Credit shocks and cycles: a Bayesian calibration approach

Roland Meeks*

May 8, 2006

Abstract

This paper asks how well a general equilibrium agency cost model describes the dynamic relationship between credit variables and the business cycle. A Bayesian VAR is used to obtain probability intervals for empirical correlations. The agency cost model is found to predict the leading, countercyclical correlation of spreads with output when shocks arising from the credit market contribute to output fluctuations. The contribution of technology shocks is held at conventional RBC levels. Sensitivity analysis shows that moderate prior calibration uncertainty leads to significant dispersion in predicted correlations. Most predictive uncertainty arises from a single parameter.

JEL classification: C11, C32, E32, E37, E44

*Nuffield College, University of Oxford. Email: roland.meeks@nuffield.ox.ac.uk. Financial support from the Economic and Social Research Council and the British Academy Post Doctoral Fellowship is gratefully acknowledged.

1 Introduction

The financial health of a corporation can be gauged in part by its credit spread, namely the cost of its borrowing in excess of a safe alternative. Credit spreads are understood to be caused by agency costs that arise when borrowers and lenders have asymmetric information. These costs may vary in intensity with the state of the business cycle, linking macroeconomic and corporate risk. This link was formalised in the dynamic general equilibrium agency cost model of Bernanke and Gertler (1989) and Carlstrom and Fuerst (1997), and is the main focus of the current paper.

A stylised account of the mechanism working between macroeconomic risk and credit spreads can be described as follows. Suppose a technology shock increases the demand for capital goods. The entrepreneurs who supply new capital would like to raise funds from households to finance increased production. Given information asymmetries, the financial contract takes the form of risky debt, a standard feature of which is fixed repayment to be made when solvent. When the entrepreneur is bankrupt, the bond holder receives a fraction of the promised payment. The possibility of leaving lenders with the bill when bankrupt leads entrepreneurs to desire excessive leverage, and intermediaries to demand a compensatory loan premium. Macroeconomic fluctuations in response to a shock are propagated by this mechanism, and the dynamics of real aggregates are made more consistent with data Carlstrom and Fuerst (1997).

Empirical studies show that macroeconomic risk is a significant factor for spreads in credit markets, see for example Gertler and Lown (1999) and Koopman and Lucas (2005). The consensus finding is that credit spreads are negatively correlated with deviations from trend output, and lead at the one year horizon, see Kwark (2002). The present paper motivates these empirical observations using the theoretical macroeconomic model of Carlstrom and Fuerst (1997). However, it is well known that their setup has the implication that credit spreads widen following a positive technology shock, leading to a procyclical spread. For example, Gomes, Yaron, and Zhang (2003) find that the model is superior to the standard adjustment cost approach in matching a number of asset pricing facts, but highlight its counterfactual implications for the cyclical movement of spreads. In a related contribution, Kwark (2002) introduces costs to adjusting investment decisions made in advance of a technology shock in order to match the pattern of observed comovements.

The contributions of this paper are as follows. First, I do not condition on technology shocks in deriving theoretical predictions for unconditional moments of data. Using a plausible set of macroeconomic risk factors, the agency cost model predicts the observed negative leading behaviour of credit spreads over the business cycle. As in Li and Sarte (2003) and Cooper and Ejarque (2000), shocks may arise directly from the financial sector. These shocks are shown to be of central importance in matching data. However, technology shocks remain the dominant source of output fluctuations, consistent with the real business cycle hypothesis, indicating that previous rejections of the agency cost model may have been premature.

Second, the consequences of global parameter uncertainty is formalised using reasonable priors for calibrated parameters as in Canova (1995), Geweke (1999a) and DeJong, Ingram, and Whiteman (1996). A significant degree of uncertainty around the point predictions is found under the baseline specification. Local sensitivity analysis identifies the persistence of credit shocks as the principal source of this uncertainty.

Third, credit and output data are modelled in a Bayesian framework, and probability distributions for both the theoretical model and the data model are calculated. Estimated Bayesian probability intervals are consistent with previous findings of negative, leading credit spreads. However, point estimates are subject to significant uncertainty in the case of the preferred measure of agency costs. In this paper, I take the view that the theory is too stylised to directly describe observed data, making estimation proceedures infeasible. Following DeJong et al. (1996) and Geweke (1999a), the model is intended to match specified moments of the data only. Success in this framework amounts to high overlap between the predictive and estimated densities.

2 The model economy

This paper builds multiple sources of macroeconomic risk known to be useful for matching theory to the data into the economy of Carlstrom and Fuerst (1997). There are four actors: households, firms, intermediaries and entrepreneurs. Firms are fully owned by households, and intermediaries are passive vehicles for the perfect diversification of idiosyncratic within-period risk. Household/firms operate in competitive markets to produce a final consumption good using capital and labour as inputs. Entrepreneurs are a separate class, with a different preference structure and access to a different technology, which transforms final goods into capital goods. There is an informational asymmetry between individual entrepreneurs and household/intermediaries which creates incentive problems that are remedied by the parties writing contracts that resemble standard debt. For further discussion, see Carlstrom and Fuerst (1997).

2.1 Households and firms

Households own capital which they rent to firms, earning rental payments r, and supply labour, which earns a competitive wage w. Household preferences over consumption c and fraction of hours worked h are given by a CES utility function with discount factor β:

$$U(c,h) = \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t a_t \frac{\left(c_t^{\chi}[1-h_t]^{1-\chi}\right)^{1-\sigma} - 1}{1-\sigma} \tag{1}$$

where the preference shock a will be assumed to follow an autoregressive process. The coefficient σ is the reciprocal of the intertemporal elasticity of substitution, and χ controls the substitutability of consumption and leisure in utility. Firms' technology is given by:

$$y_t = \theta_t k_t^{\alpha} h_t^{1-\alpha} \tag{2}$$

where k is physical capital and θ is an autoregressive shock to total factor productivity. We take the overall labour input to comprise that from households and a quantitatively trivial contribution from entrepreneurs h^e with output elasticity α^e. This will be needed by entrepreneurs who have zero net worth, in order to participate in the financial market. The cost of a unit of capital is q units of final goods. Because households are able to diversify away idiosyncratic risk via intermediaries, capital goods are received with certainty. It follows that the household's capital purchases must satisfy the Euler equation:

$$q_t = \mathrm{E}_t \beta \left(\frac{a_{t+1}}{a_t}\right) \left(\frac{c_{t+1}}{c_t}\right)^{\chi(1-\sigma)-1} \left(\frac{1-h_{t+1}}{1-h_t}\right)^{(1-\chi)(1-\sigma)} \{(1-\delta)q_{t+1} + r_{t+1}\} \tag{3}$$

where r is the rate of return on capital, and δ is the rate of depreciation. The final term on the right, in braces, is the total return (rental plus capital gain) to holding a unit of capital from this period to the next.

The labour supply decision is governed by the requirement that the returns to a marginal hour devoted to market activity are equal to the return to devoting that same hour to leisure:

$$\frac{1-\chi}{\chi}\frac{c_t}{1-h_t} = (1-\alpha)\theta_t k_t^{\alpha} h_t^{-\alpha}. \tag{4}$$

Notice that the preference shock acts symmetrically on the marginal utility of consumption and leisure, a positive shock acting to increase the demand for both, without altering the intra-temporal margin between them.

