Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Idiosyncratic investment risk and business cycles

Jonathan E. Goldberg

2014-05

Idiosyncratic investment risk and business cycles

Jonathan Goldberg*

January 10, 2014

Abstract

I show that, due to imperfect risk sharing, aggregate shocks to uncertainty about idiosyncratic return on investment generate economic contractions with elevated risk premia and a decrease in the risk-free rate. I present a tractable real business cycle model in which firms experience idiosyncratic shocks, to which managers are at least partially exposed; the distribution of these shocks is time-varying and stochastic. I show that the path for aggregate quantities, the price of physical capital, and the equity premium are the same as in a model without idiosyncratic risk, but with time-preference shocks. That is, in response to an increase in idiosyncratic uncertainty, the response of these variables is the same as if there were no idiosyncratic uncertainty but managers were suddenly reluctant to invest. However, time-preference and idiosyncratic uncertainty shocks are not isomorphic: an increase in idiosyncratic uncertainty leads to greater demand for precautionary saving and hence a decrease in the risk-free rate; in contrast, an increase in impatience has the opposite effect. In addition, with an idiosyncratic uncertainty shock, investment in physical capital can remain low even after the stock market and firm profitability recover, because managers cannot fully transfer idiosyncratic risk to diversified investors. Thus, shocks to idiosyncratic investment risk can explain, qualitatively, the aftermath of financial panics – elevated risk premia, a sharp and persistent decrease in investment, and a decrease in the risk-free rate. In a calibration, an increase in idiosyncratic investment risk similar to that experienced during the Great Recession leads firms to invest as if their cost of capital were 10 percentage points higher than the cost of capital implied by financial markets, and to a large decrease in the real risk-free rate.

Keywords: Incomplete markets, idiosyncratic risk, business cycles, equity premium, risk-free rate

JEL codes: D52, E44, G11

*Federal Reserve Board. I thank Anna Orlik, Vasia Panousi, Skander Van den Heuvel, and Francisco Vazquez-Grande for helpful comments. Any views expressed here are those of the author and need not represent the views of the Federal Reserve Board or its staff. Contact: jonathan.goldberg@frb.gov

For a given firm, uncertainty about idiosyncratic returns varies over time. Moreover, across firms, this time variation in idiosyncratic uncertainty has an aggregate component. And because a firm's managers are at least partially exposed to the firm's idiosyncratic risks, a shock to idiosyncratic uncertainty can affect a firm's investment decisions and its managers' consumption and savings behavior.

In this paper, I study how aggregate shocks to idiosyncratic investment risk affect business cycles and risk premia. To do so, I develop a tractable real business cycle model in which managers face moral hazard and hence are exposed to firm-specific shocks. Each period, firms experience an idiosyncratic shock to their return on investment. In a given period, firm-level return shocks are independent and identically distributed across firms. However, the distribution of these idiosyncratic shocks is itself an aggregate shock. Thus, the model has two aggregate shocks: a standard labor-augmenting productivity shock that is common across firms; and an aggregate shock to idiosyncratic uncertainty.

When there is no aggregate uncertainty, comparing across steady-states, I show that an increase in idiosyncratic uncertainty leads to a decrease in aggregate capital, consumption, employment, and the risk-free rate; also, the wedge between the the expected return to physical capital and the risk-free rate increases. More generally, when idiosyncratic uncertainty and aggregate productivity are stochastic, I show that the path for aggregate quantities, the price of physical capital, and the equity premium are the same as in a model without idiosyncratic risk, but with a time-varying discount factor. That is, in response to an increase in idiosyncratic uncertainty, the response of these variables is the same as if there were no idiosyncratic uncertainty but managers were suddenly impatient. Intuitively, when idiosyncratic risk increases, the risk-adjusted return to investing in physical capital falls, because managers are risk averse and investing in physical capital requires bearing idiosyncratic risk. Thus, aggregate quantities and the equity premium behave as if there were no investment risk, but managers had become impatient and thus reluctant to invest.

These results are important because a number of recent papers have used time-preference shocks to explain asset-pricing puzzles (Albuquerque, Eichenbaum and Rebelo (2012)) and business-cycle dynamics (Hall (2013), Christiano, Eichenbaum and Rebelo (2011), Smets and Wouters (2003)). My results imply that if time-preference shocks in a relatively standard real business cycle model can explain the dynamics of aggregate quantities and the equity premium, then, in a model where managers are unable to pledge a fraction of firm profits, there exists a stochastic process for uncertainty about idiosyncratic return on investment that also explains these dynamics.

However, shocks to idiosyncratic uncertainty are not isomorphic to time-preference shocks: an increase in idiosyncratic uncertainty leads to a decrease in the risk-free rate, whereas the time-preference shock that generates the same path for aggregate quantities leads to an increase in the risk-free rate. The reason is that an increase in idiosyncratic uncertainty leads to a greater demand for precautionary saving, and hence lower returns on the risk-free asset, whereas an impatience shock implies higher expected returns on all assets, including risk-free bonds. Thus, in several papers that use time-preference shocks, a time-preference shock that causes a drop in the stock market or an economic contraction leads to an increase in the (real) risk-free rate, which limits the ability of time-preference shocks to explain the joint behavior of the risk-free rate, aggregate quantities and risk premia.

For example, Albuquerque, Eichenbaum and Rebelo (2012) show that a simple Lucas tree economy with time-preference shocks can explain several asset-pricing puzzles. In their model, a "bad" shock – one that leads to an increase in the equity premium – is an increase in impatience, which leads to higher expected returns for financial assets and correspondingly an increase in the risk-free rate. This is consistent with the positive correlation in their data between equity returns and their measure of the risk-free rate. However, this also suggests that a standard time-preference shock, at least in a Lucas tree economy or real business cycle model, cannot do a good job of explaining the qualitative behavior of financial variables during and after the 2008-2009 financial crisis – characterized by elevated risk premia and a *decrease* in the risk-free rate.

To explain financial and macroeconomic dynamics during and after the financial crisis, other papers, including Christiano, Eichenbaum and Rebelo (2011), have used a time-preference shock with the opposite sign: a decrease in impatience. In a model without nominal rigidities, following such a shock, the risk-free rate falls and investment increases. In contrast, with nominal rigidities and a zero-lower-bound on nominal rates, the real interest rate can increase sharply, leading to a decrease in investment, as in Christiano, Eichenbaum and Rebelo (2011). However, this is again at odds with the decrease in the risk-free rate during and after the financial crisis.

Instead, the uncertainty shock in my paper can explain the discount rate shock that Hall (2013) uses to explain employment dynamics after the financial crisis: Hall (2013) considers a discount-rate shock that increases the required return on risky investments even as the risk-free rate decreases. Thus, my paper addresses the difficulty of standard macroeconomic models to match counter-cyclical equity premia and the decrease in risk-free rates associated with financial panics, by demonstrating that a shock to idiosyncratic investment risk decreases the risk-free rate and affects aggregate quantities, the equity premium, and the price of capital as if there were an impatience shock.

Another difference between a time-preference shock and a shock to idiosyncratic uncertainty is that, with idiosyncratic uncertainty, there is a breakdown in the standard Q-theory result that the return on firm assets is equal to the return on a financial claim on firm assets. Instead, the return on investment is greater than, rather than equal to, the return on financial claims on firms. This wedge is required to compensate managers for bearing idiosyncratic risk. I show that this wedge increases in response to an idiosyncratic uncertainty shock. Thus, with a shock to idiosyncratic uncertainty, investment in physical capital may appear "too low" given stock-market valuations. Panousi and Papanikolaou (2012) find that when idiosyncratic risk rises, firm investment falls, and more so when managers own a larger fraction of the firm. Panousi and Papanikolaou (2012) also find that, during the financial crisis, firms with higher fractions of managerial ownership reduced investment (as a share of existing capital stock) by 6 percentage points more than firms with a more diversified shareholder base. Buera and Moll (2012) provide a simple measure of the ratio of the return on physical capital to the risk-free rate, and show that it increased during the financial crisis. Smets and Wouters (2007) and Galí, Smets and Wouters (2011) consider a "risk premium" shock that generates a wedge between the return on physical capital and the return on government bonds, and find that this shock is important for explaining short-run movements in output, employment and the Federal Funds rate. Galí, Smets and Wouters (2012) show that output and the labor market recovered more quickly after pre-1990s recessions than after the three most recent recessions, and attribute much of the difference to this "risk premium" shock.

