1 Introduction

The empirical finance literature has provided substantial evidence that risk premia are time-varying (e.g. Campbell and Shiller (1988), Fama and French (1989), Ferson and Harvey (1991), Cochrane (2005)). Yet, standard business cycle models such as the real business cycle model, or the DSGE models used for monetary policy analysis, largely fail to generate the level and the cyclicality of risk premia. This seems an important neglect, since empirical work suggests a tight connection between risk premia and economic activity. For instance, Philippon (2008) and Gilchrist and Zakrajsek (2007) show that corporate bonds spreads are highly correlated with real physical investment, both in the time series and in the cross-section. A large research, summarized in Backus, Routledge and Zin (2008), shows that the stock market, the term premium, and (negatively) the short rate all lead the cycle.[1]

I propose to introduce time-varying risk premia in a standard business cycle model, through a small, stochastically time-varying risk of a "disaster", following the work of Rietz (1988), Barro (2006), and Gabaix (2007). Existing work has so far confined itself to endowment economies, and hence does not consider the feedback from time-varying risk premia to macroeconomic activity. This risk of an economic disaster could be a strictly rational expectation, or more generally it could reflect a time-varying belief which may differ from the objective probability: it is the "perceived probability" of disaster which matters. Of course in reality this change in probability of disaster may be an endogenous variable and not an exogenous shock. But it is useful to understand the effect of an increase in aggregate risk premia on the macroeconomy. This simple modeling device captures the idea that aggregate uncertainty is sometimes high, i.e. people sometimes worry about the possibility of a deep recession, as seems to be the case for instance in the Fall of 2008. It also captures the idea that there are some asset price changes which are not obviously related to current or future TFP, i.e. "bubbles", and which in turn affect the macroeconomy.

Introducing time-varying risk premia requires solving a model using nonlinear methods, i.e. going beyond the first-order approximation and considering "higher order terms". Researchers disagree on the importance of these higher order terms, and a fairly common view is that they are irrelevant for macroeconomic quantities (e.g. Tallarini (2000), Campanale et al. (2007)). My results show, however, that these terms can have a significant effect on macroeconomic dynamics.

The proposal studies two channels through which time-varying risk premia (here caused by an increase in the probability of disaster) affect macroeconomic aggregates.[2] The first channel is a cost-of-capital effect. When the elasticity of substitution of consumption is larger than unity, an increase in the risk premium leads to a decrease in investment, employment and output. These business cycle dynamics occur with no change in total factor productivity. Under some conditions the increase in probability of disaster is observationally equivalent to a preference shock, which is interesting since these shocks appear to be important in accounting for the data (e.g. Smets and Wouters (2003)). The simple model is also, at

[1]Schwert (1989) and Bloom (2008) also show that stock market volatility negatively leads economic activity.

[2]Many of the results obtained here would be qualitatively similar if I was assuming a time-varying volatility of TFP shocks, or an exogenously time-varying risk aversion. The time-varying risk of disaster is essentially a convenient, tractable way of generating time-varying risk premia.

least qualitatively, and potentially quantitatively, consistent with the lead-lag relationships between asset prices and the macroeconomy mentioned above.

The second channel is that changes in risk-premia affect the willingness to engage in risky investments. Economic activity turns to lower risk, lower expected return projects, which has the effect of lowering aggregate productivity and output. This reallocation effect has interesting micro-implications, for which I provide some support in the proposal.

There are several interesting extensions. First, one could introduce a collateral channel: increases in risk premia would reduce the value of the collateral and thus affect physical investment. Second, it would be interesting to embed these time-varying risk premia in a standard New Keynesian framework. Another interesting addition would be to incorporate defaultable debt. This work could be connected to the literature on Depressions (e.g., Kehoe and Prescott (2007)), and could also consider alternative modeling of the dynamics of disasters (e.g. persistence in low growth regimes, recoveries following disasters, and learning about the disaster state or about the disaster probability).

The proposal is organized as follows: Section 2 reviews the literature; Section 3 embeds a time-varying risk of disaster in a simple business cycle model, to study the cost-of-capital channel both theoretically and quantitatively. Section 4 illustrates the reallocation channel in a simple model and provides some cross-sectional evidence which supports it.

2 Literature Review

This proposal is mostly related to three strands of literature. First, a large literature in finance builds and estimates models which attempt to match not only the equity premium and the risk-free rate, but also the predictability of returns and potentially the term structure. Two prominent examples are Bansal and Yaron (2004) and Campbell and Cochrane (1999). However, this literature is limited to endowment economies, and hence is of limited use to analyze the business cycle or to study policy questions.