2.2 Entrepreneurs and intermediaries

Intermediaries take deposits from households and make intra-period loans to the mass of entrepreneurs. Each entrepreneur operates a production technology that carries some idiosyncratic risk, so that for every i units of consumption goods that are invested, ωi units of capital goods are produced, where the distribution function of ω is denoted $\phi(\omega)$, and $\mathrm{E}_\omega \omega = 1$. The financial contract is obtained under the assumption of risk neutrality on the part of the contracting parties, as the loan is within-period and there is no aggregate risk over its life. Expected entrepreneurial revenue is maximised subject to an expected break-even condition for the intermediaries. Entrepreneurs capture the surplus, with intermediaries as residual claimants. Formally:

$$\max q_t i_t f(\varpi_t) \quad \text{subject to} \quad q_t i_t g_t(\varpi_t) - (i_t - n_t) \geq 0 \tag{5}$$

where i is total investment, n is net worth and:

$$f(\varpi_t) = \int_{\varpi_t}^{\infty} \omega\phi(\omega)\mathrm{d}\omega - [1-\Phi(\varpi_t)]\varpi_t \tag{6}$$

and:

$$g_t(\varpi_t) = \int_0^{\varpi_t} \omega\phi(\omega)\mathrm{d}\omega - \Phi(\varpi_t)\mu_t + [1-\Phi(\varpi_t)]\varpi_t \tag{7}$$

are the expected shares of the project revenue going to the entrepreneur and intermediary respectively. Here ϖ is the breakeven level of ω, the smallest value of the idiosyncratic shock consistent with an entrepreneur being able to repay his loan. The integral in (6) is the expectation of the idiosyncratic shock, conditional upon the entrepreneur being solvent, multiplied by the probability of being solvent. Likewise, the integral in (7) is the expectation of the shock conditional upon the entrepreneur being bankrupt, times the probability that he is bankrupt. The fraction of output lost in bankruptcy is μ, which is assumed to be an autocorrelated stochastic process with mean μ_*. Innovations to μ will be termed 'credit shocks'. Variations in μ can be regarded as arising from shocks to intermediaries' monitoring technology, with $\mu_t > \mu_*$ being an adverse shock, in the sense that more product is destroyed during the act of recovering the bankrupt firm's assets ('monitoring').

The marginal addition to expected entrepreneurial revenue from an increase in net worth is measured by the shadow price on the intermediary participation constraint. From the first order condition for investment, this marginal benefit is $\lambda_t = q_t f(\varpi_t)/[1 - q_t g(\varpi_t)]$, which is seen to be the expected 'return on internal funds', and is increasing in the price of capital. The optimal breakeven value ϖ satisfies:

$$q_t i_t f'(\varpi_t) + \lambda_t q_t i_t g'(\varpi_t) = 0, \tag{8}$$

which upon substitution for the shadow price on net worth yields the efficiency condition:

$$f'(\varpi_t) = -\frac{q_t f(\varpi_t)}{1 - q_t g(\varpi_t)} g'(\varpi_t) \tag{9}$$

The expected reduction in entrepreneurial revenue from reallocating a marginal unit of funds from internal to external should match the expected increase in intermediaries' revenue times the return on internal funds.

As the constraint on intermediaries' expected profits is binding, $qig(\varpi) = i - n$, we can use the efficiency condition (9) to write the investment supply function:

$$i_t = -\frac{n_t f'(\varpi_t)}{q_t f(\varpi_t) g'(\varpi_t)} = \frac{n_t}{1 - q_t g(\varpi_t)}. \tag{10}$$

As individual investment is linear in individual net worth, aggregate investment is linear in aggregate net worth, and thus (noting in particular that $g(\varpi)$ is a constant with respect to ω):

$$I_t = \int_0^\infty i_t \phi(\omega) \mathrm{d}\omega - \int_0^{\varpi_t} \mu_t i \phi(\omega) \mathrm{d}\omega = i[1 - \Phi(\varpi_t)\mu_t] \tag{11}$$

Since some output of capital goods is destroyed by the monitoring process, expected aggregate investment is less than the expected aggregate output of all entrepreneurs. The aggregate law of motion for capital is thus:

$$k_{t+1} = (1 - \delta)k_t + i_t[1 - \Phi(\varpi_t)\mu_t], \tag{12}$$

It is this investment 'wedge' that causes investment and consumption goods to differ in price. A reduction in bankruptcy costs μ has a similar reduced form effect to an increased marginal efficiency of investment (Greenwood et al., 1988), a correspondence that was also exploited in the context of credit market shocks by Cooper and Ejarque (2000). In the current model, the risk attaching to credit shocks is time-varying, with shock of given magnitude having more effect when the probability of bankruptcy $\Phi(\varpi_t)$ is high. Households are insulated from the direct effects of such fluctations, which here affect only entrepreneurs and intermediaries. However, they must bear the resultant fluctuations in capital prices. Note that a standard real business cycle model is a special case when agency costs shrink to zero.

It is optimal for the entrepreneur to put his entire net worth at stake, investing all of his assets in the risky production process, since he is risk neutral and expected returns are at least as great as the alternative. The entrepreneur's net worth is given by income earned from participating in the production of consumption goods $w_t^e h_t^e$ and from the value this period of last periods undepreciated capital holdings z_t plus the rental income earned from capital this period. His net worth is therefore given by:

$$n_t = w_t^e h_t^e + z_t(r_t + q_t[1-\delta]); \tag{13}$$

His income from production this period is expected to be $q_t i_t f(\varpi_t)$. If the entrepreneur is solvent he makes a choice between consumption e_t and capital accumulation, otherwise the terms of his contract give all output to the intermediary. His optimisation problem is now to maximise a linear utility function, which exhibits extra 'impatience' compared to households (an assumption which prevents the entrepreneur from becoming self-financing):

$$U = \mathrm{E}_0 \sum_{t=0}^{\infty} (\beta\gamma)^t e_t \qquad \text{where } \gamma \in (0,1)$$

subject to the budget constraint:

$$e_t + q_t z_{t+1} = \{w_t^e + z_t(r_t + q_t[1-\delta])\} \frac{q_t f(\varpi_t)}{1 - q_t g(\varpi_t)} \tag{14}$$

where the last term is seen to be the return on internal funds. The resulting optimum depends only on the relative payoffs of capital today versus capital tomorrow:

$$\beta\gamma \mathrm{E}_t(r_{t+1} + q_{t+1}[1-\delta]) \frac{q_{t+1} f(\varpi_{t+1})}{1 - q_{t+1} g(\varpi_{t+1})} = q_t. \tag{15}$$

The the expected market return and the expected return on internal funds must therefore move reciprocally. Finally, entrepreneurs supply their entire time endowment inelastically to market production. Bankrupt entrepreneurs carry zero capital into the next period, and their net worth consists of the labour income they will accrue during market activity.

3 Analysis

We will now quantify the predictions that the Carlstrom and Fuerst agency cost model has for the credit cycle, and compare them to the data on the US economy from 1961-2005. Three functions are examined that diagnose the descriptive capability of the agency cost model for the behaviour of credit in the cycle: the cross correlation functions between the external funds premium and output, between corporate leverage and output, and between the external funds premium and leverage.

3.1 Theory-based predictions

Our first task is to find distributions for the statistics of interest predicted by the theoretical model. The procedure I follow is essentially the one recommended by Canova (1995). Calibration uncertainty is represented by specifying independent prior distributions over the parameters. I take 2,500 draws from the joint prior, specified in Table 1; for each draw solve for the recursive equilibrium; calculate the cross correlation function based on the resulting theoretical law of motion; finally, apply a kernel smoothing algorithm to approximate the density function at each point of the cross correlation function and find quantiles.

Prior distributions must be assigned to fifteen parameters, falling into two groups. The first group are those parameters standard in the RBC literature, which determine the income shares of labour and capital, the desireability of intertemporal consumption smoothing, and so on. The second group comprises the additional parameters required for the agency cost part of the model. Here, there is less direct guidance from current simulation practice or from estimation studies, and so indirect evidence is sought.

