative results that follow are true, whether we use the

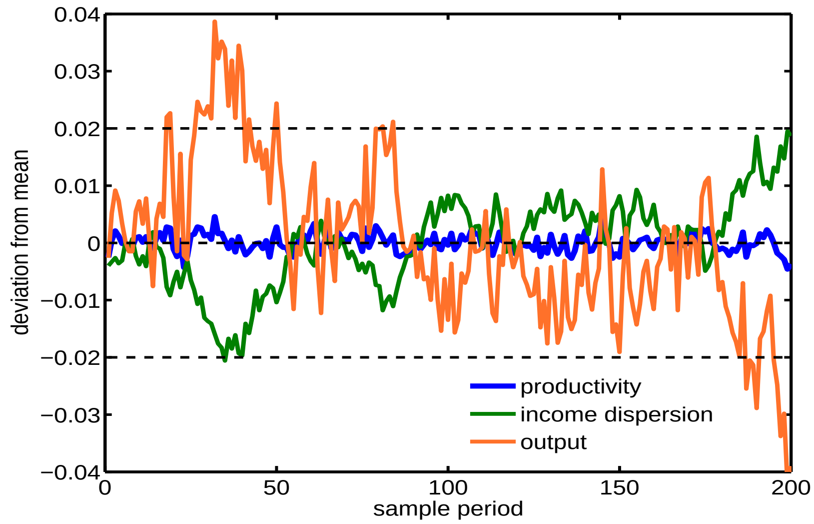

Figure 5: Productivity shocks are amplified (result 1)
The three lines plot output $\log(y)$, productivity $\log(z)$, and income dispersion σ as deviations from their population means. Data come from a simulation of the model with the productivity process re-calibrated to match the properties of output in the data.

original, volatile productivity process, or the new, less-volatile one. We report results based on this new, lower-variance calibration, in order to facilitate comparisons between the model and data.

Result 2. *Prices are counter-cyclical and smooth.*

Figure 6 illustrates the time-series behavior of prices, output and income dispersion. The correlation of log output with log prices in this simulated model is -0.94. The standard deviation of log prices is 14% of the standard deviation of log output. By comparison, simulating the behavior of the competitive price p_c [as defined in equation (12)] shows that prices in the competitive economy are smooth (standard deviation of 77% of the standard deviation of output) and less counter-cyclical (correlation with output is -0.16). GDP is smoother than, and almost perfectly correlated with productivity.[5]

Another relevant benchmark is a model with perfect information. If a buyer's type m_i is perfectly observable, then a seller will set the price $p_i = \overline{p}(m_i)$ to extract all of the surplus [as in equation (2)]. With perfect information about buyers' types, the correlation of log output and log average price becomes 0.41, and the standard deviation of log prices is 1.19 times

[5] In a symmetric version of the competitive benchmark with $A = 1$ and $\sigma = 0$, prices would be constant (independent of z). If $\sigma > 0$, as in our calibration, the competitive model delivers smooth, counter-cyclical prices. But that model does not deliver pro-cyclical market efficiency.

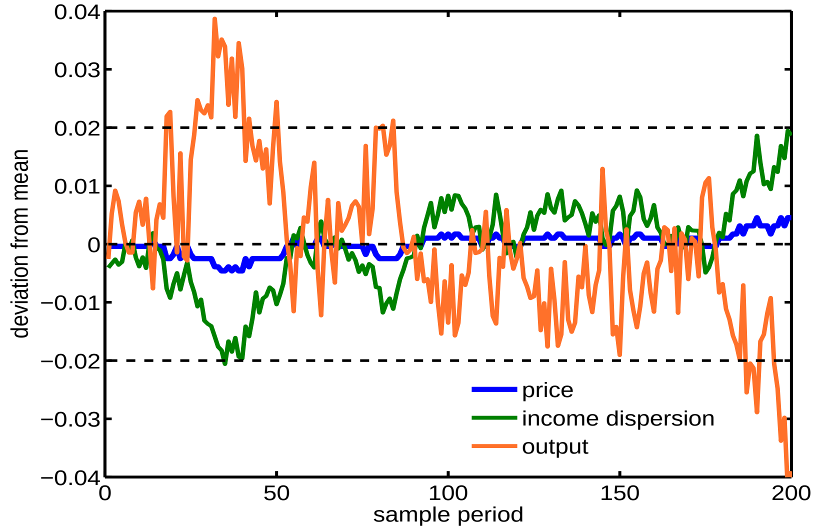

Figure 6: PRICES ARE COUNTER-CYCLICAL AND SMOOTH (RESULT 2).
The three lines plot output, $\log(y)$, income dispersion, σ, and price, $\log(p)$, from a simulation of the calibrated model. All series are shown as deviations from their population means.

the standard deviation of log output. Compared to this perfect information benchmark, the prices in our model are more counter-cyclical and less volatile.

Cyclical fluctuations in price p capture terms of trade variation in an exchange between a buyer and a seller. For theoretical purposes, the nature of the buyers' and sellers' goods does not matter. But to compare model and data, we need to decide what kind of price will be the basis for comparison. We interpret the sellers' good as a general purpose consumption good and interpret the buyers' good as money. In a successful exchange, sellers hand over δ units of the consumption good for p units of money. Therefore, the appropriate counterpart to our p is a price index which measures the money cost of a bundle of goods. The price index we use is the implicit price deflator for total GDP, but our main findings would be unchanged if instead we used an implicit price deflator for aggregate consumption or if we used a fixed-basket consumer price index.

In quarterly macroeconomic aggregates (1947:1-1996:4), the correlation of the log GDP deflator with log GDP is -0.54. The standard deviation of the log GDP deflator is 0.55 times the log standard deviation of GDP (Stock and Watson 1999). While the model's price correlation is too counter-cyclical at -0.94, it comes much closer to matching the data than the strong positive correlation in the perfect information model. The low standard deviation of the model's prices is also closer to the data. But the model makes prices *too smooth.*

The low volatility of prices in our model can be interpreted as a form of real price rigidity.

