Solving Heterogeneous-Agent Models with Financial Frictions: A Continuous-Time Approach.

September 6, 2014

Notes prepared for Yuliy's lectures at "Princeton Initiative: Macro, Money and Finance," based on work with Markus.[1]

The goal of this lecture is to

1. develop techniques of solving heterogeneous-agent economies with financial frictions in continuous time and

2. address, through model elements, the concepts related to financial stability.

In particular, we will build models that can help us think about (1) undercapitalized sectors, (2) endogenous risk, (3) tail risk, (4) asset illiquidity, (5) endogenous leverage, (6) crisis probability, (7) inefficiencies of financial crises and (8) the effects of policies.

Technically, an equilibrium is defined as a map from any history of macro shocks to the current state of the economy, described by asset prices, asset allocation, and the agents' actions (production decisions, asset trades, etc). Such a map is an equilibrium if

1. all agents behave to maximize utility and

2. markets clear.

[1] I'd like to thank Ji Huang and Greg Phelan for helpful comments regarding computation.

The technical goal of this lecture is to translate these two sets of conditions into an equilibrium characterization.

Conceptually, we will replicate two important results from the linearized versions of classic models of Bernanke, Gertler and Gilchrist (1999) and Kiyotaki and Moore (1997), that (1) temporary macro shocks can have a *persistent* effect on economic activity by making borrowers "undercapitalized" and (2) price movements *amplify* shocks. We will also be able to take advantage of the tractability that continuous time offers and study a host of new properties of fully solved equilibria. In particular, we will observe that:

1. Endogenous risk is not the same at all times: it stays hidden in normal times, but it materializes in crisis times. Thus, the dynamics of an economy with financial frictions are highly nonlinear. Endogenous risk is tail risk. Endogenous risk depends on the illiquidity of assets, and it affects the severity of crises.

2. The leverage of borrowers, who may become undercapitalized, is endogenous. It responds to the magnitude of fundamental (exogenous) macro shocks and the level of financial innovations that enable better risk management. Leverage responds to a much lesser extent to the presence of endogenous tail risk. Equilibrium leverage in normal times is a key determinant of the probability of crises.

Below, I start first with a particularly simple model to illustrate how equilibrium conditions - utility maximization and market clearing - translate into an equilibrium characterization. This simple model trivializes most of the issues we are after, e.g. the model has no price effects or endogenous risk. We do get some interesting takeaways, such as that risk premia spike up in crises. After establishing the conceptual framework for what an equilibrium is, we move on to tackle more complex models.

A Simple Model.

This model is a variation of that from Basak and Cuoco (1998). The economy has a risky asset in positive net supply and a risk-free asset in zero net supply. There are two types of agents - experts and households. Only experts can hold the risky asset - households can only lend to experts at the risk-free rate r_t, determined endogenously in equilibrium. The friction is that experts can finance their holdings of the risky asset only through debt - by selling short the risk-free asset to households. That is, experts cannot issue equity. We assume that all agents are small, and behave as price-takers. That is, unlike in microstructure models with noise traders, agents have no price impact.

Risky Asset and Returns. Here is a model of production, which is slightly more general than what we need at this point, but which can be used throughout this set of notes. Consider a productive asset (capital) in the amount k_t, which evolves according to

$$\frac{dk_t}{k_t} = (\Phi(\iota_t) - \delta)\,dt + \sigma\,dZ_t, \tag{1}$$

where ι_t is reinvestment rate per unit of capital and $\Phi(\iota_t)$ is an investment function with adjustment costs, such that $\Phi(0) = 0$, $\Phi' > 0$ and $\Phi'' \leq 0$. Thus, in the absence of investment, capital simply depreciates at rate δ. The concavity of Φ reflects decreasing returns to scale, and for negative values of ι, corresponds to *technological illiquidity* - the marginal cost of capital depends on the rate of investment/disinvestment.

Net of investment, capital generates consumption output at the rate of

$$(a - \iota_t)k_t\,dt,$$

where a is a productivity parameter.