A great deal of attention in the business cycle literature has focussed on σ, the coefficient of household risk aversion, with a wide range of values entertained. Households in our economy do not bear extra risk compared to a model with no

Table 1: Prior distribution of parameters

	Quantile								
	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90
σ	1.147	1.240	1.317	1.394	1.473	1.564	1.675	1.828	2.0909
β	0.986	0.987	0.988	0.989	0.990	0.991	0.992	0.993	0.994
α	0.340	0.345	0.350	0.355	0.360	0.365	0.370	0.375	0.380
δ	0.016	0.017	0.018	0.019	0.020	0.021	0.022	0.023	0.024
$\mu_*\dagger$	0.170	0.228	0.275	0.316	0.357	0.399	0.443	0.493	0.558
$\sigma_\omega\dagger$	0.252	0.326	0.391	0.455	0.516	0.576	0.634	0.692	0.758
ρ_θ	0.915	0.925	0.933	0.940	0.946	0.953	0.959	0.967	0.978
ρ_a	0.862	0.877	0.887	0.896	0.903	0.911	0.919	0.928	0.941
ρ_μ	0.933	0.940	0.945	0.949	0.954	0.958	0.962	0.967	0.974
σ_θ	0.002	0.002	0.002	0.003	0.003	0.003	0.004	0.004	0.005
σ_a	0.002	0.003	0.004	0.005	0.005	0.006	0.006	0.007	0.008
σ_μ	0.008	0.009	0.010	0.011	0.012	0.012	0.013	0.014	0.016

Note: Quantiles are estimated from simulated independent marginal prior distributions.
† Joint distribution is partially restricted by feasible range of γ.

agency problem, so a range of variation that reflects current practice in the literature is chosen. Values at or close to unity are common, and values greater than three are rare. The household discount rate β is pinned down to a narrow range by estimates of the real interest rate. Not having a particular stand on what this number should be, a flat prior on a narrow range centered on a baseline value of 4% per annum is chosen. The production function parameter α represents the share of income going to all capital, public and private. I allow it to take values uniformly on the interval $[.335, .385]$. The depreciation rate of installed capital δ has a mean of 8% per annum, and may is uniform on an interval between 6 and 10%.

The parameter μ_* is the mean fraction of assets destroyed by monitoring, or the costs of bankruptcy, including indirect costs such as business reorganisation. The variance of idiosyncratic risk is σ_ω^2, and the steady state breakeven point ϖ_*. Together, μ_*, σ_ω and $\Phi(\varpi_*)$ determine the steady state return on internal funds, and its reciprocal the entrepreneurial discount factor γ. As $\gamma < 1$, an indirect restriction is placed on their feasible range. To calibrate, take the density function for entrepreneurial technology $\phi(\omega)$ to be log-normal with mean one, so if we define $\tau = \log(1 + \sigma_\omega^2)$, then $\omega \sim \mathrm{LN}(-\tau^2/2, \tau^2)$. This information is translated into observables by noting that the premium on external finance is in the steady state $q_*(1+r_*)-1$, and given μ_*, the price of capital q_* is a function of ϖ_* and σ_ω, via the distribution function $\Phi(\varpi_*)$. For tractability, $\Phi(\varpi_*)$ will be fixed so the quarterly bankruptcy rate is 0.974%, whilst μ_* and σ_ω will be permitted to take values in a range that keeps the return on internal funds positive (and thus $\gamma < 1$). Carlstrom and Fuerst quote bounds for the cost of bankruptcy of between 20% and 36% of assets, on the basis of a comparison of the value of the firm as a going concern with its liquidation value. The implied finance premium ranges from 1.6% to 8% per annum.

The persistence of the forcing processes ρ_θ, ρ_a and ρ_μ are known to be important for model dynamics (see Canova, 1995 for evidence in the RBC setting), so some care must be exercised in setting reasonable ranges of variation. This study adopts values based on the direct estimates found in Ireland (2003) and DeJong et al. (2000) (the latter being Bayesian estimates), for which estimated standard errors are also available[1]. Technology shocks are estimated to be highly persistent AR(1) processes, with an autoregressive coefficient in a narrow range around $\rho_\theta = .95$. Evidence on the persistence of demand shocks from Ireland's study is weaker, as he reports subsample instabilities in its estimated value. Nevertheless, they are likely to have been reasonably persistent over our sample period, and consequently I chose $\rho_a = .90$. The persistence of credit market shocks may be established indirectly from estimated investment shock processes, due to the correspondence between these and credit market shocks in (12). Both studies estimate a value in the region of $\rho_\mu = .95$, with a small standard error of 1.6% in the case of DeJong et al., but with considerably less accuracy in the case of Ireland.

Conditional volatilities of the forcing processes σ_θ, σ_a and σ_μ are set using indirect evidence from variance decompositions. Li and Sarte (2003) find the contribution of credit shocks to the forecast variance decomposition of US manufacturing output to be 18% after ten years. Ireland finds a wide range of possible shares for investment shocks in the variance decomposition of aggregate output. Given this uncertainty, a wide range of variation in σ_μ is allowed for. Under the baseline calibration, shocks to credit markets account for roughly a fifth of output variance, and technology shocks account for roughly two thirds, fractionally below that claimed in the RBC

[1]Estimates are taken the flexible price version of Ireland's model. Additional shocks in his model are seen to be quantitatively unimportant. I assume the consistency of his estimates.

Table 2: Output variance decomposition bounds

		Variance share			
Parameter	range	tech	demand	credit	σ_y
base	-	.68	.09	.24	.020
ρ_θ	[.90, .99]	[.53, .94]	[.13, .02]	[.35, .05]	[.017, .046]
ρ_a	[.84, .96]	[.70, .64]	[.06, .14]	[.24, .22]	[.020, .021]
$\rho_\mu^\dagger$	[.90, .99]	[.74, .43]	[.09, .05]	[.17, .52]	[.020, .026]
σ_θ	[.001, .005]	[.31, .84]	[.18, .04]	[.51, .12]	[.014, .028]
σ_a	[.002, .008]	[.73, .61]	[.02, .18]	[.25, .21]	[.020, .022]
σ_μ	[.007, .016]	[.79, .57]	[.10, .07]	[.11, .36]	[.019, .022]

† Upper tail truncated. The table shows the effect of setting the named parameter at each bound, whilst other parameters are held at the mid-points of their bounds. Totals may not sum to 1 due to rounding. σ_μ scaled by 102.

literature. Household preference ('demand') shocks are set to account for residual variability, making output volatility equal to that in our data sample, 2% per annum.

To verify that the priors on the forcing processes have reasonable implications, Table 2 reports bounds for the decomposition of output variance, when persistence and conditional volatility parameters are at the extremes of their range[2]. In all cases, the share of output variation due to technology shocks remains within a range that contains a consensus figure of 70%. Technology shocks dominate, and output becomes excessively volatile, when their persistance is high. Output volatility is otherwise reasonable. Demand shocks are never more than a fifth of output variance, and credit shocks are on average less than one quarter of output variance. The predictive cross correlation distributions are therefore based on a set of empirically supported priors that have sensible consequences for aggregate fluctuations. Given the modest deviation from the technology-driven business cycle assumption of used in previous tests of the model, we will see below that the model delivers a much improved description of commonly examined features of the data.

[2]It is necessary to rescale the estimated standard deviation of credit shocks by coefficient $\Phi(\varpi)$ which is approximately 102, due to the different units attaching to the shock under the current specification.