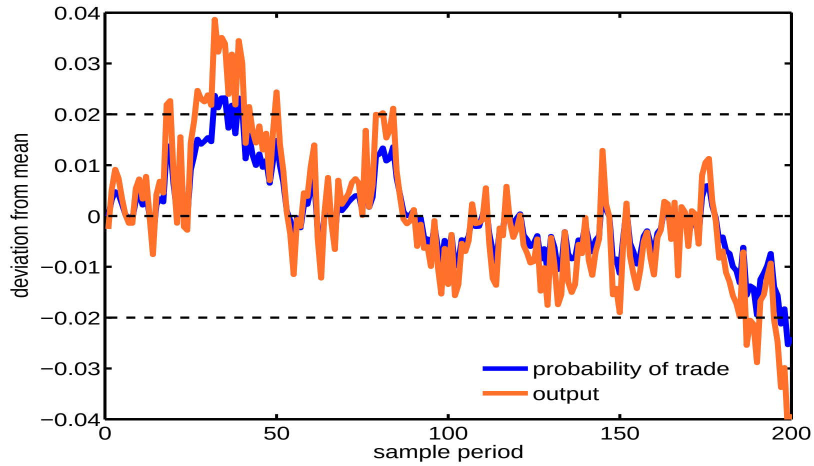

Figure 7: Market efficiency is pro-cyclical and volatile (result 3).
The two lines plot the probability of trade (the measure of market efficiency) and output, $\log(y)$, from a simulation of the calibrated model. Both series are shown as deviations from their population means.

While traditional models of sticky prices have assumed exogenous constraints on price-setting or access to timely information, our real rigidity arises endogenously from observed features of macro data. But our model and sticky-price models share the feature that recessions are times when markets are less efficient because sellers fail to lower prices.

Market efficiency Recessions are times when resources are not put to their most productive uses. Goods take longer to sell (Bils and Kahn 2000), workers are unemployed, and prices do not fall enough to clear markets. Our model produces recessions that are not just declines in productivity; they are also times when markets allocate goods less efficiently. Since almost every transaction is Pareto-improving, the probability of trade $\Pr(m_i \geq \underline{m}(p))$ measures market efficiency.

Result 3. *Market efficiency is pro-cyclical and volatile.*

Figure 7 illustrates the simulated time-series behavior of market efficiency. In contrast, the perfect information benchmark, or any model with a centralized market, predicts that trade always occurs, making its correlation and standard deviation zero. Table 1 summarizes the results from our model.

Reinterpreting results: A model of the real wage Another way to map the model's price into the data is to think of price as the inverse of the real wage, following Rotemberg and

	Prices		Efficiency		Markups	
	Corr with log output	Standard deviation	Corr with log output	Standard deviation	Corr with log output	Standard deviation
Our model	-0.94	0.14	0.99	0.65	-0.82	0.07
Perfect info	0.41	1.19	0.00	0.00	0.88	0.68
Competitive	-0.16	0.77	0.00	0.00	0.00	0.00
Constant σ	0.00	0.00	1.00	0.44	-1.00	0.005
Data	-0.54	0.55				

Table 1: Summary statistics for the model, benchmark models and the data. Standard deviation measures the variables standard deviation, relative to that of log output (real GDP). See section 5.1 for a discussion of markups.

Woodford (1999). Appendix C shows how the model could be modified so that buyers supply labor to sellers, who own the production technology. In exchange for their labor, the buyers receive goods. The rate of exchange of goods for labor, the inverse of the real wage, behaves just like the price in this model. In the data, the quarterly correlation between the inverse of the real wage and real GDP is -0.25 (we measure the real wage as real compensation per hour; Rotemberg and Woodford (1999) give a range of -0.21 to -0.54 for other measures) while the standard deviation of log real wages is 0.60 times the standard deviation of log real GDP. By contrast, the model produces a correlation with log output of 0.91 and a standard deviation of log real wages that is 0.23 times the standard deviation of log output. Our benchmark model gives price that is too smooth and too counter-cyclical relative to the data. And similarly, our model of the real wage gives wages that are too smooth and too pro-cyclical.

4.2 The role of counter-cyclical income dispersion

One thing these results do not tell us is how much of the model's effect come from changes in income dispersion and how much could be generated by changing productivity with a constant amount of information asymmetry. To distinguish these two effects, we reran the simulated model, with one change: income dispersion was constant, equal to the average level of income dispersion in the full model. In the modified model prices become almost constant, with standard deviation and output correlation that are indistinguishable from zero. This reveals that all of the price movement in the model is coming from changes in income dispersion. The reason that prices do not move is illustrated in figure 4. Output falls in a region where prices are flat. Changing parameters such as risk aversion, the mean level of output, or the variance of output shocks would make this result less extreme.

Constant dispersion lowers the volatility of the probability of trade and leaves it almost perfectly pro-cyclical. The reason is that higher productivity increases the gains from trade and the probability of trade. Since productivity is the only force left in the model and it moves both the probability of trade and GDP, the two variables become perfectly correlated. Because of the variation in the probability of trade, this model does still amplify the effect of TFP shocks on output. But, by shutting down the effect of productivity on income dispersion, which in turn affects output, most of the amplification is lost (see table 1). These results suggest that it is not not the details of decentralized trade, but rather counter-cyclical income dispersion that drives our results.

4.3 Economic behavior in a depression

Since the model's relationships between productivity, price and dispersion are nonlinear, the qualitative effect of a big shock can be quite different from a small one. We now illustrate the workings of the model far away from its calibrated parameter values.

We consider a depression-sized shock to productivity. Cole and Ohanian (2004) report that US real gross national product per adult fell 13% below trend in 1930 and reached a trough of 39% below trend in 1933. To model such an event, we simulate many periods and then introduce a productivity shock that reduces GDP 39% below its long-run level. Following this shock, productivity follows the stochastic process given in (8).