Suppose that the price per unit of capital q_t follows a Brownian law of motion of the form

$$\frac{dq_t}{q_t} = \mu_t^q\,dt + \sigma_t^q\,dZ_t, \tag{2}$$

which, of course, is endogenous in equilibrium. Then, using Ito's lemma, an investment in capital generates capital gains at rate

$$\frac{d(k_t q_t)}{k_t q_t} = (\Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q)\,dt + (\sigma + \sigma_t^q)\,dZ_t.$$

Then capital earns the return of

$$dr_t^k = \underbrace{\frac{a - \iota_t}{q_t}\,dt}_{\text{dividend yield}} + \underbrace{(\Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q)\,dt + (\sigma + \sigma_t^q)\,dZ_t}_{\frac{d(k_t q_t)}{k_t q_t}\text{, the capital gains rate}}. \tag{3}$$

Thus, generally a part of the risk from holding capital is fundamental, $\sigma\, dZ_t$, and a part is endogenous, $\sigma_t^q\, dZ_t$.

Note that the rate internal investment ι_t does not affect the risk of capital. The optimal investment rate that maximizes the expected return satisfies the first-order condition

$$\Phi'(\iota_t) = \frac{1}{q_t}. \tag{4}$$

Net Worth, Consumption and Preferences. Denote the aggregate amount of capital by K_t, so that the aggregate net worth in the economy is $q_t K_t$. If N_t is the aggregate net worth of experts, then the aggregate net worth of households is $q_t K_t - N_t$.

For *tractability*, all agents are assumed to have logarithmic utility with discount rate ρ, of the form

$$E\left[\int_0^\infty e^{-\rho t} \log c_t\, dt\right],$$

where c_t is consumption at time t. Logarithmic utility has two convenient properties, which help reduce the number of equations that characterize equilibrium. First, for agents with log utility

$$\text{consumption} = \rho \cdot \text{net worth} \tag{5}$$

that is, they always consume a fixed fraction of wealth regardless of the risk-free rate or risky investment opportunities. Second, the allocation of wealth between the risky and the risk-free asset is characterized by the equation

$$\text{volatility of wealth} = \text{Sharpe ratio of risky investment}, \tag{6}$$

where the volatility of wealth is measured in percent.[2]

We use equations (5) and (6) to formalize equilibrium conditions, and characterize equilibrium.

Definition. *An equilibrium is a map from histories of macro shocks* $\{Z_s,\ s \leq t\}$ *to the price of capital* q_t*, risk-free rate* r_t*, as well as asset holdings and consumption choices of all agents, such that*

[2]For example, if the annual volatility of S&P 500 is 15% and the risk premium is 3% (so that the Sharpe ratio is 3%/15% = 0.2), then a log utility agent wants to hold a portfolio with volatility 0.2 = 20%. This corresponds to a weight of 1.33 on S&P 500, and -0.33 on the risk-free asset.

1. *agents behave to maximize utility and*

2. *markets clear*

To find an equilibrium, we need to write down equations that processes q_t, r_t, etc. have to satisfy, and from those, characterize how these processes evolve with the realizations of shocks Z. It is convenient to express this relationship using a state variable. Here the relevant state variable, which describes the distribution of wealth, is the fraction of wealth owned by the experts,

$$\eta_t = \frac{N_t}{q_t K_t} \in [0, 1].$$

When η_t drops, experts become more constrained, and so small values of η_t correspond to a crisis regime.

So, how can we solve for the equilibrium? In two steps! First, we use the equilibrium conditions, i.e. utility and maximization and market clearing, to write down restrictions q_t and r_t need to satisfy. In this simple model, we will be able to express q_t and r_t as functions of η_t. Second, we need to derive the law of motion of η_t.

Step 1: The Equilibrium Conditions. First, from condition (5), the aggregate consumption of all agents is $\rho q_t K_t$, and aggregate output is $(a - \iota(q))K_t$, where investment ι is an increasing function of q defined by (4). From market clearing for consumption goods, these must be equal, and so

$$\rho q_t = a - \iota(q_t) \tag{7}$$

is the equilibrium price of the risky capital. The aggregate consumption of experts must be $\rho N_t = \rho \eta_t q_t K_t$, and the aggregate consumption of households is $\rho(1 - \eta_t) q_t K_t$. Condition (7) alone leads to a *constant* value of the price q.