This paper is closely related to the literature on disaster risk, especially Gourio (2012) and Gourio (2013). These papers model disasters as a series of shocks to aggregate productivity and capital quality. In Gourio (2012), there is a representative firm and financial markets are complete. In this setting, if disasters last exactly one period and involve equal-sized and permanent reductions in productivity and capital, the response of aggregate quantities to an increase in the risk of a disaster is the same as if there was no disaster risk, but agents became suddenly impatient; also, the risk-free rate falls. Thus, the response of aggregate quantities and the risk-free rate to an increased risk of disaster in his model is qualitatively similar to the response of these variables to a shock to idiosyncratic uncertainty in my model. However, our papers differ in a number of ways. First, disaster-risk shocks are different than shocks to idiosyncratic uncertainty; disaster risk concerns the level of (or uncertainty about) a coupled shock to aggregate productivity and aggregate capital quality, whereas idiosyncratic investment risk concerns only uncertainty about firm-level return on investment. Second, in Gourio (2012), markets are complete and Q-theory holds, whereas in my model, financial contracting is limited by moral hazard and without this limitation, the uncertainty shock would not matter. In addition, moral hazard gives rise to a wedge between the returns on physical capital and a financial claim on those returns; this wedge can be measured, as in Panousi and Papanikolaou (2012) and Buera and Moll (2012), to potentially distinguish between the two models. Third, Gourio's disaster shocks are motivated by catastrophic wars and natural disasters; these disasters are rarely observed and hence difficult to learn about; and, even over a fairly long sample, observed business cycles patterns and risk premia will be driven not only by changes in the probability of disaster, but also by the number of rare disasters that actually occur in the period. In contrast, panel data on stock returns, private business income and consumption provide information about the types of idiosyncratic risk studied here.

This paper is also closely related to the literature on idiosyncratic investment risk, including Angeletos (2007), Angeletos and Panousi (2009), Angeletos and Panousi (2011) and Panousi (2012). One key difference between my paper and this earlier literature on idiosyncratic investment risk is that these papers assume that idiosyncratic uncertainty and all other aggregate exogenous variables are deterministic, even if they are time-varying. In contrast, I allow idiosyncratic uncertainty and aggregate productivity to follow a stochastic process, permitting a characterization of risk premia and business cycle dynamics.

I calibrate the model to be consistent with estimates of the volatility of idiosyncratic consumption and idiosyncratic returns of public and private firms. In the calibration, following a 50 percent increase in the standard deviation of idiosyncratic return on investment, aggregate quantities and the equity premium respond as if there were a negative discount-factor shock (e.g., an "impatience" shock) of 70 basis points. However, the risk-free rate declines by 5 percentage points, relative to the change in the risk-free rate caused by such a time-preference shock. In addition, the investment wedge – the spread between the return on a firm's physical capital and the return on financial claims on the firm – increases from 5 percentage points to 10 percentage points. If there were no aggregate uncertainty, the increase in idiosyncratic investment risk would lead to a decrease in the steady-state risk-free rate from 1 percent to -3 percent. These results show that, due to imperfect risk sharing, shocks to idiosyncratic uncertainty of the size contemplated in Bloom (2009) and Gilchrist, Sim and Zakrajšek (2013) can lead to very large shocks to the real risk-free rate, of similar magnitude to the exogenous "shock to the natural rate of interest" studied in Eggertsson and Woodford (2003) and Christiano, Eichenbaum and Rebelo (2011).

1 Model

Overview. Time is discrete, indexed by $t \in \{0, ..., \infty\}$. There is a continuum of infinitely-lived firms, indexed by i, that produce consumption goods using capital goods. Each firm is run by a manager.

At the beginning of period t, manager i has an expected capital stock $k_{t-1}^{e,i}$. The manager then experiences an idiosyncratic capital-quality shock s_t^i, which is drawn according to cumulative distribution function P_{t-1}. The manager also learns the aggregate technology shock z_t and the uncertainty shock P_t. Next, the manager hires labor in a competitive market, produces, and pays d_t^i to creditors. Finally, the manager chooses how much to consume, c_t^i, an expected next-period capital stock, $k_t^{e,i}$, and a portfolio of state-contingent debts, d_{t+1}^i. The extent to which managers can offload idiosyncratic risk through the portfolio of state-contingent debts is limited by moral hazard.

Denote manager i's history of idiosyncratic shocks by $s^{i,t} = \{s_0^i, s_1^i, ..., s_t^i\}$, the history of aggregate shocks by $h^t = (P^t, z^t)$ and the history of idiosyncratic and aggregate shocks by $h^{i,t} = \left\{s^{i,t}, h^t\right\}$.

Technology. Each period, firms experience idiosyncratic capital-quality shocks: although manager i chooses expected capital $k_t^{e,i}$ in period t, the manager's actual capital stock in period $t+1$ is

$$k_{t+1}^i = s_{t+1}^i k_t^{e,i} \tag{1}$$

where $E_t[s_{t+1}^i] = 1$ and $s_{t+1}^i \in \mathbf{S} = \left[s^{min}, s^{max}\right]$ has a cumulative distribution function $P_t \in \mathbf{P}$. The corresponding density function p_t is continuous and positive over $\mathbf{S}$.

The manager's output in period $t+1$ is

$$y_{t+1}^i = F(k_{t+1}^i, z_{t+1} l_{t+1}^i)$$

where F is a neoclassical production technology and l_{t+1}^i is labor hired by manager i. The price of capital goods in period $t+1$ is p_{t+1}^K. The firm's period-$t+1$ assets are:

$$a_{t+1}^i = y_{t+1}^i - \omega_{t+1} l_{t+1}^i + (1-\delta) p_{t+1}^K k_{t+1}^i. \tag{2}$$

The shock s_{t+1}^i is independent across firms and across time. P_t and z_t follow a Markov process.

Financial markets. In period t, the manager-owner can trade a full set of Arrow-Debreu securities with period-$t+1$ payouts that depend on the history of aggregate and idiosyncratic shocks, $h^{i,t+1}$. Specifically, at history $h^{i,t}$, the manager sells a portfolio of Arrow-Debreu securities that represent a promise to pay d_{t+1}^i next period. The proceeds from this sale are $E_t\left[q_{t+1} d_{t+1}^i\right]$, where q_{t+1} is a state-price density. Because the capital-quality shocks are idiosyncratic, q_{t+1} will depend only on the history of aggregate shocks, h^{t+1}.

Although the promises are state-contingent, markets are incomplete because of moral hazard. In particular, at history $h^{i,t+1}$, manager i can abscond with $(1-\theta)$ share of the firm's period-$t+1$ assets. Thus,

$$d_{t+1}^i \leq \theta a_{t+1}^i.$$

I assume that $\theta \in (0, 1 - \frac{1}{s^{\max}})$. If θ were greater than $1 - \frac{1}{s^{\max}}$, then there would be complete risk sharing in general equilibrium and the model would reduce to a standard real-business cycle model. If θ were equal to zero, then in general equilibrium, no risk sharing would be possible.

Preferences. I assume that managers have Epstein-Zin preferences with constant elasticity of intertemporal substitution and constant relative risk aversion. That is, associated with a stochastic consumption stream $\{c_t^i\}_{t=0}^{\infty}$ is a stochastic utility stream $\{v_t^i\}_{t=0}^{\infty}$ that satisfies the following recursion:

$$v_t^i = U^{-1}\left[U(c_t^i) + \beta U\left(\mathbb{CE}_t\left[v_{t+1}^i\right]\right)\right] \tag{3}$$

where $\beta < 1$ and $\mathbb{CE}_t(v^i_{t+1}) = \Upsilon^{-1}\left(E_t\left[\Upsilon(v^i_{t+1})\right]\right)$ is the certainty-equivalent of v^i_{t+1} conditional on $h^{i,t}$. Υ and U are given by:

$$\Upsilon(c) = c^{1-\gamma} \text{ and } U(c) = c^{1-\frac{1}{\epsilon}}. \tag{4}$$

Note that $\gamma > 0$ denotes the coefficient of relative risk aversion and $\epsilon > 0$ denotes the elasticity of intertemporal substitution (EIS).[1]

Budgets. Manager i's budget constraint at state $h^{i,t}$ is:

$$c^i_t + p^K_t k^{e,i}_t \leq w^i_t + E_t\left[q_{t+1} d^i_{t+1}\right]$$

where

$$w^i_t = a^i_t - d^i_t.$$

Consumption and expected capital cannot be negative: $c^i_t > 0$ and $k^{e,i}_t > 0$.

Capital goods. Capital-goods firms participate in a perfectly competitive capital-goods market. In period t, capital-goods firm j purchases $\phi(\frac{I_t}{K_{t-1}})I^j_t$ consumption goods and transforms them into I^j_t capital goods, where ϕ is continuous and increasing and $\phi(\delta) = 1$. Aggregate capital satisfies:

$$K_t = (1-\delta)K_{t-1} + I_{t-1}$$

and, in equilibrium, the price of capital satisfies $p^K_t = \phi(\frac{I_t}{K_{t-1}})$.

In order to have expected capital goods of $k^{e,i}_t$, manager i must purchase $k^{e,i}_t$ capital goods in period t.

Workers and limited participation. There is a representative worker that does not participate in financial markets. The worker's preferences over consumption and labor are given by:

$$u_t = U^{-1}\left[U(c_t - \frac{\zeta}{1+v} l^{1+v}_t) + \beta U\left(\mathbb{CE}_t\left[u_{t+1}\right]\right)\right]$$

where Υ and U are given by (4). The mass of workers is normalized to $L = 1$.