3.2 Empirical analysis

In this section, estimation of empirical cross correlation functions for ouput growth, leverage growth, credit spreads, and other macroeconomic variables is described. I follow DeJong et al. (1996) and estimate a Bayesian vector autoregression (BVAR) to derive probability intervals for these statistics. Denote by $\mathbf{y} = \mathrm{vec}(\mathbf{y}_{p+1}^T)$ the $m(T-p)$ vector of observations from time $p+1$ through T. Then the pth order, m dimensional VAR is written as:

$$\begin{aligned}
\mathbf{y} &= \mathbf{Z}\beta_V + \mathbf{e} \quad \text{where} \quad \mathbf{e}|(\beta_V, \mathbf{\Sigma}_V, \mathcal{Z}) \sim N(\mathbf{0}, \Sigma_V \otimes \mathbf{I}_{T-p}) \\
\mathbf{Z} &= \mathbf{I}_m \otimes [\mathbf{1} \quad \mathbf{y}_p^{T-1} \quad \mathbf{y}_{p-1}^{T-2} \quad \dots \quad \mathbf{y}_1^{T-p}] \\
\text{Priors:} & \qquad \beta_V \sim N(\underline{\beta_V}, \underline{H}_\beta^{-1}) \quad H \sim \mathrm{Wi}(\underline{S}^{-1}, \underline{\nu})
\end{aligned} \tag{16}$$

where β_V is a $m(1+mp)\times 1$ vector of parameters and $\mathbf{\Sigma}_V$ is an $m\times m$ inverse precision matrix for the disturbance vector $\mathcal{E} = [\varepsilon_1 \dots \varepsilon_m]'$. The prior mean $\underline{\beta}_V$ is set to zero in all cases except the first autocorrelation of output and leverage growth, which is set to 0.8. Prior precision $\underline{H}$ is controlled using two hyperparameters, denoted tightness w_1 and symmetry w_2. The former shrinks all elements of β_V closer to their prior mean, the latter leaves precision for the lagged dependent variable of an equation unchanged, but symmetrically shrinks lags of other variables to their prior mean. Precision also increases with lag length, geometrically weighting long lags to be closer to their prior means. Prior precision $\underline{S}$ is assumed diagonal in squared OLS regression standard errors from a third order autoregression. The density function for cross correlations is constructed from the BVAR, subject to stationarity, by drawing from the posterior using a Gibbs sampler, and using a kernel smoothing algorithm to obtain an estimate. For details of the posterior sampling algorithm, see Geweke (1999b).

The data is as follows. Aggregate output, consumption and investment are measured by the annual change in the logarithm of real GDP, non-durable consumption plus services and private domestic investment respectively. Aggregate leverage growth is the annual change in the logarithm of the ratio of financial liabilities to financial assets taken from the Flow of Funds data. Financial liabilities include corporate paper (short term debt issued by companies) and bank loans. This measure tracks total liabilities closely. Financial assets include bank deposits, mortgages and the paper of other firms[3]. The majority of past studies use a corporate bond spread to proxy the external finance premium. They differ in whether the spread is taken over T-Bills, a longer dated Treasury issue, or a safe corporate bond. I use two different spreads over a safe corporate issue, to avoid maturity mismatch and periods of excessive volatility in T-Bill yields. The first is a bond rated Baa by Moody's, which are in the middle category of investment grade bonds, described by them as 'subject to moderate credit risk ... and may possess certain speculative characteristics'. The second is a below investment grade, or 'junk', corporate bond. Gertler and Lown (1999, p. 135) argue that this spread is likely to be 'closely correlated with the premium on external funds that...purely bank-dependent borrowers face'. Although this measure is preferred, the relatively recent inception of the market for high yield debt means that data is available only from 1980 onwards[4]. In the other case observations run from 1961:4 through 2005:4.

4 Results

Table 3 compares mean cross correlation functions predicted from the agency cost model with estimated quantities from observed data. The blocks contain information on four different cross correlation functions. In each case, expected values for the theory-based prediction under the base prior and the technology-shock driven prior are followed by empirical means using two alternative measures of the credit spread.

The first row of Table 3 shows the multiple-shock version of the theory model to predict a negative unconditional correlation between credit spreads and output growth across the cycle. Spreads lead the output cycle with a mean correlation of -.317 at one year, indicating that wide spreads are predicted to correlate with low future output growth. The correlation function has a minimum of -.428, indicating that strong current economic conditions are predicted to correlate with narrow current borrowing spreads. For comparison, mean correlations under the technology driven credit cycle hypothesis are also shown in Table 3. The mean prediction for

[3]Stationarity inducing transformations of the data are required for consistency with the theory model. For both leverage and output, HP filtering was tried as an alternative to differencing, in order to capture the relevant fluctations about trend. The resulting pattern of cross correlations was similar throughout.

[4]Because of the short data sample, I experimented with tighter and more informative priors. However, results were not sensitive within reasonable bounds.

the credit spread is a positive correlation with economic growth at leads, with a close to zero correlation at lags.

The picture for the multiple-shock theory model is similar to that seen in the Baa-spread data, where the mean correlation is also negative, and also falls to a trough around $s = 1$ at -.497. The junk spread has a stronger negative lead at one year with a correlation of -.517, but is close to uncorrelated contemporaneously and with lags of output. The finding of negative co-cyclicality is in line with the estimates reported by Koopman and Lucas (2005).

Table 3: Cross correlations: mean credit cycle

					spread/output				
s	-4	-3	-2	-1	0	1	2	3	4
Model (Base)	-.317	-.324	-.335	-.360	-.428	-.280	-.214	-.186	-.171
Model (Tech)	.316	.326	.334	.331	.304	.103	.029	.003	-.003
Data Baa-Aaa	-.116	-.223	-.332	-.424	-.483	-.497	-.487	-.455	-.405
Data Junk-Aaa	-.517	-.506	-.458	-.374	-.276	-.184	-.099	-.029	.021
					leverage/output				
s	-4	-3	-2	-1	0	1	2	3	4
Model (Base)	-.417	-.419	-.413	-.391	-.327	-.262	-.231	-.214	-.202
Model (Tech)	-.730	-.758	-.786	-.816	-.846	-.891	-.864	-.813	-.755
Data Baa-Aaa	-.110	-.196	-.222	-.194	-.145	-.099	-.061	-.032	-.013
Data Junk-Aaa	-.168	-.267	-.280	-.211	-.125	-.064	-.021	.005	.016
					spread/leverage				
s	-4	-3	-2	-1	0	1	2	3	4
Model (Base)	.691	.725	.763	.812	.888	.930	.912	.891	.855
Model (Tech)	-.401	-.412	-.392	-.290	.034	.040	.042	.039	.034
Data Baa-Aaa	.021	.040	.067	.103	.145	.203	.207	.172	.136
Data Junk-Aaa	.041	.076	.121	.168	.211	.195	.149	.095	.044

Note: Means of the theory-based predictive distribution under the base (multiple-shock) prior ('base'), the technology-driven prior ('tech') and the means of the posterior distribution implied by a 4th order vector autoregression. See text for details. $s < 0$ leads; $s > 0$ lags; e.g. spread at time t has theory-predicted correlation with output at time $t - s$ of $-.317$ for $s = -4$.

The two dominant effects underlying the pattern of comovements seen in the theory-based predictions are technology and intermediation cost shocks. A positive technology shock has the conventional effect of increasing the demand for investment. Supply from entrepreneurs (10) is limited by a sharp rise in the cost of borrowing, as their net worth is predetermined. High capital prices increase their return on internal funds, and thus their desire to accumulate capital and reduce leverage. As investment supply shifts outward, capital prices and premia fall, and the return on internal funds returns to normal levels. A positive demand shock has the conventional effect of causing substitution towards consumption and leisure, reducing investment demand and the finance premium.

The effects of a credit shock are as follows. When intermediation becomes more expensive, the burden of extra cost falls on intermediaries themselves, as the contract has them as residual claimants (5). If we think only of the intra-temporal problem for a moment, ignoring the general equilibrium component of the model, we can identify two effects. First, the investment supply curve shifts outwards. Holding net worth n and capital prices q constant, $\partial i/\partial \mu = qn\Phi/(1-qg)^2 > 0$. Second, the return on internal funds is raised, as own net worth is relatively more valuable, and market returns are lowered by (15). Entrepreneurs are induced to accumulate capital, which causes their net worth to rise, reducing leverage. As every unit of resource put into the investment good technology is now expected to yield fewer units of capital, so the price of capital must rise to cover these costs. From the household side, investment becomes less attractive relative to consumption. Households therefore prefer to wait until agency costs fall before investing. A boom in consumption, which reduces hours and so output follows. The immediate increase in credit premia is reversed as entrepreneurial net worth increases, and agency costs are ameliorated. The correlation pattern in the base model is more similar to that seen in the data than that of the technology model because credit shocks raise default premia even as investment demand falls. As a relatively small contribution from credit shocks is required, negative diagnostics based on an assumption of technology-driven credit cycles were overly pessimistic.