During and in the immediate aftermath of a depression, productivity, not income dispersion is the dominant force. The cyclical properties of market efficiency do not change. Although prices are less counter-cyclical – output correlation is -0.71 instead of -0.94 in normal times – markups are more counter-cyclical – their output correlation is -0.91 instead of -0.82. Covariances change because prices are more sensitive to productivity when it is very low (see figure 3). For low enough productivity, this relationship swamps the income dispersion effect. Yet, the basic character of the model survives.

One feature that does change dramatically is output persistence. Output returns back to slightly below its normal level quickly (low persistence), then takes a long time to completely return to trend (high persistence). The low persistence is inherited from productivity, the dominant driving force. After productivity has recovered, income dispersion dominates and persistence rises. The initial recovery is almost complete before income dispersion takes over because the income innovation volatility has only two regimes — this limits its possible effect on output. A third regime with much larger income shocks would slow the recovery sooner and allow for a more protracted depression.

5 Supporting evidence

5.1 Matching counter-cyclical markups

We characterized recessions as times when firms pursue low-volume, high-margin sales strategies. To measure this effect in our model, we define a markup to be the percentage difference between the price in the competitive market economy (12) and the price in our decentralized model. Since the competitive price p_c is a price per unit and our model's price is for δ units, we normalize our price by δ to make it comparable:

$$\text{markup} := \log\left(\frac{p^*}{p_c\delta}\right). \tag{13}$$

Result 4. *Markups are counter-cyclical and smooth*

The correlation of markups and log output in the simulated model is -0.82. The standard deviation of markups is 0.07 times the standard deviation of log output. In the perfect information model, the correlation with log output would be 0.88 while the standard deviation would be about 0.68 times that of output. In a perfectly competitive market, the markup is always zero, by definition.

In the data, counter-cyclical markups have been documented by and 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, and by Portier (1995) with French data. Bils (1987) quantifies this effect by inferring firms' marginal costs. He finds that markups increase an average of 4% in recession. To ask if our model comes close to matching this fact, we compute the average markup of firms for the bottom 14% of output realizations. The 14% corresponds to the fraction of quarters the postwar US economy has spent in an NBER-recession. We compare this recession markup to the average markup for the top 86% of output realizations. We find that markups are 3.5%

higher in recessions.

The 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 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 willingness to pay dispersion σ falls.

5.2 Testing the model with state-level panel data

The model's mechanism relies crucially on an information friction, measured by income dispersion, that induces sellers to raise prices. Using panel data for 49 U.S. states, we find that income dispersion and prices do have a significant positive relationship. Furthermore, the data also confirm the model's predictions for how the correlation between prices and income dispersion varies according to the level of income dispersion and productivity.

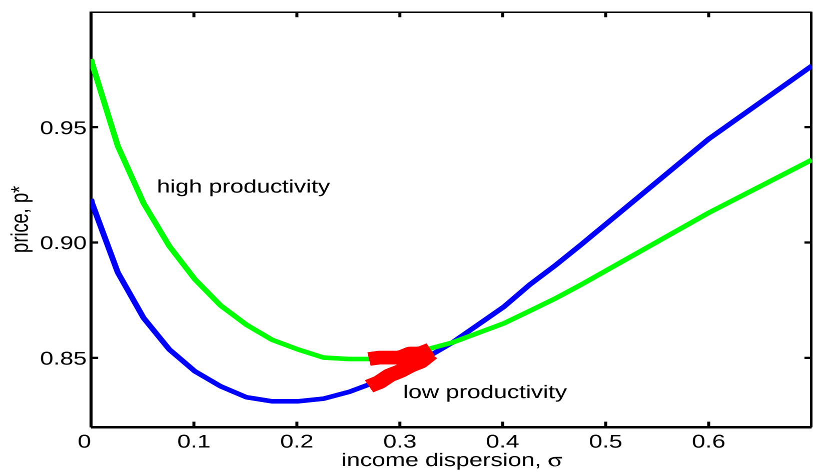

Figure 8: THE JOINT EFFECT OF DISPERSION AND PRODUCTIVITY ON PRICE. The two lines plot the model-predicted price level for varying levels or income dispersion. The line labeled 'low productivity' is for $z = 1$. 'High productivity' is $z = 1.1$. The two-standard deviation bands for income dispersion are shaded.

The model predictions depend on the calibrated range of parameters. Figure 8 plots the relationship between dispersion and prices, for two different productivity levels, and highlights the calibrated range. Four predictions emerge. First, higher dispersion raises price. Second, price is convex in dispersion, so dispersion squared should also be positively related to price. Third, an increase in productivity flattens out the relationship between dispersion and price – an increase in an interaction term should lower price. Finally, higher productivity raises price. To make the effect visible, the figure shows two productivity levels that are 30 standard deviations apart. In practice, the effect of productivity on price is much smaller than the effect of dispersion (see figure 4).

To measure a state's income 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 standard deviation. State price levels come from Del Negro (1998). As a proxy for state productivity, we use the average real salary per job in the state. All data is detrended (further details in appendix B).

Income dispersion (σ)	Squared dispersion (σ^2)	Interaction of productivity & dispersion ($\log(z)\sigma$)	Productivity ($\log(z)$)
0.17** (0.04)			
0.17** (0.04)	1.17* (0.55)		
0.12** (0.04)	1.63** (0.56)	-0.85** (0.22)	
0.16** (0.04)	1.33* (0.56)	-0.75** (0.22)	-0.07** (0.02)

Table 2: Effects of dispersion and productivity on price, in state-level panel data. Standard errors in parentheses. A * denotes coefficients significant at the 5% level, ** at the 1% level. Each regression includes state fixed effects.