Equilibrium. An equilibrium is:

1. a mapping of the aggregate history h^t into a state-price density q_t, wages ω_t, the price of capital goods p^K_t, aggregate capital K_t, aggregate consumption C_t, aggregate labor L_t, and aggregate debt D_t;

2. a mapping for each manager i from $h^{i,t}$ into consumption c^i_t, capital k^i_t, repayment d^i_t and labor-demand l^i_t; and

3. a mapping for the representative worker from ω_t into labor supply l_t

such that:

1. the plans $\left\{c^i_t, k^{e,i}_t, d^i_{t+1}, l^i_t\right\}^\infty_{t=0}$ maximize the utility of each entrepreneur, taking prices $\left\{q_t, \omega_t, p^K_t\right\}^\infty_{t=0}$ as given;

2. the plan l_t maximizes the utility of the representative worker, taking ω_t as given, and aggregate capital K_t is consistent with profit maximization by capital goods firms;

[1] In general, I assume that the EIS is greater than one, but I also characterize the dynamics if this is not the case.

3. financial, labor, capital goods and consumption goods markets clear; and
4. aggregate quantities are determined by individual policies (i.e., aggregate managerial consumption $C_t^M = \int c_t^i di$ and aggregate capital $K_t = \int k_t^{e,i} di$).

The initial condition of the economy is given by the distribution of capital goods k_0^i and debt d_0^i across firms and an initial aggregate state $h^0 = (P_0, z_0)$.

The risk-free rate, the return on equity, and the public-company model of financial markets. Because there exist a full set of Arrow-Debreu securities, it is possible to price any financial claim.

The risk-free rate between period t and $t+1$ is given by

$$R_t^{rf} = E_t \left[q_{t+1}\right]^{-1}$$

I define the return on equity as the return between period t and $t+1$ on a financial claim on aggregate firm assets. That is,

$$R_{t+1}^{equity} = \frac{\frac{Y_{t+1}-\omega_{t+1}L_{t+1}}{K_t} + (1-\delta)p_{t+1}^K}{E_t[q_{t+1}\left(\frac{Y_{t+1}-\omega_{t+1}L_{t+1}}{K_t} + (1-\delta)p_{t+1}^K\right)]}. \tag{5}$$

I define the aggregate return to investing as:

$$R_{t+1} = \frac{\frac{Y_{t+1}-\omega_{t+1}L_{t+1}}{K_t} + (1-\delta)p_{t+1}^K}{p_t^K} \tag{6}$$

Note that, due to idiosyncratic risk, the standard Q-theory result that $R_{t+1} = R_{t+1}^{equity}$ will not hold. Thus, (5) differs from the return on equity in production-based asset pricing papers such as Boldrin, Christiano and Fisher (2001), in which the price of a financial claim to aggregate firm assets is $p_t^K K_t$ and hence (5) is equal to (6). One contribution of the paper is to characterize how the investment wedge, $\frac{R_{t+1}}{R_{t+1}^{equity}} = E_t\left[q_{t+1}R_{t+1}\right]$, and the equity premium, $\frac{E_t\left[R_{t+1}^{equity}\right]}{R_t^{rf}}$, are related to idiosyncratic risk, parametrized by P_t.

Remark 1. An alternative way to model financial contracting is to envision the creation of publicly traded, limited-liability equity claims on all future cash flows of a firm, when contracting between the managers and the equity holders is subject to the same moral hazard problem as in the sequential-trading "entrepreneurial" setup above. Suppose that, in period 0, managers and investors meet and create a limited-liability publicly traded company. Each contributes wealth to create the company. The company is owned by the investors and signs a contract with the manager to provide a stream of consumption, given by c_t^i, and to make a stream of investments in physical capital, given by $k_t^{e,i}$, where c_t^i and $k_t^{e,i}$ depend on the history of aggregate and idiosyncratic shocks. In this model, c_t^i resembles managerial compensation. The managerial moral-hazard constraint is that the manager's lifetime utility v_{t+1}^i must be greater than or equal to the outside option of absconding with $(1-\theta)$ share of the firm's assets, a_{t+1}^i, and re-contracting with a new set of equity holders. This alternative model results in the same equilibrium policies, aggregate quantities and state price density as in the sequential-trading "entrepreneurial" setup above. Moreover, in general equilibrium, the return on the aggregate limited-liability equity claims is given by (5).

Remark 2. The return on equity is the return on an unlevered financial claim to aggregate firm assets. However, it is possible, as in Boldrin, Christiano and Fisher (2001), to study the return on a levered claim to aggregate firm assets:

$$R_{t+1}^{levered} = R_{t+1}^{equity} + \lambda \left(R_{t+1}^{equity} - R_t^{rf} \right)$$

where λ measures leverage. The expected excess return on the levered claim is given by:

$$\frac{E_t[R_{t+1}^{levered}]}{R_t^{rf}} = (1+\lambda)\frac{E_t[R_{t+1}^{equity}]}{R_t^{rf}} - \lambda$$

2 Equilibrium characterization

2.1 Partial equilibrium

In the model, the firm's end-of-period assets and labor demand are linear in the manager's capital, due to constant returns to scale in technology and the ability to adjust labor demand according to the realization of the idiosyncratic shock:

$$a_{t+1}^i = R_{t+1} p_t^K k_{t+1}^i \text{ and } l_{t+1}^i = l_{t+1} k_{t+1}^i$$

where $l_{t+1} = \arg\max_l \left(F(1, z_{t+1} l) - \omega_{t+1} l \right)$ and $R_{t+1} = \frac{1}{p_t^K} \left(F(1, z_{t+1} l_{t+1}) - \omega_{t+1} l_{t+1} + (1-\delta) p_{t+1}^K \right)$.

Thus, the firm's problem can be written recursively as:

$$V(w_t^i; t) = \max_{k_t^{e,i}, c_t^i, \{w_{t+1}^i\}} U^{-1} \left[U(c_t^i) + \beta U \left(\mathbb{CE}_t \left[V(w_{t+1}^i; t+1) \right] \right) \right] \tag{7}$$

subject to the budget constraint

$$E_t \left[q_{t+1} w_{t+1}^i \right] \leq w_t^i - c_t^i + E_t \left[q_{t+1} R_{t+1} - 1 \right] p_t^K k_t^{e,i}$$

and, for each s_{t+1}^i, the limited-enforcement constraints

$$w_{t+1}^i \geq (1-\theta) R_{t+1} p_t^K k_{t+1}^i \tag{8}$$

Below, I guess and verify that V is linear in the manager's wealth. Thus, conditional on aggregate history h^{t+1} and idiosyncratic history $h^{i,t}$, manager i's wealth w_{t+1}^i will equal some constant amount following every idiosyncratic shock s_{t+1}^i below a threshold s_t^{i*}, and for every state s_{t+1}^i greater than this threshold, wealth will be determined by (8). That is, w_{t+1}^i will be equal to the greater of a fixed amount or the minimum wealth level consistent with repayment. I denote the fixed amount by n_{t+1}^i, so that

$$w_{t+1}^i = \max \left\{ n_{t+1}^i, (1-\theta) R_{t+1} p_t^K k_{t+1}^i \right\}. \tag{9}$$

Using this intuition, firms' optimal decisions for given prices can be characterized.

Lemma 1. *Given prices, optimal consumption c_t^i, expected capital $k_t^{e,i}$, and minimum payoff n_{t+1}^i are linear in wealth w_t^i*

$$\begin{aligned} \frac{c_t^i}{w_t^i} &= \tilde{c}_t \\ \frac{p_t^K k_t^{e,i}}{w_t^i} &= (1-\tilde{c}_t)\kappa_t \\ \frac{n_{t+1}^i}{w_t^i} &= (1-\tilde{c}_t)\eta_{t+1} \end{aligned} \tag{10}$$

and the consumption-wealth ratio $\tilde{c}_t$ satisfies:

$$\frac{1}{\tilde{c}_t} = (\beta U(\rho_t))^{\epsilon} + 1 \tag{11}$$

where

$$\rho_t = \mathbb{CE}_t \left[\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}} \max\{\eta_{t+1}, (1-\theta) s_{t+1}^i R_{t+1} \kappa_t\} \right] \tag{12}$$

and

$$\{\eta_{t+1}, \kappa_t\} = \arg\max_{\{\eta\},\kappa} \mathbb{CE}_t \left[\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}} \max\{\eta, (1-\theta) s_{t+1}^i R_{t+1} \kappa\} \right] \tag{13}$$

subject to

$$E_t \left[q_{t+1} \max\{\eta, (1-\theta) s_{t+1}^i R_{t+1} \kappa\} \right] \leq 1 + E_t \left[q_{t+1} R_{t+1} - 1 \right] \kappa \tag{14}$$

Lemma 1 simplifies the manager's problem by transforming it into a canonical portfolio choice problem in which there is a single asset, with return

$$R_{t+1}^i \equiv \max\{\eta_{t+1}, (1-\theta) s_{t+1}^i R_{t+1} \kappa_t\} \tag{15}$$

where η_{t+1} and κ_t are given by (13). The return on the single asset reflects the optimal mix of investments in physical capital and financial assets, given prices and the limited enforcement constraints.

Correspondingly, as in a canonical portfolio choice problem with Epstein Zin preferences, the lifetime utility function V_t that solves the firm's problem is given by:

$$V_t^i = \tilde{c}_t^{\frac{1}{1-\epsilon}} w_t^i$$

Thus, ρ_t is the period-t risk-adjusted return to saving, where return to saving is measured in units of the marginal lifetime utility of wealth, $\frac{dV_{t+1}^i}{dw_{t+1}^i} = \tilde{c}_{t+1}^{\frac{1}{1-\epsilon}}$. That is,

$$\rho_t = \mathbb{CE}_t \left[\frac{dV_{t+1}^i}{dw_{t+1}^i} R_{t+1}^i \right]$$

Managers' consumption choice is described by the Euler equation (11); a higher risk-adjusted return ρ_t is consistent with lower consumption and greater investment if the elasticity of intertemporal substitution is greater than one.