Second, a smaller literature studies small business cycle models (i.e. they endogenize consumption, investment and output), and attempts to match not only business cycle statistics but also asset returns first and second moments. My project is closely related to these papers (A non-exhaustive list would include Jermann (1998), Tallarini (2000), Boldrin, Christiano and Fisher (2001), Lettau and Uhlig (2001), Kaltenbrunner and Lochstoer (2008), Campanele et al. (2008), Croce (2005), Gourio (2008c), Papanikolaou (2008), Kuehn (2008), Uhlig (2006), Jaccard (2008)). Most of these papers study the implications of productivity shocks, and generally consider only the mean and standard deviations of return, and not the predictability of returns. Many of these papers abstract from hours variation. Several of these papers note that quantities dynamics are unaffected by risk aversion,[3] hence it is sometimes said that matching asset pricing facts need not affect the business cycle implications of the model.[4] Recently some authors have also tried to generate time-varying risk premia in monetary models (e.g. Swanson and Rudebusch (2008a and 2008b)). The long-run target

[3]Fernandez-Villaverde et al. (2008) use perturbation methods and report that the first three terms are independent of risk aversion (there is, of course, a steady-state adjustment).

[4]Lucas (2003) and Cochrane (2005) emphasize this "separation theorem".

is to have a medium-scale DSGE model (as in Smets and Wouters (2003) or Christiano, Eichenbaum and Evans (2005)) that is roughly consistent with asset prices.

Finally, the paper draws from the recent literature on "disasters" or rare events (Rietz (1988), Barro (2006), Barro and Ursua (2008), Gabaix (2007), Farhi and Gabaix (2008), Martin (2007), Gourio (2008a and 2008b), Julliard and Ghosh (2008), Santa Clara and Yan (2008), Wachter (2008), Weitzmann (2007)). Disasters are a powerful way to generate large risk premia. Moreover, disasters are relatively easy to embed into standard macroeconomic models, as we will see below.

The project will also relate its findings to the empirical finance literature discussed above linking risk premia and the business cycle. There has been much interest lately in the evidence that the stock market leads TFP and GDP, which has motivated introducing "news shocks" (e.g.,Beaudry and Portier (2006)), but my model suggests that this same evidence could also be rationalized by shocks to risk premia (i.e. shocks to the probability of disasters).

Last, the paper has the same flavor as Bloom (2008) in that an increase in aggregate uncertainty creates a recession, but the mechanism and the focus of the paper (asset prices) is different.

3 Time-varying risk of disaster and business cycles

I first present a simple AK model which can be solved exactly and illustrates the role of the intertemporal elasticity of substitution (subsection 1). Then, I introduce a more general model for quantitative work (subsection 2), which is then solved numerically to study its implications (subsection 3).

3.1 Analytical Example: a simple AK economy

To highlight the key mechanism, consider a simple economy with a representative consumer who has Epstein-Zin preferences, i.e. his utility V_t satisfies the recursion:

$$V_t = \left((1-\beta) C_t^{1-\gamma} + \beta E_t \left(V_{t+1}^{1-\theta} \right)^{\frac{1-\gamma}{1-\theta}} \right)^{\frac{1}{1-\gamma}}, \tag{1}$$

where C_t is consumption; note that θ measures risk aversion towards static gambles, γ is the inverse of the intertemporal elasticity of substitution (IES) and β reflects time preference.[5] This consumer has access to an AK technology:

$$Y_t = A_t K_t,$$

where Y_t is output, K_t is capital, and A_t is a stochastic technology which is assumed to follow a stationary Markov process with transition Q. The resource constraint is:

$$C_t + I_t \leq A_t K_t.$$

[5] As explained in Gourio (2008b) and Wachter (2008), the disaster model with standard power utility has counterfactual implications, which are resolved with Epstein-Zin utility.

The economy is randomly hit by disasters: in period $t+1$, with probability p_t, there is a disaster, which destroys a share b_k of the capital stock. This could be due to a war which physically destroys capital, to expropriation of capital holders (e.g. if the capital is taken away and then not used as effectively), or it could be an exogenous shock that makes a large share of the capital worthless. The law of accumulation for capital is thus:

$$\begin{aligned} K_{t+1} &= (1-\delta)K_t + I_t, \text{ if } x_{t+1} = 0, \\ &= ((1-\delta)K_t + I_t)(1-b_k), \text{ if } x_{t+1} = 1, \end{aligned}$$

where x_{t+1} is a binomial variable which is 1 with probability p_t and 0 with probability $1-p_t$. A disaster does not affect productivity A_t.[6] The probability of disaster is assumed to vary over time, but to maintain tractability I assume in this section that it is *i.i.d.*: p_t, the probability of a disaster at time $t+1$, is drawn at time t from a constant cumulative distribution function F. Moreover, I assume that p_{t+1}, A_{t+1}, and x_{t+1} are independent.