Table 2 shows that three of the four predictions are upheld in the data. (1) In states and in years when income dispersion increases, prices rise. (2) This relationship is convex. (3) A rise in productivity dampens the dispersion effect. The fourth prediction, that higher productivity increases prices is contradicted. However, the predicted size of this effect is close to zero, as is the estimated data coefficient. The thrust of the model's effect comes from the relationship between income dispersion and prices. All the predictions relating to income dispersion are upheld by the data, are robust to changes in the specification, and are precisely estimated.

More support for our mechanism comes from a study of the effect income dispersion has on car sales. Goldberg (1996) estimates that blacks' valuations for new cars are more dispersed than whites'. She then collects data on the initial offer to blacks and whites by a car salesman. The initial offer price is higher, and the probability of sale lower, for the group with more dispersed values.

5.3 Cross-sectional variation in prices of luxuries and necessities

We generalizing the model by allowing multiple goods to be traded. This serves two purposes: it demonstrates that the results are not dependent on a single-good monopoly setting, and it delivers additional testable implications for the cross-section of goods. We consider two goods, luxuries L and necessities N. The goods differ only in their relative endowments and the order in which they are traded. We compare the prices and quantities of our two goods to empirical counterparts.

Buyers value both goods equally. Their utility takes the same form as before, with relative risk aversion $\alpha > 0$:

$$\frac{1}{2} \cdot \frac{(c_L + 1)^{1-\alpha} - 1}{1-\alpha} + \frac{1}{2} \cdot \frac{(c_N + 1)^{1-\alpha} - 1}{1-\alpha} + \frac{(m' + 1)^{1-\alpha} - 1}{1-\alpha} \tag{14}$$

Each seller consumes only one of the two goods. To keep their relative preferences for goods and money equal to that of buyers, sellers' utility must be $((c_s + 1)^{1-\alpha} - 1 + (m' + 1)^{1-\alpha} - 1)/(1-\alpha)$ for the seller of good $s = L, N$.

There are two types of sellers. One set of sellers gets a small endowment z of luxuries. The other set gets a large endowment of necessities Az, for a constant $A > 1$. Buyers' endowments of money are as before.

The timing is that the luxury goods producer first sets his price p_L and then the buyer buys or rejects the offer. Then, the necessities producer sets price p_N and again, the buyer accepts or rejects it.[6] Finally, all goods and leftover money are consumed. Equilibrium is defined as before. We restrict attention to parameter values that deliver an equilibrium where $p_N < p_L$. This ensures that we don't compare data luxuries with model necessities.

Solving the model We solve the buyer's problem by backwards induction, starting with the second good purchased, the necessity. The marginal necessities buyer is indifferent between consuming δ necessities with p_N less money, and consuming no goods at all. Let $\underline{m}(p_N)$ be that buyer's income. It is the solution to $\frac{1}{2}U(\delta) + U(\underline{m} - p_N) = U(\underline{m})$, for a given necessity price p_N. Similarly, the marginal luxuries buyer is indifferent between consuming δ units of both goods at cost $p_N + p_L$, and consuming only necessities, at cost p_N. His income

[6] The reason we choose this timing convention is that simultaneous price-posting would produce multiple equilibria. Furthermore, when luxuries sellers move first, they set the higher price, which is realistic.

is $\overline{m}(p_N, p_L)$, which solves $U(\delta) + U(\overline{m} - p_N - p_L) = \frac{1}{2}U(\delta) + U(\overline{m} - p_N)$, given prices p_N and p_L.

Since the two goods are symmetric, except for luxury good's higher price, buyers' cutoff income levels are ordered. A buyer with low income $m_i < \underline{m}$ buys neither good; a buyer with medium income $m_i \in [\underline{m}, \overline{m}]$ buys only necessities because they are cheaper; a buyer with high income $m_i > \overline{m}$ buys both goods.

The probabilities of sales take the same form as in the one-good model: for necessities, $\Pr(m_i > \underline{m}) = \Phi\left[\log(z/\underline{m}(p_N))/\sigma\right]$; and for luxuries, $\Pr(m_i > \overline{m}) = \Phi\left[\log(z/\overline{m}(p_N, p_L))/\sigma\right]$. Sellers' expected utility is a constant plus a probability of trade times a utility gain from trade. The N seller chooses p_N^* to maximize $\Phi\left[\log(z/\underline{m}(p_N))/\sigma\right]\left[U(Az - \delta) + U(p_N) - U(Az)\right]$. Taking the optimal necessity price as given, the L seller chooses price p_L^* to maximize $\Phi\left[\log(z/\overline{m}(p_N, p_L))/\sigma\right]\left[U(z - \delta) + U(p_L) - U(z)\right]$. Finally, we check that luxuries have higher prices: $p_L^* > p_N^*$.

Comparing model and data The data are described in appendix B. The parameter A needs to be recalibrated because it now measures the relative size of the necessity endowment. We choose $A = 1.5$ so that buyers' average expenditure share on luxuries is 4%, as in our data. Table 3 summarizes our results and compares them to the data.

		Traded Quantities		Prices	
		Standard deviation	Corr with output	Standard deviation	Corr with output
Model	Luxuries	3.81	-0.84	0.32	-0.90
	Necessities	0.61	0.97	0.09	-0.93
Data	Luxuries	3.10	0.60	1.56	-0.18
	Necessities	0.65	0.85	0.50	-0.43

Table 3: Summary statistics for the two-good model and the data. Standard deviation measures the variables standard deviation, relative to that of log output (real GDP). Correlations are also with log output. Data is HP-detrended.

The primary objective of the paper is to increase output volatility. Expensive, infrequently purchased goods (luxuries) that are much more volatile than cheap, frequently purchased ones (necessities) fits with the data, providing modest support for the model's mechanism. Luxuries sales are about 5 times more volatile than necessities. In the model, luxuries are volatile because they depend on the mass of buyers in the far right tail of the income distribution. The size of that tail varies greatly over the cycle. It is larger in recessions, when the income distribution fans out. The fact that the upper tail of the income grows in recessions also explains why the model's luxury sales rise in recessions. The model misses badly on the correlation of luxury sales with GDP. The only small success here is that luxury sales are less cyclical than necessities, in the data and model.