To understand the manager's problem, consider manager i's pricing kernel m_{t+1}^i, where

$$m_{t+1}^i \equiv \frac{\partial v_t^i / \partial dc_{t+1}^i}{\partial v_t^i / \partial dc_t^i} = \beta \left(\frac{c_{t+1}^i}{c_t^i} \right)^{-\frac{1}{\epsilon}} \left(\frac{v_{t+1}^i}{\mathbb{CE}_t \left[v_{t+1}^i \right]} \right)^{\frac{1}{\epsilon} - \gamma}. \tag{16}$$

(16) depends only on the assumption of Epstein Zin preferences. The Euler equation (11) can be stated as the familiar asset-pricing condition:

$$E \left[m_{t+1}^i R_{t+1}^i \right] = 1 \tag{17}$$

If there are no limited-enforcement constraints (that is, $\theta = 1$), then equilibrium requires

$$m_{t+1}^i = q_{t+1}$$

and the model becomes a standard real-business cycle model. However, if $\theta < 1$, then

$$m^i_{t+1} < q_{t+1}$$

whenever the limited enforcement constraint (8) binds. That is, for $s^i_{t+1} > s^{i*}_t$, the marginal lifetime utility from an additional unit of consumption is low, relative to the market price of consumption, but reducing wealth and consumption at state s^i_{t+1} without violating the limited-enforcement constraint would require reducing $k^{e,i}_t$, thus affecting consumption at other histories $h^{i,t+1}$.

3 General equilibrium

3.1 Cross-sectional dynamics

Define manager i's share of total managerial consumption by $\frac{c^i_t}{C^M_t}$. The next result shows that risk sharing across managers is imperfect.

Lemma 2. *The (gross) growth rate of manager i's consumption share in period $t+1$ is given by:*

$$g^i_{t+1} \equiv \frac{c^i_{t+1}/C^M_{t+1}}{c^i_t/C^M_t} = \max\{\psi_t, (1-\theta)s^i_{t+1}\} \tag{18}$$

where $\psi_t < 1$ is the unique solution to

$$E[\max\{\psi_t, (1-\theta)s^i_{t+1}\}] = 1. \tag{19}$$

Lemma 2 implies that how idiosyncratic risk between period-t and period-$t+1$ is shared depends only on the distribution of period-$t+1$ idiosyncratic shocks and the extent of moral hazard, $1-\theta$. This result follows from market clearing and the linearity of each manager's investment and consumption in her wealth.

Later, I will characterize how macroeconomic and financial variables respond to a mean-preserving spread in g^i_{t+1}. A paper from the options literature, Rasmusen (2007), provides a definition of increased risk in s^i_{t+1} that is useful here.

Definition. The distribution $\widetilde{P}_t$ is pointwise riskier than P_t if s^i_{t+1} has the same mean under each distribution and if there exist s' and s'' with $s^{min} < s' < s'' < s^{max}$ such that

$$\tilde{p}_t \quad > \quad p_t \text{ for all } s^i_{t+1} \in [s^{min}, s'] \bigcup [s'', s^{max}]$$

and $\tilde{p}_t \leq p_t$ otherwise.

An increase in pointwise riskiness shifts probability mass from each point in the middle of the support to points at each extreme of the support. As the next result shows, an increase in pointwise riskiness in s^i_{t+1} is a sufficient condition for a mean-preserving spread in g^i_{t+1}.[2]

Lemma 3. *An increase in pointwise risk in P_t leads to a mean-preserving spread in the growth rate of manager i's consumption share, g^i_{t+1}, and a decrease in risk-sharing, ψ_t.*

[2] An increase in pointwise riskiness is a mean-preserving spread, but a mean-preserving spread is not necessarily an increase in pointwise riskiness. Note that a mean-preserving spread in s^i_{t+1} is not a sufficient condition for a mean-preserving spread in g^i_{t+1}: a mean-preserving spread in s^i_{t+1} that leaves the density function unchanged for all $s^i_{t+1} > s^*_t$ will have no impact on g^i_{t+1}.

According to Lemma 3, the growth rate of each manager's consumption share becomes riskier when idiosyncratic uncertainty increases. The deterioration of risk sharing has two components: the mechanical effect that P_t is riskier; as well as an endogenous, general-equilibrium effect, in the form of a strict decrease in ψ_t, the worst-case growth of manager i's consumption share. Intuitively, worst-case consumption-share growth decreases because, with greater uncertainty about idiosyncratic depreciation, managers seek to hedge all of their increased downside risk, but can only sell θ share of their increased upside risk.

3.2 Aggregate dynamics

Using the market-clearing condition $\kappa_t = 1$, one can write the idiosyncratic return to investment, defined in (15), as:

$$R^i_{t+1} = g^i_{t+1} R_{t+1} \tag{20}$$

The term g^i_{t+1} reflects idiosyncratic risk, given by (18). The term R_{t+1} is the aggregate return to investment, given by (6).

Profit maximization of consumption-goods firms and capital-goods firms, together with labor market clearing, implies

$$R_{t+1} = \frac{F_K(K_t, z_{t+1}\left(\frac{\omega_{t+1}}{\zeta}\right)^{\frac{1}{v}}) + (1-\delta)\phi(\frac{I_{t+1}}{K_t})}{\phi\left(\frac{I_t}{K_{t-1}}\right)}$$

where ω_{t+1} is the unique solution to

$$\omega_{t+1} = F_L(K_t, z_{t+1}\left(\frac{\omega_{t+1}}{\zeta}\right)^{\frac{1}{v}}) z_{t+1}. \tag{21}$$

These results permit further characterization of the dynamics of aggregate capital and consumption.

Proposition 4. *The path for aggregate capital, K_t, aggregate consumption, C_t, and the price of physical capital, p^K_t, are the same as in a model without investment risk ($\theta = 1$), but with a different discount factor $\bar{\beta}_t$ that follows a stochastic process given by:*

$$\bar{\beta}_t = \beta U\left(\mathbb{CE}_t[g^i_{t+1}]\right) \tag{22}$$

where g^i_{t+1} is the growth rate of manager i's share of aggregate managerial consumption, given by (18). A mean-preserving spread in g^i_{t+1} is associated with a decrease in the equivalent discount factor $\bar{\beta}_t$ if and only if $\epsilon > 1$.

Proof. Combining (11) and (12), we obtain the Euler equation:

$$\tilde{c}_t = \frac{1}{\left(\beta U\left(\mathbb{CE}_t\left(\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}} R^i_{t+1}\right)\right)\right)^{\epsilon} + 1}.$$

where R^i_{t+1} is given by (20). Because the growth rate of the idiosyncratic consumption share, g^i_{t+1}, is independent of h_{t+1}, we can re-write this as:

$$\tilde{c}_t = \frac{1}{\left(\beta U\left(\mathbb{CE}_t\left(g^i_{t+1}\right)\right) U\left(\mathbb{CE}_t\left(\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}} R_{t+1}\right)\right)\right)^{\epsilon} + 1}$$

That is, conditional on the aggregate return to investing R_{t+1}, the consumption share of wealth $\tilde{c}_t$ will be the same as in a model without investment risk but with a discount factor given by (22). Therefore, the equilibrium path for aggregate capital and consumption will be the same in both models and hence the path for the aggregate return to investing will be the same as well. □

If there is no aggregate uncertainty about idiosyncratic risk (e.g., if $P_t = P$ for all t), then aggregate quantities are the same as in a model without investment risk ($\theta = 1$), but with an adjusted discount factor. The direction of adjustment depends on the EIS: if managers are willing to substitute across time ($\epsilon > 1$) , a mean-preserving spread in g^i_{t+1} makes the managers more reluctant to invest. The intuition is that an increase in idiosyncratic risk leads to a decrease in the risk-adjusted return to saving; in response, managers will consume more and save less if they are willing to substitute across time.

Similarly, if there is aggregate uncertainty about idiosyncratic risk (e.g., if P_t follows a stochastic process), then aggregate quantities are the same as in a model without investment risk ($\theta = 1$), but with time-preference shocks.[3]

3.3 Asset pricing and the investment wedge

It is not the case that shocks to idiosyncratic uncertainty are isomorphic to time-preference shocks. In particular, the prices of financial assets are different than in the model without investment risk, but with time-preference shocks given by (22), even though the paths for aggregate quantities and the price of physical capital are the same.

One way to see this is to examine how the equilibrium state-price density q_{t+1} is related to the state-price density in a model without investment risk, but with time-preference shocks.