This model has one endogenous state K and two exogenous states A and p, and there is one control variable C. The Bellman equation for the representative consumer is:

$$\begin{aligned} V(K,A,p) &= \max_{C,I} \left\{ (1-\beta) C^{1-\gamma} + \beta \left(E_{p',x',A'} V(K',A',p')^{1-\theta} \right)^{\frac{1-\gamma}{1-\theta}} \right\}^{\frac{1}{1-\gamma}} \\ s.t. &: C + I \leq AK, \\ K' &= ((1-\delta)K + I)(1-x'b_k). \end{aligned}$$

Define $W(K,A,p) = V(K,A,p)^{1-\gamma}$. Then we can guess and verify that W is of the form $W(K,A,p) = K^{1-\gamma} g(A,p)$, with[7]

$$g(A,p) = \max_i \left\{ \begin{array}{c} (1-\beta)(A-i)^{1-\gamma} \\ +\beta (1-\delta+i)^{1-\gamma} \left(1-p+p(1-b_k)^{1-\theta}\right)^{\frac{1-\gamma}{1-\theta}} \left(E_{p',A'} g(A',p')^{\frac{1-\theta}{1-\gamma}} \right)^{\frac{1-\gamma}{1-\theta}} \end{array} \right\},$$

where $i = \frac{I}{K}$ is the investment rate. The first-order condition with respect to i yields, after rearranging:

$$\left(\frac{A-i}{1-\delta+i} \right)^{-\gamma} = \frac{\beta}{1-\beta} \left(1-p+p(1-b_k)^{1-\theta}\right)^{\frac{1-\gamma}{1-\theta}} \left(E_{p',A'} g(A',p')^{\frac{1-\theta}{1-\gamma}} \right)^{\frac{1-\gamma}{1-\theta}}.$$

Given the assumption that p is *i.i.d.*, the expectation of g next period is independent of the current p. Hence, assuming that risk aversion $\theta \geq 1$, i is increasing in p if and only $\gamma > 1$ i.e. the intertemporal elasticity of substitution is less than unity. The intuition for this result is as follows:[8] if p goes up, the expected risk-adjusted return on capital $\left(1-p+p(1-b_k)^{1-\theta}\right)^{\frac{1-\gamma}{1-\theta}}$ goes down since there is a higher risk of disaster. However, the

[6] In an AK model, a permanent reduction in productivity would lead to a permanent reduction in the growth rate of the economy, since permanent shocks to A affect the *growth rate* permanently.

[7] Note that if $\gamma > 1$ the max needs to be transformed into a min.

[8] This intuition is similar to that spelled out in Weil (1989) in a consumption/savings example with exogenous returns.

effect of a change in the expected return on the consumption/savings choice depends on the value of the IES, because of offsetting wealth and substitution effects. If the IES is unity (i.e. utility is log), savings are unchanged and thus the savings or investment rate does not respond to a change in the probability of disaster. But if the IES is larger than unity, the substitution effect dominates, and i is decreasing in p (under the maintained assumption that $\theta \geq 1$). Hence, an increase in the probability of disaster leads initially, in this model, to a decrease in investment and an increase in consumption (since output is unchanged on impact). In the subsequent periods, the decrease in investment leads to a decrease in the capital stock and hence in output. This simple analytical example thus shows that a change in the perceived probability of disaster can lead to a decline in investment and output. While the preceding example is revealing,[9] a serious examination of the role of beliefs regarding disasters requires a quantitative model.

3.2 Quantitative model: a RBC model with disasters

This section introduces a real business cycle model with time-varying risk of disaster and study its quantitative implications. This model extends the simple example of the previous section in the following dimensions: (a) the probability of disaster is not *i.i.d.* but can be persistent; (b) the production function is neoclassical; (c) labor is elastically supplied; (d) disasters may affect total factor productivity as well as capital; (e) there are capital adjustment costs.