The cross correlation of leverage and output growth in the data shows a mild negative leading relationship, stable and similar in mean across specifications, see the second block of Table 3. High future output is therefore correlated with both a narrow spread and low leverage this quarter. On aggregate, the value of firm borrowing rises more slowly than the value of firm assets in the upswing of an economic cycle, in spite of easier borrowing conditions. This hints that firms meet a good deal of their short term investment needs from internal funds. Given that cashflows are procyclical, demand for short-term finance is likely to be countercyclical. If firms mainly borrowed to finance higher future output, leverage would be a positive and leading correlate with GDP. The agency cost model does predict that lower leverage correlates with cyclical upswings, albeit more strongly than seen in the data.

The third block of Table 3 shows the correlation between spreads and aggregate leverage. The base agency cost model predicts a strong positive correlation at all leads and lags. In the agency cost view, wide spreads correlate with high leverage. In the aggregate data, there is close to a zero correlation between leverage and either measure of the spread. This finding suggests that agency costs are not the sole source of aggregate risk driving credit spreads, and may not be the main one. Table 4 shows mean correlations between consumption, investment and output for comparison with other business cycle models. Notice that the demand shocks help bring the mean cross correlations for macroeconomic variables closer to the mean of the data.

Table 4: Cross correlations: Macroeconomic variables

	consumption/output								
s	-4	-3	-2	-1	0	1	2	3	4
Model (Base)	.494	.505	.510	.498	.442	.581	.642	.671	.686
Model (SD)	.502	.517	.533	.553	.583	.609	.640	.669	.695
Model (Tech)	.818	.847	.877	.906	.929	.913	.906	.900	.895
Data Baa-Aaa	.221	.295	.386	.488	.596	.556	.443	.299	.186
	investment/output								
s	-4	-3	-2	-1	0	1	2	3	4
Model (Base)	.528	.555	.597	.674	.841	.608	.493	.426	.380
Model (SD)	.727	.756	.784	.805	.806	.757	.696	.634	.575
Model (Tech)	.809	.839	.869	.897	.916	.881	.828	.771	.715
Data Baa-Aaa	.163	.272	.449	.628	.767	.607	.419	.246	.109

Note: Means of the theory-based predictive distribution from the multiple-shock prior ('base'), from the technology and demand-shock driven prior (supply and demand 'SD'), from the technology-driven prior ('tech'); the mean of the posterior distribution implied by a 4th order vector autoregression using Baa-Aaa spread to measure the funds premium. See text for details; $s < 0$ leads; $s > 0$ lags.

Figure 1 shows two dimensional contour plots of cross correlation probability distributions. Contours represent probability quantiles of the distributions at 0.5, 2.5, 97.5 and 99.5%. For example, the 99% probability interval can be read off as the

Figure 1: Probability intervals for cross correlation functions

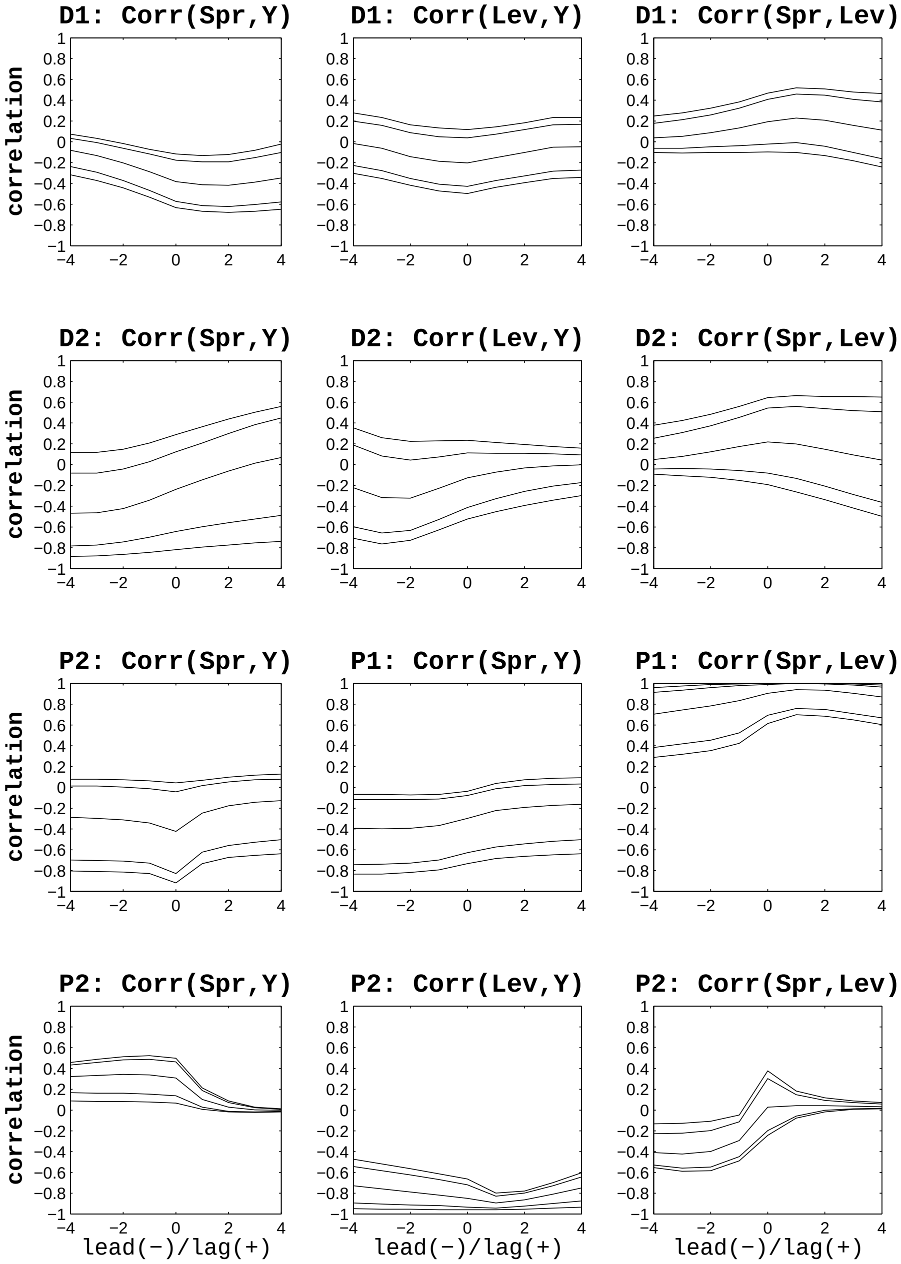

D1: Baa-spread data; D2: Junk-spread data; P1: base prior (multiple shocks); P2: technology shock prior. Contours represent probability intervals at 99% and 95%; the central line is the mean. These quantiles were obtained by simulation from the posterior distribution of each BVAR or from theory predictive distributions. See the main text for details.

area within the top and bottom contours. The central line traces the posterior mean of the distributions. The figure has four rows: rows one and two (labeled D for data) correspond to the empirical model using Baa data, and the model using Junk data. In the Baa data, the correlation between the spread and output is less than zero with at least 90% probability between $s = -2$ and $s = 4$. In the Junk data, probability intervals are wider, and closer to the prior of uncorrelatedness, due to the shorter data sample. For the remaining correlations, which are between identical variables in each specification but different samples, the pictures are predictably similiar.

The third and fourth rows of Figure 1 gives distributions for the theory model. There are two alternative priors. The first is our baseline multiple shock model (P1), and the second turns off both credit and demand shocks (P3), leaving a technology-driven credit cycle. Demand shocks of the magnitude I consider have little relevance for the credit cycle, although they are relevant for other business cycle moments (Table 4). The baseline theory model has wide probability bands for the credit spread-output correlation. The 90% probability interval runs from zero correlation to a high negative correlation. The predictions of the technology driven model for all correlations are concentrated in a narrower 99% probability interval than under the base prior. In one sense, these predictions are therefore more robust to prior uncertainty.