The model's prices are also a mixed success. Prices inherit the failings of the one good model but the relative volatility and correlations exhibit some features of the data. The model understates price volatility and overstates its correlation with output; both are faults inherited from the one-good model. But it does produce luxury prices that are less counter-cyclical than necessities. Quantitatively, the relative standard deviations of prices are close: luxury prices are 3.1 times more volatile than necessity prices in the data, and 3.6 times in the model. This effect comes from changing income dispersion. In recessions, when the density of buyers near the middle of the distribution is thinning out, more buyers are moving to the tails of the income distribution. Therefore, the elasticity of demand falls (prices rise) more for necessities bought by the middle of the income distribution than for luxuries bought by the rich.

Comparing the one-good and two-good models Not only do our results survive with multiple goods, but the additional good amplifies output slightly more. The standard deviation of output is 12 times the standard deviation of the technology shocks that drive it. In the one-good model, that ratio was 9. Likewise, counter-cyclical, smooth prices survive. The correlation of prices with GDP was -0.94 in the one-good model and is -0.90 (for luxuries) and -0.93 (for necessities) in the two-good model. For markups, this means that their properties are also unchanged by the introduction of the second good. The correlation of markups with GDP was -0.82 in the one-good model and is -0.94 and -0.87 for luxuries and necessities.

5.4 Are the welfare costs realistic?

If small information frictions cause such large disruptions to macroeconomic aggregates, should agents acquire information to reduce the friction? The answer to this question depends on the cost of asymmetric information for the seller, because only sellers have incomplete information. Sellers will differentiate customer types if the cost of the information friction is higher than the cost of acquiring information.

To compute the cost, we first compute expected utility in our calibrated model. Then, we compute expected utility in a model where buyers' incomes have the same distribution as in our model, but their income is public information. Public information benefits sellers by allowing them to extract the entire surplus from trade without the risk of setting price too high. Full information increases sellers' utility; sellers would sacrifice 35% of their profits to achieve full information. Conversely, buyers suffer when their types are known. On average, buyers would be willing to sacrifice 34% of their consumption to prevent sellers from having perfect information. This is not realistic; no seller could know every buyers' willingness to pay with certainty. Instead, they can gather information or use price discrimination to reduce, but not eliminate, uncertainty. To estimate the benefit of partial uncertainty resolution, we split the individual-specific component of income in two independent pieces: $\varepsilon_i = \sqrt{\gamma}e_i + \sqrt{1-\gamma}\tilde{e}_i$, where $e_i, \tilde{e}_i \sim \mathcal{N}(0, \sigma^2)$ and $\gamma \in [0, 1]$. The seller observes the first component e_i, but does not know the second component $\tilde{e}_i$. The degree of information symmetry is indexed by γ: setting $\gamma = 0$ returns our original model, but as $\gamma \to 1$ all income is observable and information asymmetry disappears.

Most of the sellers' benefit is associated with completely, rather than partially resolving uncertainty. To eliminate 10% of the uncertainty about buyer characteristics ($\gamma = 0.1$), sellers would forfeit 0.7% of their profits. To eliminate 50% of their uncertainty, sellers would forfeit 4.8% of their profits, and 90% elimination of uncertainty is worth 16% of profits, less than half the gains from full information. These more modest utility gains tell us that sellers will acquire information only if it is cheap or very informative.

Although the potential gains to eliminating uncertainty are high for the seller, most of this gain is a distributional effect. The aggregate welfare cost of asymmetric information is 1.6% of consumption. While this is small, it is orders of magnitude larger than the 0.008-0.1% cost estimated by Lucas (1987).

6 Conclusion

This paper presents a simple model that delivers intuitive relationships between business cycle activity and market frictions. Because the key market friction is tied to a measurable variable, the theory offers quantitative predictions for GDP and price fluctuations. A simple friction, like not knowing how much a buyer earns, can amplify business cycle fluctuations 9-fold, and can generate smooth, counter-cyclical prices and markups, of the right orders of magnitude and with realistic cross-sectional properties.

We have analyzed one information friction in a static endowment economy. The next step should be to build this friction into a standard dynamic production economy. New effects may arise where inefficient trade in inputs feeds back to affect output. Inability to sell goods because of strong information frictions would affect the demand for labor and capital. Such additional testable predictions could help to assess the predictive power of the model.

Extending the model with more intertemporal linkages could also could allow it to explain cyclical movements in inventories. The counter-cyclical movements in inventory-sales ratios and the acyclical rate of stockouts has posed a puzzle to economists (Bils 2005). One way that these two facts can be reconciled is if there are more temporary or targeted discounts in recessions. Temporarily reducing prices to attract low-income customers is a likely cause of a stockout. Perhaps recessions are a time when trade is less likely but is also highly volatile.

Can our mechanism be reconciled with the long run decline in business cycle volatility and increase in idiosyncratic volatility (Comin and Philippon 2005)? Because we take income dispersion as given we have nothing to say about why it might rise. But an augmented version of our model could explain why increased income dispersion might cause the decline in macro volatility. If firms can develop better tools for identifying customer characteristics — on-line or mid-week sales, student and senior discounts — they could target prices to a lower-variance distribution of customers. This is more valuable when customers are more heterogeneous. Because market segmentation mitigates asymmetric information frictions, it dampens the effect of cyclical changes in income dispersion. If firms in our model could choose costly market segmentation technologies, then the increase in idiosyncratic volatility would lead more firms to adopt the technologies, which would dampen business cycle fluctuations.