3.2.1 Model Setup

The representative consumer has preferences of the Epstein-Zin form, and the utility index incorporates leisure as well as consumption:

$$V_t = \left((1-\beta) u(C_t, N_t)^{1-\gamma} + \beta E_t \left(V_{t+1}^{1-\theta} \right)^{\frac{1-\gamma}{1-\theta}} \right)^{\frac{1}{1-\gamma}}, \tag{2}$$

where the per period felicity function $u(C, N)$ is assumed to have the following form:

$$u(C, N) = C^{\upsilon}(1-N)^{1-\upsilon}.$$

There is a representative firm, which produces output using a standard Cobb-Douglas production function:

$$Y_t = K_t^{\alpha} \left(z_t N_t \right)^{1-\alpha},$$

where z_t is total factor productivity (TFP), to be described below. The firm accumulates capital subject to adjustment costs:

$$\begin{aligned} K_{t+1} &= (1-\delta)K_t + \Phi\left(\frac{I_t}{K_t}\right) K_t, \text{ if } x_{t+1} = 0, \\ &= \left((1-\delta)K_t + \Phi\left(\frac{I_t}{K_t}\right) K_t \right)(1-b_k), \text{ if } x_{t+1} = 1, \end{aligned}$$

[9] This example is related to work by Epaulard and Pommeret (2003), Jones, Manuelli and Siu (2005a, 2005b), and to the earlier work of Obstfeld (1994).

where Φ is a concave function, and x_{t+1} is 1 if there is a disaster at time $t+1$ (with probability p_t) and 0 otherwise (probability $1-p_t$). The resource constraint is

$$C_t + I_t \leq Y_t.$$

Finally, we describe the shock processes. Total factor productivity is affected by the "business cycles" normal shocks ε_t as well as jumps (disasters). A disaster reduces TFP by a permanent amount b_z :

$$\begin{aligned} \log z_{t+1} &= \log z_t + \mu + \sigma\varepsilon_{t+1}, \text{ if } x_{t+1} = 0, \\ &= \log z_t + \mu + \sigma\varepsilon_{t+1} + \log(1-b_z), \text{ if } x_{t+1} = 1, \end{aligned}$$

where μ is the drift of TFP, and σ is the standard deviation of small "business cycles" shocks. Moreover, p_t follows a stationary Markov process with transition Q. I assume that $p_{t+1}, \varepsilon_{t+1}$, and x_{t+1} are independent conditional on p_t. This assumption requires that the occurrence of a disaster today does not affect the probability of a disaster tomorrow.[10]

3.2.2 Two analytical results

This model has three states: capital K, technology z and probability of disaster p, and two controls: consumption C and hours worked N. Denote $V(K,z,p)$ the value function, and define $W(K,z,p) = V(K,z,p)^{\frac{1}{1-\gamma}}$. The social planning problem can be formulated as:[11]

$$\begin{aligned} W(K,z,p) &= \max_{C,I,N} \left\{ (1-\beta)u(C,N)^{1-\gamma} + \beta \left(E_{p',z',x'} W(K',z',p')^{\frac{1-\theta}{1-\gamma}} \right)^{\frac{1-\gamma}{1-\theta}} \right\}, \\ s.t. &: C + I \leq z^{1-\alpha} K^{\alpha} N^{1-\alpha}, \\ K' &= \left((1-\delta)K + \Phi\left(\frac{I}{K}\right) K \right) (1 - x' b_k), \\ \log z' &= \log z + \mu + \sigma\varepsilon' + x' \log(1-b_z). \end{aligned}$$

A standard homogeneity argument implies that we can write $W(K,z,p) = z^{\upsilon(1-\gamma)} g(k,p)$, where $k = K/z$, and g solves an associated Bellman equation, which is then solved numerically using a mix of value and policy function iteration. Of course, a nonlinear method is crucial to analyze time-varying risk premia. From then on, it is relatively straightforward to compute the policy rules, to simulate the model, and to price assets using the stochastic discount factor: $M_{t,t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \left(\frac{V_{t+1}}{E_t\left(V_{t+1}^{1-\theta}\right)^{\frac{1}{1-\theta}}} \right)^{\gamma-\theta}$. Following Barro (2006), I will assume that government bonds are not risk-free but are subject to default risk during disasters. More precisely, if there is a disaster, then with probability q the bonds will default and the recovery rate will be r.

[10] This assumption could be wrong either way: a disaster today may indicate that the economy is entering a phase of low growth or is less resilient than thought; but on the other hand, if a disaster occurred today, and GDP fell by a large amount, it is unlikely that GDP will fall again by a large amount next year. Rather, historical evidence suggests that the economy is likely to grow above trend for a while (Gourio (2008)). Future work will study the implications of relaxing these independence assumptions.

[11] Here too, if $\gamma > 1$ the max needs to be transformed into a min.

Before turning to the quantitative analysis, it is useful to point out two simple analytical results which follow from the Bellman equation above:

Result 1: Assume that the probability of disaster is constant, and that $b_k = b_{tfp}$ i.e. TFP and capital fall by the same amount if there is a disaster.[12] Then, in a sample without disaster, the quantities implied by the model (consumption, investment, hours, output and capital) are the same as those implied by the model with no disasters ($p = 0$), but a different time discount factor $\beta^* = \beta(1 - p + p(1 - b_k)^{\upsilon(1-\theta)})^{\frac{1-\gamma}{1-\theta}}$.