An informal judgement on the performance of the model under the baseline multiple-shock prior (P1) relative to the technology-based prior (P2) can be reached by assessing the overlap between their respective density functions, and those estimated using the BVAR. This is the basis of the diagnostic criterion proposed by DeJong et al. (1996), and the formal odds ratio of Geweke (1999a). Agreement between the multiple-shock model and Baa data on the credit-output cycle is high, especially at leads. The leverage-ouput correlation also appears to show some overlap, in spite of its lower mean. A drawback of the technology-shock prior is that high probability is often assigned to regions of the sample space that are assigned a low probability by the data. That is, intervals overlap either very little, or are narrow at points where the data is more diffuse. The spread-leverage correlation is not well described under either prior, with little overlap in the probability intervals.

A prominent feature of the credit shock model (P1) compared to the technology shock model (P2) is the wide probability interval seen particularly in the tails of the cross correlation function. To determine which parameters are responsible for this dispersion, denote the cross covariance as function of the calibrated model parameters $\gamma(\mathbf{p})$. To gauge the sensitivity of $\gamma(.)$ to $\mathbf{p}$ in the neighbourhood of the baseline calibration, I compute the elasticity of the function with respect to each parameter. Table 5 details these elasticities for the cross correlation between the premium on external finance and output, for one year lead and lag. The rows of Table 5 show that

Table 5: Sensitivity elasticities

	Cross correlation between credit spread and output								
s	-4	-3	-2	-1	0	1	2	3	4
σ	-.388	-.393	-.388	-.356	-.285	-.256	-.229	-.215	-.210
β	-2.44	-1.29	1.10	6.49	15.6	6.33	-1.37	-4.91	-5.96
α	1.06	1.12	1.20	1.34	1.53	1.38	1.26	1.21	1.21
δ	-.003	-.022	-.047	-.084	-.134	-.097	-.071	-.065	-.067
μ_*	.818	.810	.794	.760	.711	.869	1.02	1.11	1.15
Φ	.860	.855	.841	.818	.811	.992	1.20	1.33	1.39
σ_ω	3.12	3.01	2.58	1.30	-1.25	-.193	.485	.752	.838
ρ_θ	-6.64	-6.53	-6.44	-6.42	-6.54	-6.64	-6.73	-6.79	-6.83
ρ_a	-.208	-.176	-.147	-.140	-.177	-.232	-.281	-.309	-0.322
ρ_μ	27.3	26.0	23.2	16.9	6.10	22.6	37.8	46.3	50.3
σ_θ	-1.01	-1.01	-.997	-.956	-.875	-.812	-.763	-.748	-.751
σ_a	-.146	-.148	-.145	-.130	-.100	-.091	-.084	-.083	-.084
σ_μ	1.15	1.16	1.14	1.09	.974	.902	.848	.831	.835

Rows give the sensitivity elasticity for the cross correlation of the credit spread at t with output at $t-s$ with respect to the named parameter. For example, a one percent increase in the CRRA coefficient σ would reduce $\text{Corr}(\text{spr}_t, y_{t+4})$ by 0.388 per cent.

the most informative parameters are those that control the unconditional variance of the technology and intermediation cost processes, the household rate of time preference and the capital share, followed by the parameters controlling idiosyncratic risk and agency costs. The variance terms matter because they determine which shocks dominate, on average, and therefore which pattern of comovements dominate, on average. The discount factor matters because it a central determinant of the optimal capital stock in the neoclassical growth model, and therefore variation in it shifts the steady state up and down a concave production function. The correlations are locally robust to variations in a majority of parameters, in particular some of the harder to calibrate finance-related parameters such as the bankruptcy probability

Φ. However, the main risk to the model's predictions attaches to the tails of the cross covariance function as ρ_μ varies, as locally its sensitivity elasticity is large.

5 Conclusion

This paper has assessed how well a macroeconomic model incorporating agency costs describes the cyclical movement of credit spreads and corporate financial data. Previous research shows that macroeconomic risk is a significant factor for credit spreads, and new evidence from a Bayesian VAR supports the finding that spreads lead the cycle with a negative correlation. It was shown that multiple plausible sources of macro risk, including in particular shocks arising in the credit market, improve the descriptive power of the model for unconditional correlations seen in the data. Technology shocks remained the dominant source of output fluctuations, as in previous studies.

Not all aspects of the credit cycle were well described. In particular, spreads appear close to uncorrelated with aggregate leverage in the data. Also, in focusing upon the credit cycle, detailed consideration of other important aspects of the business cycle were left aside. However, the model does predict positive cross correlations between consumption, investment and output, as in the data.

The analysis in this paper points to some further issues that were not the foremost concerns of papers such as Canova (1995), which first proposed the kind of global sensitivity analysis performed here. It is nowadays routine to specify prior distributions for parameters of theoretical macroeconomic models as part of a Bayesian estimation. For example, it is straightforward to check the predictive distribution of moments required in estimation, and this could aid in diagnosing low marginal likelihood values. Also, local sensitivity analysis may aid in the specification of priors by identifying the most important parameters. In particular, researchers often choose to fix certain parameters and to estimate others. Presumably the data will be most informative for those parameters with a high elasticity with respect to moments used in estimation.

References

Ben S. Bernanke and Mark Gertler. Agency costs, net worth, and business fluctuations. *American Economic Review*, 79(1):14–31, 1989.

Fabio Canova. Sensitivity analysis and model evaluation in simulated dynamic general equilibrium economies. *International Economic Review*, 36(2):477–501, May 1995.

Charles T. Carlstrom and Timothy S. Fuerst. Agency costs, net worth, and business fluctuations: A computable general equilibrium analysis. *American Economic Review*, 87(5):893–910, 1997.

Russel W. Cooper and Joao Ejarque. Financial intermediation and aggregate fluctuations: a quantitative analysis. *Macroeconomic Dynamics*, 4, December 2000.

David N. DeJong, Beth Fisher Ingram, and Charles H. Whiteman. A Bayesian approach to calibration. *Journal of Business and Economic Statistics*, 14(1):1–9, January 1996.

David N. DeJong, Beth Fisher Ingram, and Charles H. Whiteman. Keynesian impulses versus Solow residuals: identifying sources of business cycle fluctuations. *Journal of Applied Econometrics*, 15:311–329, 2000.

Mark Gertler and Cara Lown. The information in the high-yield bond spread for the business cycle: Evidence and some implications. *Oxford Review of Economic Policy*, 15(3):132–150, 1999.

John Geweke. Computational experiments and reality. Mimeo, University of Minnesota, June 1999a.

John Geweke. Using simulation methods for Bayesian econometric models: inference, development and communication. *Econometric Reviews*, 18:1–126, 1999b.

Joao F. Gomes, Amir Yaron, and Lu Zhang. Asset prices and business cycles with costly external finance. *Review of Economic Dynamics*, 6:767–788, 2003.

Jeremy Greenwood, Zvi Hercowitz, and Gregory W. Huffman. Investment, capacity utilization, and the real business cycle. *American Economic Review*, 78, No. 3: 402–417, June 1988.

Peter N. Ireland. Endogenous money or sticky prices? *Journal of Monetary Economics*, 50:1623–1648, 2003.

Siem Jan Koopman and André Lucas. Business and default cycles for credit risk. *Journal of Applied Econometrics*, 20:311–323, 2005.

Noh-Sun Kwark. Default risks, interest rate spreads, and business cycles: Explaining the interest rate as a leading indicator. *Journal of Economic Dynamics & Control*, 26:271–302, 2002.

Wenli Li and Pierre-Daniel G. Sarte. Credit market frictions and their direct effects on U.S. manufacturing fluctuations. *Journal of Economic Dynamics & Control*, 28:419–443, 2003.