A The pricing function with log utility

With log utility, the gains from trade can be written $\Delta(p,z) := \log[(p+1)(Az-\delta+1)/(Az+1)]$ while the probability of trade can be written $\mu(p,z,\sigma) := \Phi[\log(kz/(p-k))/\sigma]$ where $k := \delta/(1+\delta)$. With this notation, the first order condition for a maximum can be written:

$$\frac{\partial\Delta(p,z)}{\partial p}\frac{p}{\Delta(p,z)} = -\frac{\partial\mu(p,z,\sigma)}{\partial p}\frac{p}{\mu(p,z,\sigma)}. \tag{15}$$

This implicitly determines a pricing function $p(z,\sigma)$.

Proof of Proposition 1

Applying the implicit function theorem to (15) and using the second order condition for a maximum shows that a necessary and sufficient condition for $\partial p(z,\sigma)/\partial\sigma > 0$ is:

$$\frac{\partial^2\mu(p,z,\sigma)}{\partial\sigma\partial p}\mu(p,z,\sigma) > \frac{\partial\mu(p,z,\sigma)}{\partial\sigma}\frac{\partial\mu(p,z,\sigma)}{\partial p}. \tag{16}$$

Computing the relevant derivatives and simplifying shows that this is true if and only if:

$$\left(\frac{1}{\sigma}\log\left(\frac{p-k}{kz}\right)\right)^2 - \frac{1}{\sigma}\log\left(\frac{p-k}{kz}\right)H\left(\frac{1}{\sigma}\log\left(\frac{p-k}{kz}\right)\right) < 1, \tag{17}$$

where $H(w) := \phi(w)/(1-\Phi(w)) > 0$ denotes the standard normal hazard function for $w \in \mathbb{R}$. This satisfies $H(w) > w$, $H'(w) > 0$, $H(-\infty) = 0$ and $H(w)/w \to 1$ as $w \to \infty$. Let $\varphi(w) := w[w - H(w)]$. With a simple change of variables (17) can be written $\varphi(w) < 1$. Due to the properties of a standard normal hazard function, it is straightforward to show that $\varphi(w)$ is strictly decreasing with $\varphi(0) = 0$ and a unique $\omega := -\varphi^{-1}(1) > 0$ independent of z and σ such that $\varphi(-\omega) = 1$. We can therefore say that (17) is satisfied if and only if:

$$p(z,\sigma) > k + kz\exp(-\omega\sigma). \tag{18}$$

Now let the right side of (18) be $F(\sigma) := k + kz\exp(-\omega\sigma)$ which is strictly decreasing in σ and also satisfies $F(0) = k + kz = p(z,0)$. Therefore, as $\sigma \to 0$, we have $p(z,0) \to F(0)$ and so the necessary and sufficient condition (18) gives $\lim_{\sigma\to 0}\{\partial p(z,\sigma)/\partial\sigma\} \leq 0$. Therefore, for each z there exists a unique $\overline{\sigma}(z) > 0$ such that $p(z,\sigma) < F(\sigma)$ for all $\sigma < \overline{\sigma}(z)$, $p(z,\overline{\sigma}(z)) = F(\overline{\sigma}(z))$ and $p(z,\sigma) > F(\sigma)$ for all $\sigma > \overline{\sigma}(z)$. □

Proof of Proposition 2

Applying the implicit function theorem to (15) and using the second order condition for a maximum shows that a necessary and sufficient condition for $\partial p(z,\sigma)/\partial z > 0$ is:

$$\frac{\partial^2\mu(p,z,\sigma)}{\partial z\partial p}\frac{z}{\frac{\partial\mu(p,z,\sigma)}{\partial p}} > \frac{\partial\mu(p,z,\sigma)}{\partial z}\frac{z}{\mu(p,z,\sigma)} - \frac{\partial\Delta(z,p)}{\partial z}\frac{z}{\Delta(z,p)}. \tag{19}$$

Computing the relevant derivatives and simplifying shows that this is true if and only if:

$$\frac{1}{\sigma}\log\left(\frac{p-k}{kz}\right) < H\left(\frac{1}{\sigma}\log\left(\frac{p-k}{kz}\right)\right)\left[1 - \frac{p+1}{p-k}G(z)\right], \tag{20}$$

where $H(w) := \phi(w)/(1-\Phi(w)) > 0$ again denotes the standard normal hazard function for $w \in \mathbb{R}$ and where:

$$G(z) := \frac{\delta A z}{(Az-\delta+1)(Az+1)} \in (0,1). \tag{21}$$

$G(z)$ is maximized at $z^* = \sqrt{1-\delta}/A$ with $G'(z) < 0$ for all $z > z^*$. Assumption 1 then implies $z > z^*$ and hence $G'(z) < 0$. Now define $\psi(w) := w/H(w)$, which has the properties $\psi'(w) > 0$, $\psi(-\infty) = -\infty$, $\psi(0) = 0$, and $\psi(\infty) = 1$. We can write the necessary and sufficient condition as:

$$\psi\left(\frac{1}{\sigma}\log\left(\frac{p-k}{kz}\right)\right) < 1 - \frac{p+1}{p-k}G(z). \tag{22}$$

Notice that the left hand side of (22) is strictly increasing in p taking on the value $-\infty$ at $p = k$ and approaching 1 from below as $p \to \infty$. Similarly, the right hand side is strictly increasing in p also taking on the value $-\infty$ at $p = k$ but approaching $1 - G(z) < 1$ from below as $p \to \infty$. By the intermediate value theorem there exists a unique $\overline{p}(z,\sigma)$ such that the left hand side is less than the right hand side if and only if $p(z,\sigma) < \overline{p}(z,\sigma)$ (we let $\overline{p}(z,\sigma) = 0$ in the non-generic case where the left hand side is always greater than the right). Hence we can write the necessary and sufficient condition as:

$$\frac{\partial p(z,\sigma)}{\partial z} > 0 \Leftrightarrow p(z,\sigma) < \overline{p}(z,\sigma). \tag{23}$$