Discussion of Result 1: Of course, asset prices will be different, and in particular the equity premium will be higher. Hence, this result is in the spirit of Tallarini (2000): fixing the asset pricing properties of a RBC model may not lead to any change in the quantity dynamics.[13] The observational equivalence is broken in a long enough sample since disasters must occur; or if one can trade assets contingent on disasters, since the prices would be different under the two models. This setup is not fully satisfactory in any case since the constant probability of disaster implies (nearly) constant risk premia.

Result 2: Assume still that $b_k = b_{tfp}$, but let now p vary over time. Then, in a sample without disaster, the quantities implied by the model are the same as those implied by a model with no disasters, but with stochastic discounting (i.e. β follows an exogenous stochastic process).

Discussion of result 2: Result 2 is interesting in light of the empirical literature which suggests that "preference shocks" may be important. Chari, Kehoe and McGrattan (2008) complain that these shocks lack microfoundations. My model is much "smaller" than these medium-scale models, but I conjecture that this equivalence should hold in larger versions.

3.2.3 Calibration

Parameters are listed in Table 1. (This calibration and quantitative results are still preliminary and can likely be improved.) Many parameters are fairly standard (see e.g. Cooley and Prescott (1995)). Risk aversion is 10, but note that this is the risk aversion over the consumption-hours bundle. Since hours does not vary much when there is a disaster, this utility index is about three times less volatile than consumption. The IES is set equal to 2. One crucial element is the probability and size of disaster. As in Barro (2006), I assume that $b_k = b_{tfp} = .43$ and the probability is $p = .017$ per year on average. The second crucial element is the persistence and volatility of movements in this probability of disaster. For now I follow Gourio (2008b) and assume that this change in probability of disaster is volatile and highly persistent.

[12] This assumption simplifies the analysis: the steady-state of the economy shifts due to a change in z, but the ratio of capital to productivity is unaffected by the disaster, i.e. the economy is in the same position relative to its steady-state after the disaster and before the disaster.

[13] Gabaix (2008) proves some results with a similar flavor.

Parameter	Greek Letter	Value
Capital share	α	.34
Depreciation rate	δ	.02
Adjustment cost curvature	η	.25
Trend growth of TFP	μ	.005
Discount factor	β	.992
IES	$1/\gamma$	2
Share of consumption in utility	υ	.3
Risk aversion	θ	10
Standard deviation of TFP shock	σ	.01
Size of disaster in TFP	b_{tfp}	.43
Size of disaster for capital	b_k	.43
Probability of disaster in low prob state	p_l	.0005
Probability of disaster in high prob state	p_h	.008
Probability of transition from p_l to p_h	π	.99

Table 1: Parameter values for the baseline (quarterly) model.

3.2.4 Quantitative Implications: first and second moments of asset returns, and business cycle facts

Table 2 reports some moments obtained from model simulations for a sample without disasters. The dynamics of quantities in response to a TFP shock are similar to those of a standard model without disasters.[14] The model is able to generate a relatively large equity premium of 4% per year. (In population, i.e. in samples which have disasters, the equity premium is 3.2%. The model generates a slightly negative term premium, consistent with the evidence for indexed bonds in the US and UK. However, the model does not generate enough volatility in equity returns. The volatility is higher in a sample with disasters (7.2%). Adding some financial or operating leverage, and possibly some wage rigidities may also help here, and so do steeper adjustment costs or a lower IES, or a more volatile probability of disaster.[15] The volatility of hours is low, because of the adjustment and the labor supply specification, which may be improved.

[14]Consumption, investment and employment are procyclical, and investment is the most volatile series. The T-bill rate is procyclical, as is the equity return and Tobin's q. The equity premium is acyclical. This model generate some positive autocorrelation of consumption growth, hence the dynamics of consumption are qualitatively similar to those which are studied in Bansal and Yaron (2004). This could in principle generate larger risk premia, however, as argued by Kaltenbrunner and Lochstoer (2008), this effect is not quantitatively very important if shocks are permanent and the IES is not small.

[15]Gourio (2008b) and Wachter (2008) show that in an endowment economy, the disaster model is able to generate volatile returns provided that the probability of disaster is volatile, the IES is not low, and there is substiantial leverage.