Lemma 5. *At time t, the state-price density q_{t+1} is the same, up to a pre-determined scalar, as in a model without investment risk ($\theta = 1$) but with stochastic discount factor $\bar{\beta}_t$ given by (22). That is, if q^{β}_{t+1} is the state-price density in the model without investment risk and with time-preference shocks given by (22), then*

$$\frac{q_{t+1}}{q^{\beta}_{t+1}} = \frac{\mathbb{CE}_t[\frac{g^i_{t+1}}{\psi_t}]^{\gamma}}{\mathbb{CE}_t[g^i_{t+1}]} > 1. \tag{23}$$

Proof. For any idiosyncratic history $h^{i,t+1}$ for which the limited-enforcement constraint (8) is not binding, we have $q_{t+1} = m^i_{t+1}$. Hence, substituting into (16), we have

$$q_{t+1} = \beta\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\psi_t R_{t+1}(1-\tilde{c}_t)\right)^{-\frac{1}{\epsilon}}\left(\frac{\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}}\psi_t R_{t+1}}{\mathbb{CE}_t\left[\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}} g^i_{t+1} R_{t+1}\right]}\right)^{\frac{1}{\epsilon}-\gamma}$$

[3] Proposition 4 describes how the adjusted discount factor defined by (22) responds to a mean-preserving spread in g^i_{t+1}, an endogenous variable. One way to connect this result to an exogenous shock is to recall, from Lemma 2, that a pointwise increase in risk in the exogenous distribution P_t is a sufficient condition for a mean-preserving spread in g^i_{t+1}.

Because the growth rate of the idiosyncratic consumption share, g^i_{t+1}, is independent of h_{t+1}, we can write

$$q_{t+1} = \beta\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\psi_t R_{t+1}(1-\tilde{c}_t)\right)^{-\frac{1}{\epsilon}}\left(\frac{\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}}\psi_t R_{t+1}}{\mathbb{CE}_t\left[\tilde{c}_{t+1}^{\frac{1}{1-\epsilon}}R_{t+1}\right]\mathbb{CE}_t\left(g^i_{t+1}\right)}\right)^{\frac{1}{\epsilon}-\gamma}$$

which, together with (22), implies (23). Finally, (19) and Lemma 2 imply $\frac{q_{t+1}}{q^\beta_{t+1}} > 1$. $\square$

This result has immediate implications for the investment wedge, the risk-free rate and the equity premium.

3.3.1 Investment wedge

In a model without idiosyncratic investment risk ($\theta = 1$), the return on investment, (6), would equal the return on equity, (5). Correspondingly, we would have:

$$E_t[q_{t+1}R_{t+1}] = 1. \tag{24}$$

This no-arbitrage condition is the standard Q-theory result that financial-market prices q_{t+1} can price the return to investing in physical capital.

However, with idiosyncratic investment risk, (24) does not hold. Because investing in physical capital involves idiosyncratic risk, managers need to compensated to bear these risks. This will take the form of an investment wedge.

Proposition 6. *The investment wedge,* $\frac{R_{t+1}}{R^{equity}_{t+1}} = E_t[q_{t+1}R_{t+1}]$, *is greater than one.*

Proposition 6 is a corollary of Lemma 5. To see this, substitute from (23) to write the investment wedge as:

$$E[q_{t+1}R_{t+1}] = \frac{q_{t+1}}{q^\beta_{t+1}}E_t[q^\beta_{t+1}R_{t+1}] = \frac{q_{t+1}}{q^\beta_{t+1}} \tag{25}$$

From Proposition 4, we have that the equilibrium path for R_{t+1} is the same as in the model without investment-risk, but with time-preference shocks given by (22). Note that in the model without investment risk, the state-price density that prices financial assets q^β_{t+1} will price the return to investing in physical capital, R_{t+1}; that is, $E_t[q^\beta_{t+1}R_{t+1}] = 1$. The result then follows from Lemma 5.

3.3.2 The equity premium and the risk-free rate

Additional corollaries of Lemma 5 concern the risk-free rate and the equity premium:

Proposition 7. *The equity premium,* $\frac{E_t\left[R^{equity}_{t+1}\right]}{R^{rf}_t}$, *is the same as in the model without investment risk, but with stochastic discount factor* $\bar{\beta}_t$ *given by (22).*

Proof. From (5) and (6), the equity premium is given by:

$$\begin{aligned} \frac{E_t\left[R_{t+1}^{equity}\right]}{R_t^{rf}} &= \frac{E_t\left[R_{t+1}\right]}{E_t\left[q_{t+1}R_{t+1}\right]}E_t[q_{t+1}] \\ &= \frac{E_t\left[R_{t+1}\right]}{E_t\left[q_{t+1}^{\beta}R_{t+1}\right]}E_t[q_{t+1}^{\beta}] \end{aligned} \tag{26}$$

where the second equality follows from (23). (26) is the equity premium in the model without investment risk, but with stochastic discount factor $\bar{\beta}_t$. □

Proposition 8. *The risk-free rate, $R_t^{rf} = E_t[q_{t+1}]^{-1}$, will be lower than in the model without investment risk, but with stochastic discount factor $\bar{\beta}_t$ given by (22).*

Proof. Substitute from (23) to write the risk-free rate as:

$$R_t^{rf} = E_t\left[q_{t+1}^{\beta}\right]^{-1}\frac{q_{t+1}^{\beta}}{q_{t+1}} \tag{27}$$

In the model without investment risk, but with time-preference shocks, the risk-free rate is given by $E_t\left[q_{t+1}^{\beta}\right]^{-1}$. And, from Lemma 5, $q_{t+1} > q_{t+1}^{\beta}$. □

Thus, although aggregate quantities and the equity premium are same as in a model without investment risk, but with time-preference shocks given by (22), the risk-free rate is lower.

3.3.3 Effects of an increase in idiosyncratic uncertainty

Next, I will characterize how the investment wedge and the state-price density change in response to an increase in idiosyncratic uncertainty.

Proposition 9. *There exists a $\overline{\gamma} > 1$ such that, if $\gamma < \overline{\gamma}$, an increase in pointwise risk in P_t leads to a strict increase in the investment wedge $\frac{R_{t+1}}{R_{t+1}^{equity}} = \frac{q_{t+1}}{q_{t+1}^{\beta}}$; and a strict decrease in the ratio of the risk-free rate to the risk-free rate that would obtain in the model without investment risk, but with stochastic discount factor $\bar{\beta}_t$ given by (22).*

Proposition 9 highlights the different forces affecting asset pricing in a general equilibrium environment with idiosyncratic risk. As is well known from Hadar and Seo (1990) and Gollier (1995), in a standard partial-equilibrium portfolio problem with one safe asset and one risky asset, a mean-preserving spread in the return on the risky asset does not, in general, imply that a risk-averse agent will reduce her allocation to the risky asset. Instead, with constant relative risk aversion, Hadar and Seo (1990) show that a mean-preserving spread leads to a decrease in the allocation to that asset if $\gamma < 1$. Correspondingly, in my model, if $\gamma > 1$, it is possible to construct a mean-preserving spread in s_{t+1}^i such that the investment wedge decreases. In particular, any mean-preserving spread in s_{t+1}^i that leaves the density unchanged for all $s_{t+1}^i < s_t^*$ will result in a strict decrease in the investment wedge. However, with an increase in pointwise risk in s_{t+1}^i, there is always an endogenous increase in "downside risk," in the form of a strict decrease in ψ_t. This guarantees that if relative risk aversion is not too large, an increase in pointwise risk in s^i_{t+1} will lead to an increase in the investment wedge.[4]

It also possible to define a stricter notion of an increase in risk such that the investment wedge increases with risk, for every γ. In particular, recall that a pointwise increase in risk transfers probability mass from the middle of the distribution, with $s^i_{t+1} \in (s', s'')$, to the lower and upper parts of the distribution, with $s^i_{t+1} \in [s^{min}, s'] \bigcup [s'', s^{max}]$. Any pointwise increase in risk with $s'' < s^*_t$ leads to a strict increase in the investment wedge.[5]

3.4 Deterministic steady state

If there is no aggregate uncertainty, we can characterize the steady state of the economy:

Proposition 10. *Suppose that*

$$\epsilon > \underline{\epsilon} \equiv \left[1 + \frac{\log \beta}{\log\left(\mathbb{CE}_t[g^i_{t+1}]\right)}\right]^{-1}. \tag{28}$$

Then there exist unique steady-state values K^*, C^*, L^*, I^*, Y^*. A mean-preserving spread in g^i_{t+1} is associated with a decrease in K^*, C^*, L^*, I^*, Y^* if and only if $\epsilon > 1$.

Proposition 10 follows from the result that aggregate quantities are the same as in an economy without investment risk, but with a different discount factor given by (22). If $\epsilon > \underline{\epsilon}$, the equivalent discount factor is less than one and a unique steady state exists. Moreover, steady-state aggregate capital, consumption and employment decrease to a lower steady state following an increase in idiosyncratic risk if and only if $\epsilon > 1$. To see this, note that aggregate quantities respond to an increase in idiosyncratic risk as if managers were more impatient, if and only if $\epsilon > 1$.

We can also further characterize the risk-free rate. If $\epsilon > \underline{\epsilon}$, there is a unique steady-state value for the risk-free rate, given by:

$$R^{rf*} = \frac{1}{\beta}\psi^\gamma_t \mathbb{CE}_t[g^i_{t+1}]^{\frac{1}{\epsilon}-\gamma} < \frac{1}{\beta} \tag{29}$$

Proposition 8 implies the inequality in (29). The analogous result to Proposition 9 is that if relative risk aversion γ is not too high, an increase in pointwise risk in P_t leads to a strict decrease in the risk-free rate. However, whereas the threshold in Proposition 9 was $\overline{\gamma} > 1$, here it is $\overline{\gamma} > \frac{1}{\epsilon}$. Similarly, any pointwise increase in risk with $s'' < s^*_t$ leads to a strict decrease in the risk-free rate.