Another application of the implicit function theorem shows that the cutoff function $\overline{p}(z,\sigma)$ is increasing in z, $\partial\overline{p}(z,\sigma)/\partial z > 0$ with boundaries $\overline{p}(\delta/A,\sigma) \geq k$ and $\overline{p}(\infty,\sigma) = \infty$. Hence for each $\sigma > 0$ there exists a $\overline{z}(\sigma) \geq \delta/A > 0$ such that $\partial p(z,\sigma)/\partial z > 0$ if and only if $z > \overline{z}(\sigma)$. □

B Calibration details

Income dispersion The quarterly persistence and standard deviation of income are derived from the annual estimates of Storesletten, Telmer, and Yaron (2004) as follows: $\varphi = 0.952^{1/4}$, 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 $Q := 1/(1+\varphi+\varphi^2+\varphi^3) = 0.2546$. We deviate from Storesletten, Telmer, and Yaron (2004) in one additional way: Our income shock changes variance, depending on whether productivity is above its mean or not. The variance of their income shock depends on whether output is above its mean. Because of the high correlation of productivity and output, these two processes are virtually indistinguishable. We make this change because conditioning the variance on output, an endogenous variable, is significantly more complicated.

Productivity The productivity statistics come from time series that are HP filtered with a smoothing parameter of 1600.

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 10000 runs of the simulation.

Two-good model data Income dispersion is the standard deviation of log income from the U.S. Census. The data is annual from 1966-2001 and is detrended using an HP filter.

The price and quantity data come from national income and product accounts. They are quarterly, seasonally-adjusted PCE personal consumption expenditures, by types of product (1959:1-2006:2). Following Poterba (2000), we use jewelry and watches and personal boats and planes to represent luxury expenditures. To ensure that our results are not specific to luxury durables, we include two services as well: watch and jewelry repair and foreign travel. Just like durables are more volatile and cyclical than services, so are their luxury counterparts. But both luxury categories are more volatile than their non-luxury counterparts. Volatilities of the luxury goods vary between 2-7 times as volatile as GDP. Their correlations with GDP range between 27% and 54%. Together, the four goods categories represent 4% of personal expenditures. The remaining categories represent necessities. While this leaves necessities contaminated, their properties are simply those of aggregate expenditures. This is true in the data and in the model.

State-level income dispersion The state data come from the County Wage and Salary Summary (CA34), produced by the Bureau of Economic Analysis' regional economic accounts. The data are reported annually from 1969-2004. Due to missing data for Arkansas, we use 49 states. The District of Columbia is excluded because computing dispersion is impossible with only one county.

State prices are the state consumer price indices computed by Del Negro (1998). They are annual from 1969-1995. We remove the national trends in the CPI, in salary levels and in salary dispersion. The national trend is the average across states, with each state weighted by its total number of jobs.

C An alternative model with labor

In this appendix we develop an alternative formulation of our model that shows how our mechanism generates procyclical but smooth fluctuations in real wages.

There is a large but equal number of workers and producers. They have identical preferences over consumption c and leisure ℓ given by $U(c)+U(\ell)$ where $U(x) = [(x+1)^{1-\alpha}-1]/(1-\alpha)$ for $x = c, \ell$. Producers are identical and have a technology for turning labor time into consumption goods: one unit of time produces Az units of goods. Workers have heterogeneous endowments of labor time. Worker i has time endowment z_i where $\log(z_i) = \log(z)+\varepsilon_i$ and $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$. Every worker is matched with a producer. Since z_i is not observable to producers ex ante, employment contracts merely specify a given shift-length. Producers post a wage that is a take-it-or-leave-it offer of w goods in return for δ units of time.

We could alternatively model this situation by assuming that workers have heterogeneous *effective* time endowments, but this alternative formulation ends up obscuring the formal parallel with our benchmark model. In particular, in our setting employed workers have residual leisure $\ell_i = z_i - \delta$ that depends additively on δ. By contrast, if workers have

z_i effective time but producers employ δ units of real time, then employed workers have residual leisure $\ell_i = z_i(1-\delta)$ which depends multiplicatively on δ.

Worker i accepts a wage of w if and only if:

$$U(w) + U(z_i - \delta) \geq U(0) + U(z_i) \tag{24}$$

This implicitly defines a reservation wage $\underline{w}(z_i)$ that represents the lowest wage a worker with time z_i will accept.

If a wage offer is accepted, the producer has utility $U(Az\delta - w) + U(1)$. If the wage is not accepted, the producer has $U(0) + U(1)$. (We assume that producers inelastically supply one unit of labor). Producers do not know the type z_i of the worker they are matched with, but they do know the distribution of z_i they are sampling from. Production takes place if and only if a wage $w \geq \underline{w}(z_i)$ is offered. Therefore, producers choose a wage w^* that solves:

$$w^* \in \arg\max_{w \geq 0} \left\{ \Pr\left(w \geq \underline{w}(z_i)\right) U(Az\delta - w) \right\}. \tag{25}$$

We solve this problem numerically and calibrate it as in section 3. The only difference is that we now need to set $A = 1.37$ to achieve an average probability of trade of about 0.65. As shown in figure 9, the wage-setting model looks like the mirror-image of our benchmark price-setting model (figure 4). For our calibration, the sensitivity of the wage to aggregate productivity is ambiguous but a rise in income dispersion (as in a recession) unambiguously reduces the wage while a fall in income dispersion (as in a boom) increases the wage.