Parameter	Model	Data
Mean return on pure risk-free asset	-1.4%	NA
Mean return on short-term gov't bond	0.8%	0.9%
Mean real term premium	-0.5%	NA
Mean excess return on unlevered equity	4.0%	5.4%
Std dev. of return on short-term gov't bond	1.8%	1.5%
Std dev. of return on unlevered equity	1.6%	19.4%
$\sigma(\Delta \log C)/\sigma(\Delta \log Y)$	0.78	0.52
$\sigma(\Delta \log N)/\sigma(\Delta \log Y)$	0.25	0.90
$\sigma(\Delta \log I)/\sigma(\Delta \log Y)$	1.96	2.20

Table 2: Annual Moments implied by the model.

3.2.5 The effect of an increase in the probability of disaster

We can now perform the key experiment of an upward "shock" to the probability of disaster, which leads to an increase in risk premia. Figure 1 plots the (nonlinear) impulse response function to such a shock.[16] Investment falls, and consumption rises, as in the analytical example of section 3.1, since the elasticity of substitution is assumed to be greater than unity. Employment falls too, through an intertemporal substitution effect: the risk-adjusted return to savings is low and thus working today is less attractive. Hence, output drops because both employment and the capital stock fall, even though there is no change in current or future total factor productivity. This is one of the main result: this shock to the "perceived risk" leads to a recession. After impact, total resources available shrink, and so does consumption. These results are robust to changes in parameter values, except of course for the IES which crucially determines the sign of the responses. The size of adjustment costs affects the magnitude of the response of investment and hours.[17] Regarding asset prices, we see (Figure 2) that following the shock, and the onset of the recession, the risk premium on equity increases, the yield curve becomes upward-sloping, and the short rate falls. These are all patterns which are typical of a recession.

3.2.6 Asset prices lead the business cycle

Indeed, a very interesting feature of the model is its ability to replicate the stylized facts that the stock market and the term premium lead the business cycle, and the short rate leads negatively the business cycle. Figure 3 presents the facts by displaying the monthly cross-correlogram of industrial production and employment, with the T-bill rate, the term spread, the market return and the market excess return (blue line; 2-SE bands in black). As emphasized by Backus, Routledge and Zin (2008), standard models fail at replicating this important feature of the data: they imply that the correlation is essentially zero except on

[16] The figure plots the path implied by the model, starting in steady-state, if, at $t = 6$, the economy shifts from the low probability of disaster to the high probability of disaster. For clarity, there are no further shocks to the probability of disaster, no realized disaster, and no "normal shocks" ε.

[17] The model predicts some negative comovement between C and I, which is reminiscent of Barro and King (1984), but the quantitative significance of this point depends on the labor supply specification.

impact. This is also true in my model when the driving source of fluctuations is TFP shocks. But shocks to the probability of disasters generate a pattern similar to the data, as shown in Figure 3 which superposes (in red) the model implications to the data.[18].

The intuition is simple: a shift from high to low probability of disasters will increase the risk-free rate, and will also lead to a business cycle boom as shown in Figure 2. Similarly, the term spread will become positive, since long-term yields move less than short-term yields; and the risk premium falls, so the return on the stock market was high on average before the shift.

4 Changes in Risk-Taking, Reallocation, and Macroeconomic Implications

A change in the aggregate risk affects macroeconomic aggregates also by affecting the willingness to take on risk. Faced with an increase in the probability of an economic disaster, investors shift resources to technologies and projects which are less exposed to disasters. In doing so, they move the economy alongside a risk/return frontier, and pick project which are less risky but also have lower expected returns. As a result, the expected output of the economy falls, and so does productivity.

The first subsection presents a simple analytical model to illustrate this effect. Because this argument relies merely on changes in the aggregate risk and not especially on the disaster formulation, I use a model with a standard lognormal shock which may be more familiar. The second section presents some preliminary microeconomic evidence which supports the effect.

While the idea of risk-return trade-off is of course familiar from the portfolio choice literature, there has been little work examining (theoretically or empirically) the risk-return trade off in real investment, i.e. the choice of exposure to aggregate shocks.[19]

4.1 A two-period toy model

Consider the following two-period, partial equilibrium model. In the first period, agents decide how to allocate a fixed amount of capital K between two different technologies. In the second period, the economy is hit by an aggregate shock z, and output is produced by each of the two technologies.

For simplicity, I will assume that the aggregate shock z is lognormally distributed, $E \log z = \mu$, and $V \log z = \sigma^2$, and that the discount factor is $M(z) = e^{-\gamma z}$.[20] As a result, the log risk-free rate is $\log R_f = \gamma\mu - \frac{\gamma^2\sigma^2}{2}$.

The two technologies are summarized by the production functions:

$$Y_i(z, K_i) = z^{\lambda_i} e^{-\frac{\lambda_i^2\sigma^2}{2} - \lambda_i\mu} K_i^{\alpha},$$

[18]Note that the model is now calibrated to monthly data. For clarity, the figure is drawn assuming that the only shocks are shocks to the probability of disaster (i.e. there are not TFP shocks).