Example. Suppose the investment function takes the form

$$\Phi(\iota) = \frac{\log(\kappa\iota + 1)}{\kappa},$$

where κ is the adjustment cost parameter. Then $\Phi'(0) = 1$. Higher κ makes function Φ more concave, and as $\kappa \to 0$, $\Phi(\iota) \to \iota$, a fully elastic investment function with on adjustment costs. Then optimal investment rate is $\iota = (q-1)/\kappa$, and the market-clearing condition (7) leads to the price of

$$q = \frac{a + 1/\kappa}{r + 1/\kappa}.$$

The price converges to 1 as $\kappa \to 0$, i.e. the investment technology is fully elastic. The price q converges to a/r as $\kappa \to \infty$.

Second, we can use condition (6) for experts to figure out the equilibrium risk-free rate. We look at the return on risky and risk-free assets to compute the Sharpe ratio of risky investments. We look at balance sheets of experts to compute the volatility of their wealth. Then we use equation (6) to get the risk-free rate.

Because q is constant, risky asset earns the return of

$$dr_t^k = \underbrace{\frac{a-\iota}{q}\,dt}_{\rho,\text{ dividend yield}} + \underbrace{(\Phi(\iota)-\delta)\,dt + \sigma\,dZ_t}_{\text{capital gains rate}},$$

and risk-free asset earns r_t so the Sharpe ratio of risky investment is

$$\frac{\rho + \Phi(\iota) - \delta - r_t}{\sigma}.$$

Because experts must hold all the risky capital in the economy, with value $q_t K_t$ (households cannot hold capital), and absorb risk through net worth N_t, the volatility of their net worth is

$$\frac{q_t K_t}{N_t}\sigma = \frac{\sigma}{\eta_t}.$$

Using (6),

$$\frac{\sigma}{\eta_t} = \frac{\rho + \Phi(\iota) - \delta - r_t}{\sigma} \quad \Rightarrow \quad r_t = \rho + \Phi(\iota) - \delta - \frac{\sigma^2}{\eta_t}. \tag{8}$$

Step 2: The Law of Motion of η_t. To finish deriving the equilibrium, we need to describe how shocks Z affect the state variable $\eta_t = N_t/(q_t K_t)$. First, since η_t is a ratio, the following stochastic differentiation formula for ratios will be helpful for us:

ITO'S FORMULA FOR A RATIO. Suppose two processes X_t and Y_t follow

$$\frac{dX_t}{X_t} = \mu_t^X\,dt + \sigma_t^X\,dZ_t \quad \text{and} \quad \frac{dY_t}{Y_t} = \mu_t^Y\,dt + \sigma_t^Y\,dZ_t.$$

Then

$$\frac{d(X_t/Y_t)}{X_t/Y_t} = (\mu_t^X - \mu_t^Y + (\sigma_t^Y)^2 - \sigma_t^X\sigma_t^Y)\,dt + (\sigma_t^X - \sigma_t^Y)\,dZ_t. \tag{9}$$

Second, it convenient to express the laws of motion of the numerator and denominator of η_t by focusing total risk of each part, and the Sharpe ratio given by (8). Specifically,

$$\frac{dN_t}{N_t} = r_t\,dt + \underbrace{\frac{\sigma}{\eta_t}}_{\text{risk}}\,\underbrace{\frac{\sigma}{\eta_t}}_{\text{Sharpe}}\,dt + \frac{\sigma}{\eta_t}\,dZ_t - \underbrace{\rho\,dt}_{\text{consumption}} \quad \text{and}$$

$$\frac{d(q_tK_t)}{q_tK_t} = r_t\,dt + \underbrace{\sigma}_{\text{risk}}\,\underbrace{\frac{\sigma}{\eta_t}}_{\text{Sharpe}}\,dt + \sigma\,dZ_t - \underbrace{\rho\,dt}_{\text{dividend yield}}.$$

In the latter equation, we subtract the dividend yield from the total return on capital to obtain the capital gains rate.