Appendices (Not for Publication)

A Linearisation

In order to obtain a numerical solution for the law of motion, I log-linearise around the steady state. As a reminder, table 6 lists the variables. The full set of linearised relations, defining the constant $\Upsilon = \alpha k_*^{\alpha-1} h_*^{1-\alpha-\alpha_e}$ as the steady state real interest rate, follow (recall that entrepreneurs supply their entire time endowment inelastically).

$$
\begin{aligned}
0 &= -\hat{y}_t + \hat{\theta}_t + \alpha \hat{k}_t + (1-\alpha-\alpha_e)\hat{h}_t && (17)\\
0 &= -\hat{r}_t + \hat{\theta}_t + (1-\alpha-\alpha_e)\hat{h}_t - (1-\alpha)\hat{k}_t && (18)\\
0 &= R_*^d (I_* - n_*)\hat{R}_t^d + I_*(R_*^d - \varpi_* q_*)\hat{I}_t - R_*^d n_* \hat{n}_t - I_* \varpi_* q_* (\hat{\varpi}_t + \hat{q}_t) && (19)\\
0 &= -c_* \hat{c}_t - e_* \hat{e}_t - I_* \hat{I}_t + y_* \hat{y}_t && (20)
\end{aligned}
$$

These are respectively the production function, definition of the capital rental rate, definition of the premium on external funds, and the aggregate budget constraint. The law of motion for aggregate capital is given by

$$k_* \hat{k}_{t+1} = (1-\delta) k_* \hat{k}_t + I_*(1 - \Phi(\varpi_*)\mu_*)\hat{I}_t - I_* \mu \phi(\varpi_*)\varpi_* \hat{\varpi}_t - I_* \Phi(\varpi_*)\mu_* \hat{\mu}_t \quad (21)$$

Notice that as $\mu \to 0$ in (21), the expression collapses to the standard log-linear accumulation equation. Notice also that this is similar to the Greenwood et al. (1988) shock to the marginal efficiency of capital when agency costs are held constant.

The central Euler equations governing household and entrepreneurial capital accumulation are

$$
\begin{aligned}
&\mathrm{E}_t \beta (q_*[1-\delta] + \Upsilon)\{(\chi[1-\sigma] - 1)\hat{c}_{t+1} + a_{t+1}\} \\
&\quad - \beta (q_*[1-\delta] + \Upsilon)(1-\chi)(1-\sigma)\frac{h_*}{1-h_*}\hat{h}_{t+1} + \beta q_*(1-\delta)\hat{q}_{t+1} + \beta \Upsilon \hat{r}_{t+1} = \\
&\qquad q_* \hat{q}_t + \beta (q_*[1-\delta] + \Upsilon)\{(\chi[1-\sigma] - 1)\hat{c}_t + a_t\} \\
&\qquad\quad - \beta (q_*[1-\delta] + \Upsilon)(1-\chi)(1-\sigma)\frac{h_*}{1-h_*}\hat{h}_t \quad (22)
\end{aligned}
$$

Table 6: Model variables

Variable	Description
y	output
c	household consumption
k	aggregate capital
h	household hours worked
q	price of capital goods
I	investment
r	return on capital (equity)
R^d	excess return on debt (external funds premium)
ϖ	break-even profitability
e	entrepreneurial consumption
n	entrepreneurial net worth
z	entrepreneurial capital
θ	technology shock
a	household preference shock
μ	intermediation cost shock

and

$$\begin{aligned}&[\beta\gamma\{q_* f(\varpi_*)(1-\delta)+[q_*(1-\delta)+\Upsilon]f(\varpi_*)\}+q_* g(\varpi_*)]\,\hat{q}_{t+1}\\ &\qquad[\beta\gamma\{q_*(1-\delta)+\Upsilon\}f'(\varpi_*)+q_* g'(\varpi_*)]\,\varpi_*\hat{\varpi}_{t+1}+\beta\gamma\Upsilon f(\varpi_*)\hat{r}_{t+1}\\ &\qquad\qquad[1-q_* g_*(\varpi_*)]\hat{b}_{t+1}+\Phi(\varpi_*)\mu_*\hat{\mu}_{t+1}=[1-q_* g(\varpi_*)]\{\hat{q}_t+\hat{b}_t\}.\end{aligned} \quad (23)$$

where b is a shock to entrepreneurial preferences not used in the main paper. These Euler equations are the only places where the two preference shocks impact. The shocks have the effect of raising the demand for consumption and leisure, in the case of the household, and of consumption rather than capital accumulation in the case of the entrepreneur.

The efficiency condition for the financing contract (or alternatively, the link between bankruptcies in the capital goods sector and capital prices) is

$$\begin{aligned}0=\hat{q}_t-q_*&\left[\left(\frac{\phi(\varpi_*)}{1-\Phi(\varpi_*)}\right)^2\mu_* f(\varpi_*)+\frac{\phi(\varpi_*)\mu_* f'(\varpi_*)}{1-\Phi(\varpi_*)}+\right.\\ &\left.\frac{\phi(\varpi_*)\mu_* f(\varpi_*)}{1-\Phi(\varpi_*)}+\mu_*\phi(\varpi_*)\right]\varpi_*\hat{\varpi}_t-q_*\left[\Phi(\varpi_*)\mu_*+\frac{\phi(\varpi_*)\mu_* f(\varpi_*)}{1-\Phi(\varpi_*)}\right]\hat{\mu}_t.\end{aligned} \quad (24)$$

Table 7: Timing assumptions

Stage	Description
1	All aggregate shocks to preferences and technologies are realised.
2	Household and entrepreneurial labour supply decisions; firm capital rental and hours decisions; production of consumption good.
3	If household wishes to add to its capital stock, it must give consumption goods to the intermediary for investment in entrepreneurs' capital good production technology.
4	Entrepreneurs invest all of their own funds plus borrowings in the capital good technology; an idiosyncratic shock is realised, and each entrepreneur observes his own output.
5	If solvent, the entrepreneur repays his loan; if bankrupt, the intermediary captures the residual output after some fraction of the investment is lost via monitoring.
6	Solvent entrepreneurs make their choice between consumption and capital accumulation.

Notice that an implication of this expression is that as agency costs shrink ($\mu \to 0$), capital prices approach a constant, and steady state capital prices approach unity, in which case the household Euler equation (22) collapses to the standard RBC Euler equation.

Finally, the investment 'supply curve', the evolution of entrepreneurial net worth and entrepreneurs' budget constraint are respectively

$$\begin{aligned} 0 &= I_*[1 - q_* g(\varpi_*)]\hat{I}_t - I_* q_* g(\varpi_*)\hat{q}_t - I_* q_* g'(\varpi_*)\varpi_*\hat{\varpi}_t - n_*\hat{n}_t \\ &\quad + I_* q_* \Phi(\varpi_*)\mu_*\hat{\mu}_t \qquad (25) \\ 0 &= -n_*\hat{n}_t + z_*(q_*[1-\delta] + \Upsilon)\hat{z}_t + z_* q_*(1-\delta)\hat{q}_t + z_*\Upsilon\hat{r}_t \qquad (26) \\ z_* q_* \hat{z}_{t+1} &= q_*(I_* f(\varpi_*) - z_*)\hat{q}_t + I_* q_* f(\varpi_*)\hat{I}_t + I_* q_* f'(\varpi_*)\varpi_*\hat{\varpi}_t - e_*\hat{e}_t. \qquad (27) \end{aligned}$$

To summarise, I have detailed the locally valid log-linear approximation to the behavioural relations of the model which may be solved numerically, once values have been assigned to parameters.

B Analytic cross correlation functions

This section details a computationally simple method for calculating the cross correlation function for the theory model. Traditionally, moment calculation was performed by simulating a large data set using draws from the assumed model distribution. This is computationally expensive and introduces sampling variation which we want to avoid when, for example, calculating numerical derivatives.