[19]See however Angeletos et al. (2007), Ramey and Ramey (1995), and Obstfeld (1994) for related ideas.

[20]Heuristically, assume that z is proportional to aggregate consumption, and γ is a measure of risk aversion.

for $i = 1, 2$. These technologies have decreasing returns ($\alpha < 1$) and have different exposures to the shock z if $\lambda_1 \neq \lambda_2$. Decreasing returns imply that the return on capital endogenous. High λ technologies are more risky, but the expected productivity is the same if the amount of capital invested in each technology is the same: $E_z[Y_i(z, K_i)] = K_i^\alpha$. Hence, if there is no risk, we would simply allocate capital equally across the two technologies. With aggregate risk however, agents will invest less in the more risky technology, up to the point where the expected *discounted* marginal products are equal; the expected marginal products will differ, with one technology have a higher expected return, to compensate for its higher risk.

Mathematically, assuming that capital depreciates at rate δ, each firm $i = 1, 2$ picks K_i to maximize expected discounted profit:

$$\max_{K_i} \left\{ E_z \left(M(z) \left(Y_i(z, K_i) + (1-\delta) K_i \right) \right) \right\} - K_i,$$

and simple computations show that this is equivalent to:

$$\max_{K_i} \left\{ \frac{K_i^\alpha e^{-\lambda_i \gamma \sigma^2} + (1-\delta) K_i}{R^f} - K_i \right\},$$

i.e. agents discount the future undepreciated capital stock at the risk-free rate and discount the future expected profits using a risk premium $e^{-\lambda_i \gamma \sigma^2}$. The first order condition is thus:

$$\alpha K_i^{\alpha - 1} e^{-\lambda_i \gamma \sigma^2} + 1 - \delta = R^f, \tag{3}$$

which is a standard user cost rule determining capital demand in sector i. The user cost incorporates a risk premium which depends on the risk of the technology λ_i.

To determine the aggregate effects of this effect, recall that $K = K_1 + K_2$ is the aggregate capital available at time 1. Equation (3) implies the following allocation of capital:

$$K_1 = \frac{K}{1 + \exp\left(\frac{(\lambda_1 - \lambda_2)\gamma\sigma^2}{1-\alpha} \right)}.$$

Note that if *either* $\lambda_1 = \lambda_2$ *or* $\gamma = 0$ *or* $\sigma = 0$, then we have $K_1 = K_2 = \frac{K}{2}$: capital is allocated equally so as to equate the expected marginal product of capital across sectors. Aggregate output in period 2 is then:

$$Y(z) = \sum_{i=1}^{2} Y_i(z) = \sum_{i=1}^{2} z^{\lambda_i} e^{-\frac{\lambda_i^2 \sigma^2}{2} - \lambda_i \mu} K_i^\alpha,$$

and expected output is $E_z(Y(z)) = \sum_{i=1}^{2} K_i^\alpha$. Denote by $Y^* = 2^{1-\alpha} K^\alpha$ the aggregate output if there is no risk, and define $TFP^* = \frac{Y}{K^\alpha} = 2^{1-\alpha}$ the aggregate total factor productivity without risk. A simple second order approximation reveals that, in general:

$$TFP \simeq TFP^* \times \left(1 - \frac{(\lambda_1 - \lambda_2)^2 \gamma^2 \alpha^2 \sigma^4}{1 - \alpha} \right),$$

which shows that TFP will be smaller in an economy with aggregate risk than in an economy with risk. The presence of risk introduces a "distortion" by creating an unequal allocation of capital which lowers aggregate TFP. The output Y is governed by a similar formula (since total capital is fixed in this example). A higher risk aversion γ or a higher variance of the shock σ will increase the distortion. Note that if $\lambda_1 = \lambda_2$, there is no distortion and $TFP = TFP^*$ (up to the third order): the distortion is created by the *difference* in risk exposures.

It remains to be seen, however, if this effect is large. The first-order effect is zero, because when risk is small, the misallocation is very small and marginal products are nearly equated. The next section gives some simple empirical evidence that supports this mechanism.[21]

4.2 Cross-sectional evidence

In support of the reallocation mechanism discussed in the previous section, I show that when the aggregate risk premium is higher, investment shifts relatively towards the firms which are less risky. Here riskiness is to be understood as exposure to aggregate shocks.