Using the formula for the ratio,

$$\frac{d\eta_t}{\eta_t} = (r_t + \sigma^2/\eta_t^2 - \rho - r_t - \sigma^2/\eta_t + \rho + \sigma^2 - \sigma^2/\eta_t)\,dt + (\sigma/\eta_t - \sigma)\,dZ_t$$

$$= \frac{(1-\eta_t)^2}{\eta_t^2}\sigma^2\,dt + (1-\eta_t)\sigma\,dZ_t. \tag{10}$$

Observations. A few observations about what happens in equilibrium. Variable η_t fluctuates with macro shocks - a positive shock increases the wealth allocation of experts, because experts are levered. A negative shock erodes η_t, and experts require a higher risk premium to hold risky assets. Experts are convinced to keep holding risky assets by the increasing Sharpe ratio

$$\frac{\sigma}{\eta_t} = \frac{\rho + \Phi(\iota) - \delta - r_t}{\sigma},$$

which goes to ∞ as η_t goes to 0. Strangely, this is achieved due to the risk-free rate $r_t = \rho + \Phi(\iota) - \delta - \sigma^2/\eta_t$ going to $-\infty$, rather than due to depressed price of the risky asset. Because q_t is constant, there is no endogenous risk, no amplification and no volatility effects. Therefore, in this model, assumptions that allow such a simple solution also eliminate any price effects that we are so interested in. We have to work harder to get those effects. [3]

[3] Besides the absence of price effects, another problem with this model is that in the

However, now at least we have seen how equilibrium conditions can be translated into formulas that describe how the economy behaves. Next, we can consider more complicated models, in which the price of the risky asset q_t reacts to shocks. We also develop methodology that allows for agents to have more complicated preferences, for nontrivial distribution of assets among agents, and for investment.

Optimal Portfolio Choice.

Consider an agent, whose marginal utility of wealth θ_t follows

$$\frac{d\theta_t}{\theta_t} = \mu_t^\theta \, dt + \sigma_t^\theta \, dZ_t. \tag{11}$$

Then the agent's stochastic discount factor (SDF) at time t is $e^{-\rho s}\theta_{t+s}/\theta_t$ for payoff received at time $t+s \geq t$, where ρ is the rate at which the agent discounts utility. The SDF can be used to price assets: for an asset with return

$$dr_t^A = \mu_t^A \, dt + \sigma_t^A \, dZ_t,$$

the following asset-pricing relationship has to hold

$$0 = \mu_t^\theta - \rho + \mu_t^A + \sigma_t^A \sigma_t^\theta. \tag{12}$$

This relationship ensures that if wealth n_t is invested in asset A, so that $dn_t/n_t = dr_t^A$, then $n_{t+s}e^{-\rho s}\theta_{t+s}/\theta_t$ is a martingale. Equation (12) is very important and used often in analysis of continuous-time heterogeneous-agent models.

One consequence of (12) is that the Sharpe ratio of any risky investment available to this agent is given by $-\sigma_t^\theta$. Indeed, applying (12) to the risk-free asset with return r_t, we obtain

$$0 = \mu_t^\theta - \rho + r_t.$$

long run expert sector becomes so large that it overwhelms the whole economy. To see this, not that the drift of η_t is always positive. This feature is typical of models in which one group of agents has an advantage over another group - in this case only experts can invest in the risky asset. It is possible to prevent expert sector from becoming too large through an additional assumption. For example, Bernanke, Gertler and Gilchrist (1999) assume that experts are randomly hit by a shock that makes them households. Alternatively, if experts have a higher discount rate than households, then greater consumption rate prevents expert sector from becoming too large.

Subtracting this from (12) and dividing by risk, we obtain

$$\underbrace{\frac{\mu_t^A - r_t}{\sigma_t^A}}_{\text{Sharpe ratio}} = -\sigma_t^\theta. \tag{13}$$

Example 1. Let us see how equation (6) for a log utility agent follows from a more general relationship (13). Note that the agent's marginal utility is $\theta_t = 1/c_t$, where consumption c_t is proportional to net worth according to (5). Therefore, if the volatility of net worth is σ_t^n, then $\sigma_t^\theta = -\sigma_t^n$. Substituting this in (13), we get

$$\frac{\mu_t^A - r_t}{\sigma_t^A} = \sigma_t^n,$$

where the left hand side is the Sharpe ratio, and the right hand side is the volatility of net worth.