A solution to the model is a set of paths for the n endogenous variables $\mathbf{y}$ and m states $\mathbf{x}$ that satisfy all behavioural relations and constraints for any realisation of the k exogenous variables $\mathbf{z}$, such that all markets clear. As the model is non-linear, we study behaviour in the neighbourhood of the steady state. I use the MSV procedure to obtain a law of motion of the form:

$$\mathbf{x}_{t+1} = \mathbf{P}\mathbf{x}_t + \mathbf{Q}\mathbf{z}_t \tag{28}$$

$$\mathbf{y}_t = \mathbf{M}\mathbf{x}_t + \mathbf{N}\mathbf{z}_t \tag{29}$$

$$\mathbf{z}_{t+1} = \mathbf{R}\mathbf{z}_t + \epsilon_t \quad \text{where} \quad \mathrm{E}\epsilon_t = 0 \text{ and } \mathrm{E}\epsilon_t\epsilon_t' = \Sigma. \tag{30}$$

Once this law is in hand, any desired moments can be calculated analytically. The log-linearised version of the model is given in Appendix A.

We start by rewriting the state space model (28)–(30) in VAR form, by stacking the relevant matrices as follows:

$$\begin{pmatrix} \mathbf{y}_t \\ \mathbf{x}_t \\ \mathbf{z}_t \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{n\times k} & \mathbf{MP} & \mathbf{MQ}+\mathbf{NR} \\ \mathbf{0}_{m\times n} & \mathbf{P} & \mathbf{Q} \\ \mathbf{0}_{k\times n} & \mathbf{0}_{k\times m} & \mathbf{R} \end{pmatrix} \begin{pmatrix} \mathbf{y}_{t-1} \\ \mathbf{x}_{t-1} \\ \mathbf{z}_{t-1} \end{pmatrix} + \begin{pmatrix} \mathbf{N} \\ \mathbf{0}_{m\times k} \\ \mathbf{I}_k \end{pmatrix} \epsilon_t$$

If we define ξ, $\mathbf{\Lambda}$ and $\mathbf{u}$ in the obvious way, then we can write this as

$$(\mathbf{I} - \mathbf{\Lambda} L)\xi_t = \mathbf{u}_t \quad \text{where} \quad \mathrm{E}\mathbf{u}\mathbf{u}' = \mathbf{\Omega}.$$

Inverting the polynomial in the lag operator and expanding yields the MA representation

$$\xi_t = (\mathbf{I} + \mathbf{\Lambda} L + \mathbf{\Lambda}^2 L^2 + \mathbf{\Lambda}^3 L^3 + ...)\mathbf{u}_t.$$

This corresponds to the absolutely summable sequence of MA coefficients given in Hamilton (1994, §10.2) whereupon the sth autocovariance is

$$\mathbf{\Gamma}_s = \sum_{v=0}^{\infty} \mathbf{\Lambda}^{s+v}\mathbf{\Omega}(\mathbf{\Lambda}^v)'$$

and the sth and $s-1$th autocovariance are related by the simple recursion

$$\mathbf{\Gamma}_s = \mathbf{\Lambda}\mathbf{\Gamma}_{s-1},$$

with the relationship between leads and lags being

$$\mathbf{\Gamma}_{-s} = \mathbf{\Gamma}'_s.$$

Once we have calculated $\mathbf{\Gamma}_0$, it is therefore computationally cheap to obtain any other autocovariance desired. Using the sparse matrix structure in Matlab saves computer memory. I found that terms in $\mathbf{\Lambda}^J$ were numerically less than 1e-15 for an expansion of J =1.5e3.

C BVAR setup

We briefly describe the construction a model of an observed vector time series $\mathcal{Y}^o$ of dimension $T \times m$ using a BVAR. We will formulate the model as a special case of the SUR model described in the BACC manual. The notation $\mathcal{Y}_\tau^T$ will refer to observations on $\mathcal{Y}^o$ in rows τ up to T. The calligraphic script is to remind us that this is observed sample data, rather than data simulated from our theory model.

We will henceforth consider the specification of a VAR of dimension m and order p. Define the $T-p \times 1+mp$ matrix $\mathcal{X}$ as

$$\mathcal{X} = [\mathbf{1} \quad \mathcal{Y}_p^{T-1} \quad \mathcal{Y}_{p-1}^{T-2} \quad \dots \quad \mathcal{Y}_1^{T-p}] \tag{31}$$

where $\mathbf{1}$ is a $T-p$ vector of 1s. The $m(T-p) \times m(1+mp)$ matrix $\mathcal{Z}$ is then defined as

$$\begin{align} \mathcal{Z} &= \mathbf{I}_m \otimes \mathcal{X} \tag{32} \\ &= \begin{pmatrix} \mathcal{X} & 0 & \dots & 0 \\ 0 & \mathcal{X} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathcal{X} \end{pmatrix} \tag{33} \end{align}$$

Finally, vectorise the data matrix, defining $\mathcal{Y} = \text{vec}(\mathcal{Y}_{p+1}^T)$. We may then write the system as

$$\mathcal{Y} = \mathcal{Z}\beta_V + \mathcal{E} \quad \text{where} \quad \mathcal{E}|(\beta_V, \mathbf{\Sigma}_V, \mathcal{Z}) \sim N(\mathbf{0}, \Sigma_V \otimes \mathbf{I}_{T-p}) \tag{34}$$

where β_V is a $m(1+mp) \times 1$ vector of parameters and $\mathbf{\Sigma}_V$ is an $m \times m$ inverse precision matrix for the disturbance vector $\mathcal{E} = [\varepsilon_1 \dots \varepsilon_m]'$. The posterior density over β_V and $\mathbf{\Sigma}_V$ is calculated using the methods described in Geweke (1999b).

The empirical cross correlation function is calculated in the manner described in Appendix 5.A, which requires only that the VAR(p) of order m be rewritten in VAR(1) form. Supposing that

$$\mathcal{Y}_t = B_1\mathcal{Y}_{t-1} + B_2\mathcal{Y}_{t-2} + \ldots + B_p\mathcal{Y}_{t-p} + \mathcal{E}_t,$$

then the system in VAR(1) form becomes

$$\begin{pmatrix} \mathcal{Y}_t \\ \mathcal{Y}_{t-1} \\ \vdots \\ \mathcal{Y}_{t-p+2} \\ \mathcal{Y}_{t-p+1} \end{pmatrix} = \begin{pmatrix} \mathbf{B}_1 & \mathbf{B}_2 & \dots & \mathbf{B}_{p-1} & \mathbf{B}_p \\ \mathbf{I}_m & 0 & \dots & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 \\ 0 & & \dots & \mathbf{I}_m & 0 \end{pmatrix} \begin{pmatrix} \mathcal{Y}_{t-1} \\ \mathcal{Y}_{t-2} \\ \vdots \\ \mathcal{Y}_{t-p+1} \\ \mathcal{Y}_{t-p} \end{pmatrix} + \begin{pmatrix} \mathcal{E}_t \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}$$

Setting $\mathbf{\Lambda}$ to be the coefficient matrix on the right of this equation, we may then proceed as before, subject to the restriction that no root lies outside the unit circle.

Prior precision of variable j in equation i at lag l is given by:

$$S(i,j,l) = \frac{w_1 g(l) f(i,j) s_i}{s_j} \tag{35}$$

where w_1 is overall tightness, $g(l)$ is a scalar polynomial in l with geometric coefficients, $f(i,i) = 1$ and $f(i,j) = w_2$ for $i \neq j$ and s_i is the inverse precision of an autoregression in the ith variable.

D Data sources

Leverage - FRB Release Z1, March 2006, Table B102.
Junk bond spread - Salamon Smith Barney (Courtesy of Cara Lown).
Moody's Aaa bond yield - Datastream.
Output, Consumption, Investment - BEA, 1st Quarter 2006.

Figure 2: Data plots