4 Quantitative implications of idiosyncratic investment risk

The calibration approach here is to characterize the financial and macroeconomic response to a shock to idiosyncratic uncertainty, while making as few as possible parametric assumptions.

[4] Of course, for the same reason, an increase in pointwise risk in s^i_{t+1} will lead to an increase in the ratio of the risk-free rate to the risk-free rate that would obtain in the model with time-preference shocks and the same path for aggregate quantities.

[5] To see this, note that such a pointwise increase in risk leads to an increase in $\mathbb{CE}_t[\gamma_t^{-1} g^i_{t+1}]^\gamma$, the numerator of the investment wedge (23), and a decrease in the denominator, $\mathbb{CE}_t[g^i_{t+1}]$.

Specifically, for several $P_t \in \mathbf{P}$, I calculate the distribution of idiosyncratic consumption growth, g^i_{t+1}, and the investment wedge, $\frac{R_{t+1}}{R^{equity}_{t+1}}$. I also calculate, for several $P_t \in \mathbf{P}$, the discount factor $\bar{\beta}_t$ such that aggregate quantities and the equity premium in the model are the same as in a model without idiosyncratic investment risk, but with discount factor $\bar{\beta}_t$. In doing so, I do not need additional assumptions about the technologies for producing consumption and capital goods. Also, I will not have to specify the Markov process for the technology shock z_t and the idiosyncratic risk shock P_t. That is, the results presented are such that, for any given Markov process for $P_t \in \mathbf{P}$, one can calculate the Markov process for $\bar{\beta}_t$. Although this approach offers only a partial characterization of the dynamics of the economy, the results will be consistent with, for example: an AR(1) process for the (log) standard deviation of idiosyncratic shocks (as in Gilchrist, Sim and Zakrajšek (2013)); or an AR(1) process for the *growth* rate of the standard deviation of idiosyncratic shocks (which would give rise to "quantity-equivalent" time-preference shocks similar to those in Albuquerque, Eichenbaum and Rebelo (2012)). Also, previous theoretical results will be informative about how the calibration results would vary with different parametric assumptions: for example, the investment wedge does not depend on the EIS, and the "quantity-equivalent" discount factor is decreasing in the EIS.

In addition, I calculate the marginal product of capital and the risk-free rate that would prevail in the steady state of the model when there is no aggregate uncertainty.

4.1 Parameter choice

I assume that s^i_{t+1} follows a Pareto distribution with standard deviation $\sigma_t \in (0, \infty)$. As before, I assume $E_t[s^i_{t+1}] = 1$. Thus, the tail parameter α_t satisfies

$$\alpha_t = 1 + \sqrt{1 + \sigma_t^{-2}} > 2$$

and the cumulative distribution function P_t is given by

$$P_t(s^i_{t+1}) = 1 - \left(1 - \alpha_t^{-1}\right)^{\alpha_t} \left(s^i_{t+1}\right)^{-\alpha_t}$$

for any $s^i_{t+1} \geq 1 - \alpha_t^{-1}$. Below, I repeat the analysis under the alternative assumption that s^i_{t+1} follows a log-normal distribution.

In the literature on idiosyncratic investment risk, Panousi (2012) uses a constant value of $\sigma = 0.3$, while Roussanov (2010) sets $\sigma = 0.45$. Using panel data on private-business income, DeBacker et al. (2012) also suggest $\sigma = 0.45$. One can also use idiosyncratic stock-market returns to calibrate σ_t, consistent with the public-company implementation of financial markets discussed in Section 2. Goyal and Santa-Clara (2003) find that the average volatility of individual stock returns is 16 percent per month, suggestive of annual volatility of individual stock returns between 50 and 60 percent. Thus, below I consider $\sigma_t \in [0.15, 0.6]$. The corresponding range for the Pareto tail parameter is $\alpha_t \in [2.9, 7.7]$.

To calibrate the risk-sharing technology (parametrized by θ in this model), the literature on idiosyncratic investment risk typically looks to data on the volatility of idiosyncratic consumption growth. Of course, disaggregated consumption data is subject to a variety of limitations: there is significant measurement error; households in the Consumer Expenditure Survey (CEX) are followed for only a short period of time; and idiosyncratic consumption may be driven by deterministic factors such as age and aggregate consumption. To deal with these issues, Blundell, Pistaferri and Preston (2008) impute consumption using the Panel

Study of Income Dynamics and the CEX, and regress imputed log annual consumption on year and year-of-birth dummies and a range of family characteristics. The first difference of these residuals corresponds, in my model, to the log of the consumption-share growth rate, $\log(g_{t+1}^i)$. Using the results in Blundell, Pistaferri and Preston (2008), a reasonable estimate of the cross-sectional standard deviation of this first difference is 13 percent per annum.[6] Blundell, Pistaferri and Preston (2008) study consumption dynamics for the entire population. In this paper, the focus is on investors and managers. Jacobs and Wang (2004) estimate the standard deviation of idiosyncratic consumption growth across all households and also limiting the sample only to asset holders. Averaging across time, they find that the mean standard deviation across all households is roughly similar to the mean standard deviation across asset holders. However, controlling for age and education, Jacobs and Wang find that the standard deviation of idiosyncratic consumption growth across asset holders is about 50 percent larger than the standard deviation across all households.

A number of calibration exercises make explicit or implicit choices about the volatility of the idiosyncratic consumption share. De Santis (2007) assumes that the cross-sectional standard deviation of $\log(g_{t+1}^i)$ has a mean of 10 percent. In Panousi (2012), the standard deviation of idiosyncratic consumption growth, in the steady state, is 7 percent. As a benchmark, I choose a value of $\theta = 0.3$, which results in a cross-sectional standard deviation of the log idiosyncratic consumption share growth rate equal to 8 percent if $\sigma_t = 0.3$.

For the elasticity of intertemporal substitution, I choose $\epsilon = 2$. Gourio (2012) chooses the same value. An EIS greater than one is required for an increase in idiosyncratic uncertainty to generate an economic contraction. I also consider how a lower EIS effects the results. I set the discount factor β equal to 0.95 and the coefficient of relative risk aversion γ equal to 4.

4.2 Effects of idiosyncratic uncertainty on financial and macroeconomic variables

Table 1 shows how different levels of idiosyncratic uncertainty affect risk sharing, macroeconomic dynamics and risk premia. The first column corresponds to the benchmark case, with $\sigma_t = 0.3$. The second and third columns correspond to a 50 percent and a 100 percent increase in the standard deviation of idiosyncratic shocks, as in Bloom (2009). The final column represents a 50 percent decrease in the standard deviation of idiosyncratic shocks.

Relative to the benchmark case, a 50 percent increase in σ_t leads to a roughly proportional increase in the volatility of idiosyncratic consumption. In the benchmark case, the lower bound of idiosyncratic consumption-share growth, ψ_t, is fairly high: following the worst possible shock, which corresponds to a decrease in capital of about 20 percent, a manager's consumption share falls only 2 percent.[7] With a 50 percent increase in σ_t, a manager's consumption share falls twice as much (that is, about 4 percent) following a bad shock.

The macroeconomic and financial effects of idiosyncratic risk are summarized in the next rows. With a 50 percent increase in σ_t, the response of aggregate quantities and the

[6] Blundell, Pistaferri and Preston (2008) use data from 1980 to 1992. In Table 4 of their paper, they show variance and first-order autocovariance of the residuals from regression described here. To find a first-order measure of the variance that is not contaminated by imputation and other errors, they suggest subtracting twice the absolute value of the first-order auto-covariance from the variance. Due to data availability, this is possible only for 8 years. The average of this adjusted variance measure over these 8 years is 0.13^2.

[7] The worst possible shock s_{t+1}^i corresponds to $1 - \alpha_t^{-1}$.

equity premium are the same as in a model without investment risk, but with a discount-factor shock of approximately 70 basis points. That is, aggregate quantities and the equity premium respond as if there were no investment risk, but the discount factor decreased from 0.945 to 0.938. At the same time, the investment wedge – the spread between the return on a firm's physical capital and the return on financial claims on the firm – increases from 5 percentage points to 10 percentage points. That is, in the benchmark case, firms invest as if their cost of capital were 5 percentage points higher than it actually is; with a 50 percent increase in idiosyncratic uncertainty, that wedge increases to 10 percentage points. Equivalently, the risk-free rate following the uncertainty shock will be about 5 percentage points lower than in a model without investment risk, but with a discount-factor shock that generates the same response for aggregate quantities.

The final row shows the steady-state risk-free rate that would obtain if there were no aggregate uncertainty. In the benchmark case, the risk-free rate would be about 1 percentage point per year. With a 50 percent increase in σ_t, the risk-free rate would be -3 percentage points per year. These results show that, due to imperfect risk sharing, shocks to idiosyncratic uncertainty of the size contemplated in Bloom (2009) can lead to very large shocks to the real risk-free rate, of similar magnitude to the exogenous "shock to the natural rate of interest" studied in Eggertsson and Woodford (2003) and Christiano, Eichenbaum and Rebelo (2011).