Example 2. In general, assets can be priced from consumption of risk-averse agents. Consider an agent with CRRA utility

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma},$$

whose consumption follows

$$\frac{dc_t}{c_t} = \mu_t^c\, dt + \sigma_t^c\, dZ_t.$$

Then, by Ito's lemma, marginal utility $c^{-\gamma}$ follows

$$\frac{d(c_t^{-\gamma})}{c_t^{-\gamma}} = \underbrace{\left(-\gamma\mu_t^c + \frac{\gamma(\gamma+1)}{2}(\sigma_t^c)^2\right)}_{\mu_t^\theta} dt + \underbrace{(-\gamma\sigma_t^c)}_{\sigma_t^\theta} dZ_t. \tag{14}$$

Substituting this into (13), we obtain the following relationship for the pricing of any risky asset relative to the risk-free asset:

$$\frac{\mu_t^A - r_t}{\sigma_t^A} = \gamma\sigma_t^c. \tag{15}$$

A Model with Price Effects and Instabilities.

We now illustrate how these principles can be used to solve a more complex model, which we borrow and extend from Brunnermeier and Sannikov (2014). We will be able to get a number of important takeaways from the model:

1. Equilibrium dynamics is characterized by a relatively stable steady state, where the system spends most of the time, and a crisis regime. In the steady state, experts are adequately capitalized and risk premia fall. The experts' consumption offsets their earnings - hence the steady state is formed. Experts have the capacity to absorb most macro shocks, hence prices near the steady state are quite stable. However, an unusually long sequence of negative shocks causes experts to suffer significant losses, and pushes the equilibrium into a crisis regime. In the crisis regime, experts are undercapitalized and constrained. Shocks affect their demand for assets, and thus affect prices of the assets that experts hold. This creates feedback effects, creating endogenous risk.

2. High volatility during crisis times may push the system in a very depressed region, where experts' net worth is close to 0. If that happens, it takes a long time for the economy to recover. Thus, the system spends a considerable amount of time far away from the steady state. The stationary distribution may be bimodal.

3. Endogenous risk during crises makes assets more correlated.

4. There is a volatility paradox, because risk-taking is endogenous. If the aggregate risk parameter σ becomes smaller, the economy does not become more stable. The reason is that experts allow greater leverage, and pay out profits sooner, in response to lower fundamental risk. Due to greater leverage, the economy is prone to crises even when exogenous shocks are smaller. In fact, endogenous risk during crises may actually be higher when σ is lower.

5. Financial innovations, such as securitization and derivatives hedging, that allow for more efficient risk-sharing among experts, may make the system less stable in equilibrium. The reason, again, is that risk-taking is endogenous. By diversifying idiosyncratic risks, experts tend to increase leverage, amplifying systemic risks.

The key to studying endogenous risk is to introduce capital *illiquidity* into the model. That is, we allow households to buy capital as well. However, assume that households are less productive - when they hold capital, their productivity parameter $\underline{a} < a$ is lower than that of experts. Thus, households earn the return of

$$d\underline{r}_t^k = \underbrace{\frac{\underline{a} - \iota_t}{q_t}\, dt}_{\text{dividend yield}} + \underbrace{(\Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q)\, dt + (\sigma + \sigma_t^q)\, dZ_t}_{\frac{d(q_t k_t)}{q_t k_t},\ \text{the capital gains rate}} \tag{16}$$

when they manage capital. The households' return differs from that of experts, (3), only in the dividend yield that they earn.

We may also wish to allow the model to be more general in several other ways. In particular we want to know how one may be able to solve a model in which (1) there is a force that prevents experts from "saving their way out" away from the constraints, e.g. experts could have a higher discount rate ρ than that of households, r (2) we may want to allow experts and households to have more general utility functions,[4] e.g. CRRA with risk aversion coefficient γ,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma},$$

and (3) we may want to allow experts to issue some equity, even though they cannot be 100% equity financed. Specifically, suppose that experts must retain at least a fraction $\bar{\chi} \in (0, 1]$ of equity.

To summarize, experts and household maximize, respectively

$$E\left[\int_0^\infty e^{-\rho t} u(c_t)\right], \quad \text{and} \quad E\left[\int_0^\infty e^{-rt} u(\underline{c}_t)\right].$$

We denote the fraction of capital allocated to experts by $\psi_t \leq 1$, and look for an equilibrium, and the fraction of equity retained by experts by $\chi_t \geq \bar{\chi}$.

We want to characterize how any history of shocks $\{Z_s, s \leq t\}$ maps to equilibrium prices q_t and r_t, asset allocations ψ_t and χ_t, and consumption so that (1) all agents maximize utility through optimal consumption and portfolio choice and (2) markets clear.