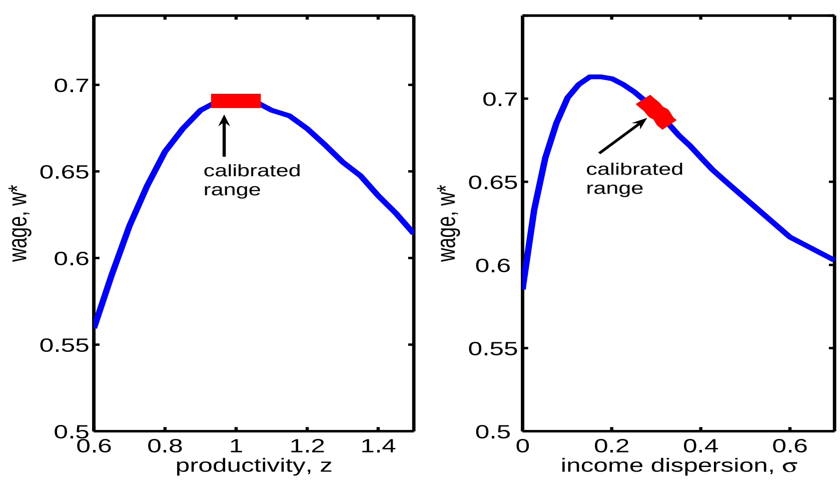

Figure 9: Calibrated parameter range in the labor model.
The direct effect of z on w^* is ambiguous, but the effect of σ on w^* is negative

With the small productivity shocks described in section 4.1, the models correlation of log real wages and log output is 0.91: Real wages are procyclical. We also find that the standard deviation of log real wages relative to log output is only 0.23: Wages are smooth. The other implications of this wage-setting model are essentially the same as in our benchmark. For example, the probability of trade is procyclical and about 0.62 times as volatile as log output. The same productivity shocks that let the price-setting model match the empirical volatility and persistence of output enable the wage- setting model to match those moments as well. The model standard deviation of log output is about 0.013 (against 0.016 in the data) with a first order autocorrelation of 0.83 (against 0.80 in the data). In short, the wage-setting model delivers essentially the same empirical implications as our benchmark price-setting model.

References

Aruoba, B., and R. Wright (2003): "Search, Money and Capital : A Neoclassical Dichotomy," *Journal of Money Credit and Banking*, 35(6), 1086–1105.

Basu, S. (1996): "Procyclical Productivity, Increasing Returns or Capital Utilization?," *Quarterly Journal of Economics*, 111, 719–51.

Beaudry, P., and F. Portier (2004): "An Exploration into Pigou's Theory of Cycles," *Journal of Monetary Economics*, 51(6), 1183–1216.

Bils, M. (1987): "The Cyclical Behavior of Marginal Cost and Price," *American Economic Review*, 77(5), 838–855.

——— (2005): "Deducing Markup Cyclicality from Stockout Behavior," Working Paper.

Bils, M., and J. Kahn (2000): "What Inventory Behavior Tells Us About Business Cycles," *American Economic Review*, 90(3), 458–481.

Burnside, C., M. Eichenbaum, and S. Rebelo (1993): "Labor Hoarding and the Business Cycle," *Journal of Political Economy*, 101, 245–273.

Caballero, R., and M. Hammour (1994): "The Cleansing Effect of Recessions," *American Economic Review*, 84, 1350–68.

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.

Cole, H., and L. Ohanian (2004): "New Deal Policies and the Persistence of the Great Depression: A General Equilibrium Analysis," *Journal of Political Economy*, 112(4), 779–816.

Comin, D., and T. Philippon (2005): "The Rise in Firm-Level Volatility: Causes and Consequences," in *NBER Macroeconomics Annual*.

Cooley, T., R. Marimon, and V. Quadrini (2004): "Aggregate Consequences of Limited Contract Enforceability," *Journal of Political Economy*, 112, 817–847.

Del Negro, M. (1998): "Aggregate Risk Sharing Across US States and Across European Countries," FRB Atlanta Working Paper.

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.

Greenwald, B., J. Stiglitz, and A. Weiss (1984): "Information Frictions in the Capital Market and Macroeconomic Fluctuations," *American Economic Review*, 74(2), 194–199.

Grossman, S., and L. Weiss (1982): "Heterogeneous Information and the Theory of the Business Cycle," *Journal of Political Economy*, 90(4), 699–727.

Guvenen, F. (2005): "An Empirical Investigation of Labor Income Processes," U. Texas Working Paper.

Jaimovich, N. (2006): "Firm Dynamics, Markup Variations, and the Business Cycle," Stanford University Working Paper.

Jaimovich, N., and S. Rebelo (2006): "Can News About the Future Drive the Business Cycle?," Working Paper.

King, R., and S. T. Rebelo (1999): "Resuscitating Real Business Cycles," in *Handbook of Marcoeconomics*, ed. by J. Taylor, and M. Woodford. Elsevier.

Lorenzoni, G. (2006): "A Theory of Demand Shocks," Working Paper.

Lucas, R. (1987): *Models of Business Cycles*. Oxford: Basil Blackwell.

Lustig, H., and S. Van Nieuwerburgh (2005): "Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective," *Journal of Finance*, 60(3), 1167–1219.

Murphy, K., A. Shleifer, and R. Vishny (1989): "Building Blocks of Market-Clearing Models," *NBER Macroeconomics Annual*, pp. 247–286.

Myerson, R., and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 28, 265–281.

Portier, F. (1995): "Business Formation and Cyclical Markups in the French Business Cycle," *Annales D'Economie et de Statistique*, 37-38, 411–440.

Poterba, J. (2000): "Stock Market Wealth and Consumption," *Journal of Economic Perspectives*, 14(2), 99–118.

Rampini, A. (2004): "Entrepreneurial Activity, Risk, and the Business Cycle," *Journal of Monetary Economics*, 51, 555–573.

Roberts, I., and J. Simon (2001): "What do Sentiment Surveys Measure?," Working Paper.

Rotemberg, J., and M. Woodford (1999): "The Cyclical Behavior of Prices and Costs," in *Handbook of Marcoeconomics*, ed. by J. Taylor, and M. Woodford. Elsevier.

Stiglitz, J., and A. Weiss (1981): "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 71(3).

Stock, J., and M. Watson (1999): "Business Cycle Fluctuations in US Macroeconomic Time Series," in *Handbook of Marcoeconomics*, 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.

Van Nieuwerburgh, S., and L. Veldkamp (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53(4), 753–772.