If the EIS were different, the distribution of idiosyncratic consumption-share growth, g^i_{t+1}, and the investment wedge would not change. Thus, the only variables reported in Table 1 that would change are the "quantity-equivalent" discount factor, $\bar{\beta}_t$, and the steady-state risk-free rate that would obtain if there were no aggregate uncertainty. From (22) and (27), one observes that, all else equal, $\bar{\beta}_t$ is decreasing in the EIS and that R^{rf*} is increasing in the EIS. As shown in Table 2, the size of the quantity-equivalent discount-factor shock corresponding to a 50 percent increase in idiosyncratic uncertainty would be somewhat smaller if the EIS were 1.5, at 50 basis points, rather than 70 basis points. The steady-state risk-free rate at the baseline level of idiosyncratic uncertainty would be 20 basis points lower, and a 50 percent increase in idiosyncratic risk would result in a slightly larger drop in the steady-state risk-free rate.

Unlike a change in the EIS, a change in relative risk aversion would affect the investment wedge, as shown in Table 3. If relative risk aversion were halved, the investment wedge would be 3 percentage points at the baseline level of idiosyncratic uncertainty, rather 5 percentage points. Moreover, a 50 percent increase in idiosyncratic uncertainty would lead to an investment wedge of 7 percentage points, rather than 10 percentage points. Similarly, the effects of idiosyncratic uncertainty on the "quantity-equivalent" discount factor and the steady-state risk-free rate would also be somewhat smaller if relative risk aversion were halved.

Finally, Table 4 shows how the results would be affected by assuming that P_t is log-normal, rather than Pareto.[8] At the baseline level of idiosyncratic uncertainty, assuming that P_t is log-normal implies that idiosyncratic consumption volatility is somewhat lower, and the lower-bound on consumption-share growth is a bit higher. Correspondingly, at the baseline level of idiosyncratic uncertainty, if P_t is log-normal, the quantity-equivalent discount factor and the steady-state risk-free rate are slightly higher, and the investment wedge is slightly lower.

However, at higher levels of idiosyncratic uncertainty, the effects of idiosyncratic uncer-

[8]That is, in Table 4, I assume: $\log s^i_{t+1}$ is normally distributed; $Var(s^i_{t+1}) = \sigma^2_t$; and $E_t[s^i_{t+1}] = 1$.

tainty are much larger if P_t is log-normal than if P_t is Pareto. Following a shift from $\sigma_t = 0.3$ to $\sigma_t = 0.45$, the volatility of idiosyncratic consumption-share growth increases 7 percentage points, and the lower-bound on the consumption-share growth rate falls 3 percentage points. Correspondingly, the quantity-equivalent discount-factor shock and the increase in the investment wedge are larger in magnitude than if P_t is Pareto.

5 Conclusion

In this paper, I have argued that shocks to uncertainty about idiosyncratic return on investment can explain the aftermath of financial crisis – elevated risk premia, a sharp and persistent decrease in investment, and a decrease in the risk-free rate. More specifically, aggregate quantities and the equity premium respond to an increase in idiosyncratic uncertainty as if there were no idiosyncratic investment risk and instead managers experienced an "impatience" time-preference shock. However, unlike an impatience shock, an increase in idiosyncratic uncertainty leads to a decrease in the risk-free rate.

Thus, a shock to idiosyncratic uncertainty is similar to the "risk premium" shock that Smets and Wouters (2007) , Galí, Smets and Wouters (2011) and Galí, Smets and Wouters (2012) find is important for explaining short-run movements in output, employment and interest rates, as well as the depth of and slow recovery from the Great Recession. In these papers, the "risk premium" shock makes it unattractive to hold physical capital. In order for the risk-free rate to fall meaningfully following such a shock, one needs the supply of financial assets to be somewhat inelastic. In Smets and Wouters (2003), this is the case because the only financial asset is a government bond, and government spending is exogenous. In my paper, following a shock to uncertainty about idiosyncratic rate of return, the supply of financial assets is somewhat inelastic because the way to create financial assets is by deploying more physical capital. Thus, as each manager seeks to shift from investing in her own physical capital to investing in diversified claims on other managers' physical capital, the other managers are trying to do the same.

References

Albuquerque, Rui, Martin S Eichenbaum, and Sergio Rebelo. 2012. "Valuation risk and asset pricing." National Bureau of Economic Research.

Angeletos, George-Marios. 2007. "Uninsured idiosyncratic investment risk and aggregate saving." *Review of Economic Dynamics*, 10(1): 1–30.

Angeletos, George-Marios, and Vasia Panousi. 2009. "Revisiting the supply side effects of government spending." *Journal of Monetary Economics*, 56(2): 137–153.

Angeletos, George-Marios, and Vasia Panousi. 2011. "Financial integration, entrepreneurial risk and global dynamics." *Journal of Economic Theory*, 146(3): 863–896.

Bloom, Nicholas. 2009. "The impact of uncertainty shocks." *Econometrica*, 77(3): 623–685.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption inequality and partial insurance." *American Economic Review*, 98(5): 1887–1921.

Boldrin, Michele, Lawrence J Christiano, and Jonas DM Fisher. 2001. "Habit persistence, asset returns, and the business cycle." *American Economic Review*, 91(1): 149–166.

Buera, Francisco J, and Benjamin Moll. 2012. "Aggregate implications of a credit crunch." National Bureau of Economic Research.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo. 2011. "When Is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119(1): 78–121.

DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. 2012. "The properties of income risk in privately held businesses." Mimeo, Federal Reserve Board.

De Santis, Massimiliano. 2007. "Individual consumption risk and the welfare cost of business cycles." *American Economic Review*, 97(4).

Eggertsson, Gauti B, and Michael Woodford. 2003. "Zero bound on interest rates and optimal monetary policy." *Brookings Papers on Economic Activity*, 2003(1): 139–233.

Galí, Jordi, Frank Smets, and Rafael Wouters. 2011. "Unemployment in an estimated new keynesian model." National Bureau of Economic Research.

Galí, Jordi, Frank Smets, and Rafael Wouters. 2012. "Slow recoveries: a structural interpretation." *Journal of Money, Credit and Banking*, 44(2): 9–30.

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajšek. 2013. "Uncertainty, Financial Frictions, and Irreversible Investment." Mimeo, Boston University.

Gollier, Christian. 1995. "The comparative statics of changes in risk revisited." *Journal of Economic Theory*, 66(2): 522–535.

Gourio, François. 2012. "Disaster Risk and Business Cycles." *American Economic Review*, 102(6): 2734–2766.

Gourio, François. 2013. "Credit Risk and Disaster Risk." *American Economic Journal: Macroeconomics*, 5(3): 1–34.

Goyal, Amit, and Pedro Santa-Clara. 2003. "Idiosyncratic risk matters!" *The Journal of Finance*, 58(3): 975–1008.

Hadar, Josef, and Tae Kun Seo. 1990. "The effects of shifts in a return distribution on optimal portfolios." *International Economic Review*, 31(3): 721–736.

Hall, Robert E. 2013. "High Discounts and High Unemployment." Mimeo, Stanford University.

Jacobs, Kris, and Kevin Q Wang. 2004. "Idiosyncratic consumption risk and the cross section of asset returns." *The Journal of Finance*, 59(5): 2211–2252.

Panousi, Vasia. 2012. "Capital taxation with entrepreneurial risk." Mimeo, Federal Reserve Board.

Panousi, Vasia, and Dimitris Papanikolaou. 2012. "Investment, idiosyncratic risk, and ownership." *The Journal of Finance*, 67(3): 1113–1148.

Rasmusen, Eric. 2007. "When Does Extra Risk Strictly Increase an Option's Value?" *Review of Financial Studies*, 20(5): 1647–1667.

Roussanov, Nikolai. 2010. "Diversification and its discontents: Idiosyncratic and entrepreneurial risk in the quest for social status." *The Journal of Finance*, 65(5): 1755–1788.

Smets, Frank, and Raf Wouters. 2003. "An estimated dynamic stochastic general equilibrium model of the euro area." *Journal of the European Economic Association*, 1(5): 1123–1175.

Smets, Frank, and Raf Wouters. 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American Economic Review*, 97(3): 586–606.

Appendix: Proofs omitted from the text

*Proof of Lemma **1***

Define

$$B(x_t^i) = \max_{k_t^{e,i},\{w^i_{,t+1}\}} \mathbb{CE}_t\left[V(w^i_{t+1}; t+1)\right] \tag{30}$$

subject to

$$E_t\left[q_{t+1}w^i_{t+1}\right] \leq x_t^i + E_t\left[q_{t+1}R_{t+1}-1\right]p_t^K k_t^{e,i}$$

and (8). Using (30), we can re-write (7) as

$$V(w_t^i;t) = \max_{x_t^i} U^{-1}\left[U(w_t^i - x_t^i) + \beta U\left(B(x_t^i;t\right)\right] \tag{31}$$

Conjecture the following solution:

$$\begin{aligned}
V(w_t^i;t) &= \xi_t w_t^i && (32)\\
\frac{c_t^i}{w_t^i} &= \tilde{c}_t && (33)\\
\frac{p_t^K k_t^{e,i}}{w_t^i} &= (1-\tilde{c}_t)\kappa_t\\
\frac{n^i_{t+1}}{w_t^i} &= (1-\tilde{c}_t)\eta_{t+1}
\end{aligned}$$

where ξ_t is an endogenous random variable.