There are two empirical challenges in testing this mechanism. First, we need to identify the more risky firms. In this section, I will assume that riskiness is reflected in discount rates and hence market prices, so I will assume that firms with low Tobin's q (high book-to-market ratio, i.e. "value" firms) are the more risky. Second, we need to measure the aggregate risk premium, in an ex-ante sense. Here I will use two different methods. First and most simply, I will use the aggregate Tobin's q as an (inverse) measure of aggregate riskiness. The idea is that, if dividend growth is not too predictable (a robust empirical finding), variations in aggregate Tobin's q are mostly due to changes in discount rates. The second method is to run a relatively standard empirical model to forecast stock returns, and use the predicted values as a measure of the ex-ante risk premium. I follow a standard specification (e.g., Campbell and Yogo (2006)), and use the dividend yield, the term spread and the short rate to forecast excess equity return:[22]

$$rp_{t+1} \overset{def}{=} R^e_{t+1} - R^f_{t+1} = \alpha + \beta \frac{D_t}{P_t} + \gamma yieldspread_t + \delta short\ rate_t + \varepsilon_{t+1}. \tag{4}$$

Using standard Compustat data (1963-2004), I then create three groups of firms, based on their book-to-market ratio. (Each year, firms are rebalanced across groups according to their book-to-market ratio, as is usual in the empirical finance literature.) It is well known that the firms with low book-to-market (high Tobin's q) have higher investment.[23] I construct the investment rate $\frac{I_{it}}{K_{it}}$ of each group, and run the following regressions:

$$\frac{I_{it}}{K_{it}} = \alpha_i^1 + \beta_i^1 Q_t + u_{it}, \tag{5}$$

[21] The toy model can clearly be generalized to a DSGE two-sector economy, with the following ingredients: Epstein-Zin utility, adjustment costs to investment in each sector, and a time-varying risk premium through either time-varying risk of disaster or time-varying volatility. Solving this model numerically reveals that this result is robust, i.e. a higher risk leads to a reallocation towards the safest technology.

[22] These two measures are similar, because the price-dividend ratio is highly correlated with the aggregate market-to-book ratio, and the dividend yield is the most important variable that predicts returns in equation (4).

[23] While the $q-$theory does not work well at the firm level, it works better at the portfolio level (see e.g. Liu, Whited and Zhang (2007)).

	low B/M ($i = 1$)	medium B/M ($i = 2$)	high B/M ($i = 3$)
$100 \times \beta_i^1$	-2.56	-0.49	-0.58
t-stat	-4.31	-0.55	-0.91
$100 \times \beta_i^2$	13.15	-9.59	-4.65
t-stat	1.15	-1.24	-0.73
$100 \times \gamma_i^1$	-7.78	-0.71	1.94
t-stat	-3.96	-0.48	1.23
$100 \times \gamma_i^2$	93.05	6.11	-10.29
t-stat	1.89	0.31	-0.38

Table 3: Estimated regression coefficients and Newey-West adjusted SE, in a regression of investment rate on the aggregate Tobin's Q or the expected equity premium.

where Q_t is the aggregate Tobin's q; and

$$\frac{I_{it}}{K_{it}} = \alpha_i^2 + \beta_i^2 \widehat{rp}_{t+1} + u_{it}, \tag{6}$$

where $\widehat{rp}_{t+1}$ is the fitted value from the regression (i.e. the risk premium estimated given date t variables). The test is that β_i^1 should be higher, and β_i^2 would be lower, for firms with higher book-to-market. I also try to test if the share of investment done in the more risky firms falls when the equity premium is higher. To test this, I run the following regressions:

$$\frac{\frac{I_{it}}{K_{it}}}{\frac{I_t}{K_t}} = \delta_i^1 + \gamma_i^1 Q_t + u_{it}, \tag{7}$$

and

$$\frac{\frac{I_{it}}{K_{it}}}{\frac{I_t}{K_t}} = \delta_i^2 + \gamma_i^2 \widehat{rp}_{t+1} + u_{it}, \tag{8}$$

Table 3 presents the results, and Figure 4 plots these regression coefficients with the 2-standard error bands on each side. Clearly, the patterns are respected: $\beta^1_{\text{low b/m}} < \beta^1_{\text{high b/m}}$: the investment of firms with high B/M increases by more (falls by less) when the aggregate Q_t increases. Similarly, higher expected excess returns on equity reduces the investment in firm with high B/M by more, i.e. $\beta^2_{\text{low b/m}} > \beta^2_{\text{high b/m}}$ The share coefficients γ_i^1 and γ_i^2 behave similarly. However, the coefficients are not always statistically significant.

Of course, these results are just a starting point, but they show that the mechanism may well be working in the right direction, and that the effects may be quantitatively significant. These results are robust to introducing a lag between investment and asset prices (time-to-build) and to some changes in specification or data.