Substituting from (9) and (32) into (30), one obtains:

$$B(x_t^i) = \max_{k_t^{e,i},\{w^i_{,t+1}\}} \mathbb{CE}_t\left[\xi_{t+1}\max\left\{n^i_{t+1}, (1-\theta)R_{t+1}s^i_{t+1}p_t^K k_t^{e,i}\right\}\right] \tag{34}$$

subject to

$$E_t\left[q_{t+1}\max\left\{n^i_{t+1}, (1-\theta)R_{t+1}s^i_{t+1}p_t^K k_t^{e,i}\right\}\right] \leq x_t^i + E_t\left[q_{t+1}R_{t+1}-1\right]p_t^K k_t^{e,i} \tag{35}$$

Dividing the first-order condition with respect to $k_t^{e,i}$ by the first-order condition with respect to n^i_{t+1}, one obtains:

$$\frac{n^i_{t+1}}{p_t^K k_t^{e,i}} = \left(\left(\frac{(1-\theta)E_t\left[q_{t+1}R_{t+1}s^i_{t+1}\mathbf{1}\{s^i_{t+1} > s^{i*}_{t+1}\}\right] - E_t\left[q_{t+1}R_{t+1}-1\right]}{(1-\theta)^{1-\gamma}E_t\left[\left(\xi_{t+1}R_{t+1}s^i_{t+1}\right)^{1-\gamma}\mathbf{1}\{s^i_{t+1} > s^{i*}_{t+1}\}\right]}\right)\frac{\xi^{1-\gamma}_{t+1}}{q_{t+1}}\right)^{\frac{1}{\gamma}}$$

where s^{i*}_{t+1} is defined by $n^i_{t+1} = (1-\theta)R_{t+1}s^{i*}_{t+1}p_t^K k_t^{e,i}$. This, together with the linearity of (35), implies that the solutions n_t^i and k_t^i to (34)-(35) are linear in x_t^i and that $B(x_t^i;t)$ is linear in x_t^i. Then, from (31), one obtains that x_t^i and c_t^i are linear in w_t^i, consistent with (33).

Next, using the envelope condition of (31), one obtains:

$$\frac{dV_t^i}{dw_t^i} = \left(\frac{c_t^i}{V_t^i}\right)^{-\frac{1}{\epsilon}} = \xi_t$$

Substituting from (32) and (33), we have

$$\xi_t = \tilde{c}_t^{\frac{1}{1-\epsilon}}$$

and thus

$$B(x_t^i) = \rho_t x_t^i. \tag{36}$$

Substituting (36) into (31) and taking the first-order condition with respect to x_t^i, one obtains (11).♢

Proof of Lemma 2

From Lemma 1, each manager's investment and consumption decisions are linear in her wealth. Thus, we have that $g_{t+1}^i = \frac{R_{t+1}^i}{R_{t+1}}$. Market clearing requires that $\kappa_t = 1$: all consumption goods that are not consumed are used as inputs to create capital goods. Financial market clearing implies that, for each h^{t+1},

$$\int_{s^{min}}^{s^{max}} R_{t+1}^i dP_t(s_{t+1}^i) = R_{t+1}.$$

Dividing both sides by R_{t+1} and denoting $\psi_t = \frac{\eta_{t+1}}{R_{t+1}}$, one obtains (19). Note that $\psi_t \in ((1-\theta)s^{\min}, 1)$ if $s^{\max} > \frac{1}{1-\theta}$ and that $\psi_t = 1$ otherwise. ♢

Proof of Lemma 3

Suppose that $\widetilde{P}_t$ is pointwise riskier than P_t. Note that $\int_{s^{min}}^{s^{max}} \max\{\psi_t, (1-\theta)s_{t+1}^i\} d\widetilde{P}_t(s_{t+1}^i) > 1$. Hence, if $\widetilde{P}_t$ is the c.d.f. of s_{t+1}^i, then ψ_t no longer solves (19). Instead, there is a unique $\tilde{\psi}_t < \psi_t$ that solves (19).

Denote the cumulative distribution function of g_{t+1}^i by $G(g_{t+1}^i; P_t)$, where g_{t+1}^i is defined by (18) and (19) and the c.d.f. of s_{t+1}^i is P_t. Then there exists a $\bar{g} \in [\psi_t, (1-\theta)s^{\max})$ such that $G(g_{t+1}^i; \widetilde{P}_t) > G(g_{t+1}^i; P_t)$ if $g_{t+1}^i \in [\tilde{\psi}_t, \bar{g})$ and $G(g_{t+1}^i; \widetilde{P}_t) \leq G(g_{t+1}^i; P_t)$ if $g_{t+1}^i \geq \bar{g}$. In addition, $G(g_{t+1}^i; \widetilde{P}_t) = G(g_{t+1}^i; P_t) = 0$ if $g_{t+1}^i < \tilde{\psi}_t$. ♢

Proof of Proposition 9

Let $h(P_t)$ denote $\mathbb{CE}_t[\frac{g_{t+1}^i}{\psi_t}]^\gamma \mathbb{CE}_t[g_{t+1}^i]^{-1}$, where the expectations are taken with respect to P_t and where g_{t+1}^i and ψ_t are determined using (18) and (19) conditional on P_t. Thus, from Lemma 5, $h(P_t) = \frac{q_{t+1}}{q_{t+1}^\beta}$. Suppose that $\tilde{P}_t$ is pointwise riskier than P_t. From Lemma 3, a shift from P_t to $\tilde{P}_t$ induces a mean-preserving spread in g_{t+1}^i and a strict decrease in ψ_t. Finally, note that $\lim_{\gamma \to 1} h(P_t) - h(\tilde{P}_t) = \psi_t^{-1} - \tilde{\psi}_t^{-1} < 0$ and that $h(P_t) - h(\tilde{P}_t) < 0$ if $\gamma < 1$.♢

Table 1: Effects of shock to idiosyncratic uncertainty

		$\sigma_t = 0.3$	$\sigma_t = 0.45$	$\sigma_t = 0.6$	$\sigma_t = 0.15$
Volatility of idiosyncratic consumption	std. dev. of $\log(g^i_{t+1})$	8.3	13.0	16.7	2.7
Lower bound on consumption-share growth	$\log(\psi_t)$	-2.0	-4.2	-6.4	-0.3
"Quantity-equivalent" discount factor	$\bar{\beta}_t$	0.945	0.938	0.930	0.949
Investment wedge	$\frac{R_{t+1}}{R^{equity}_{t+1}} - 1$	5.1	9.7	13.9	1.0
Steady-state risk-free rate if no aggregate uncertainty	$R^{rf*} - 1$	0.7	-2.8	-5.7	4.3

Note: All variables reported in percentage points, except for the "quantity-equivalent" discount factor. See text for details.

Table 2: Effects of shock to idiosyncratic uncertainty, if $\epsilon = 1.5$

		$\sigma_t = 0.3$	$\sigma_t = 0.45$	$\sigma_t = 0.6$	$\sigma_t = 0.15$
"Quantity-equivalent" discount factor	$\bar{\beta}_t$	0.947	0.942	0.937	0.950
Steady-state risk-free rate if no aggregate uncertainty	$R^{rf*} - 1$	0.5	-3.2	-6.3	4.3

Note: All variables reported in percentage points, except for the "quantity-equivalent" discount factor. See text for details.

Table 3: Effects of shock to idiosyncratic uncertainty, if $\gamma = 2$

		$\sigma_t = 0.3$	$\sigma_t = 0.45$	$\sigma_t = 0.6$	$\sigma_t = 0.15$
"Quantity-equivalent" discount factor	$\bar{\beta}_t$	0.947	0.941	0.935	0.950
Investment wedge	$\frac{R_{t+1}}{R_{t+1}^{equity}} - 1$	3.4	6.9	10.2	0.6
Steady-state risk-free rate if no aggregate uncertainty	$R^{rf*} - 1$	2.2	-0.6	-3.0	4.7

Note: All variables reported in percentage points, except for the "quantity-equivalent" discount factor. See text for details.

Table 4: Effects of shock to idiosyncratic uncertainty if s^i_{t+1} is log-normal

		$\sigma_t = 0.3$	$\sigma_t = 0.45$	$\sigma_t = 0.6$	$\sigma_t = 0.15$
Volatility of idiosyncratic consumption	std. dev. of $\log(g^i_{t+1})$	5.5	12.0	18.5	0.5
Lower bound on consumption-share growth	$\log(\psi_t)$	-1.5	-4.9	-9.4	-0.0
"Quantity-equivalent" discount factor	$\bar{\beta}_t$	0.948	0.939	0.926	0.950
Investment wedge	$\frac{R_{t+1}}{R^{equity}_{t+1}} - 1$	4.5	13.6	24.3	0.1
Steady-state risk-free rate if no aggregate uncertainty	$R^{rf*} - 1$	1.0	-6.3	-13.1	5.1

Note: All variables reported in percentage points, except for the "quantity-equivalent" discount factor. See text for details.