[4] Brunnermeier and Sannikov (2014) explicitly consider the case of risk-neutral experts and households. Experts are constrained to consume nonnegative quantities, but households can consume both positive and negative amounts. This assumption leads to the simplification that the risk-free rate in the economy always equals the households' discount rate r.

The equilibrium concept here is portfolio optimization subject to constraints (no short-selling of capital and a bound on equity issuance by experts). For example, households can invest in capital, the risk-free asset, and experts' equity, and optimize over portfolio weights on these three assets (with a nonnegative weight on capital). Thus, the solution is based on a classic problem in asset pricing. Note also that because of required returns are different between households and experts, the experts' inside equity will generally earn a different return from the equity held by households - experts will earn "management fees" that households do not earn.[5]

We will solve for the equilibrium in three steps. First, we introduce the experts' marginal utility of wealth θ_t, and use asset pricing and market clearing conditions to write down equations that stochastic laws of motion of q_t, θ_t and ψ_t must satisfy. Second, we focus on the experts' balance sheets to write down the law of motion of

$$\eta_t = \frac{N_t}{q_t K_t},$$

expert wealth as a percentage of the whole wealth in the economy, where K_t is the total amount of capital in the economy. Third, we look for a Markov equilibrium, and characterize equations for q_t, ψ_t, etc., as functions of η_t. We solve these equations numerically.

Step 1: The Equilibrium Conditions. Suppose that the experts' marginal utility of wealth θ_t follows an equation of the form (11), and that of households,

$$\frac{d\underline{\theta}_t}{\underline{\theta}_t} = \mu_t^{\underline{\theta}}\, dt + \sigma_t^{\underline{\theta}}\, dZ_t.$$

Consider first a model in which experts cannot issue any equity. Then (13) implies the following asset-pricing relationship for capital held by experts:

$$\frac{\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - r_t}{\sigma + \sigma_t^q} = -\sigma_t^{\theta}. \tag{17}$$

An analogous relationship for households is

$$\frac{\frac{\underline{a}-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - r_t}{\sigma + \sigma_t^q} \leq -\sigma_t^{\underline{\theta}}, \tag{18}$$

[5]This is not a universal assumption in literature. For example, He and Krishmanurthy (2013) assume that returns are equally split between experts and households, so that rationing is required to prevent households from demanding more expert equity than the total supply of expert equity.

with equality if $\psi_t < 1$, i.e. households hold capital in positive amount. Note that households may choose not to hold any capital, and if so, then the Sharpe ratio they would earn from capital could fall below that required by the asset-pricing relationship.

Now, let us allow for equity issuance. Then we should replace the right-hand side of (17) with

$$\chi_t(-\sigma_t^\theta) + (1-\chi_t)(-\underline{\sigma}_t^\theta). \tag{19}$$

The required return on capital held by experts depends on the equilibrium capital structure that experts use. Furthermore, if experts require a higher risk premium than households, then $\chi_t = \bar{\chi}$, i.e. experts will issue the maximum equity they can. Thus, we have[6]

$$\chi_t = \bar{\chi} \text{ if } -\sigma_t^\theta > -\underline{\sigma}_t^\theta, \quad \text{otherwise} \quad -\sigma_t^\theta = -\underline{\sigma}_t^\theta.$$

Under this condition, we can replace χ_t with $\bar{\chi}$ in (19).

It is useful to combine (17) and (18), eliminating μ_t^q and r_t. Allowing for equity issuance, we have

$$\frac{(a-\underline{a})/q_t}{\sigma+\sigma_t^q} \geq \bar{\chi}(\underline{\sigma}_t^\theta - \sigma_t^\theta), \tag{20}$$

with equality if $\psi_t < 1$.

The required risk premia can be tied to the agent's consumption processes via (14) in the CRRA case and to the agent's net worth processes in the special logarithmic case. Under the baseline risk-neutrality assumptions of Brunnermeier and Sannikov (2014), $\underline{\sigma}_t^\theta = 0$ when households are risk-neutral and financially unconstrained - i.e. they can consume negatively.

We will use these conditions to characterize q_t, ψ_t, χ_t, etc. as functions of η_t. Before we do that, though, we must derive an equation for the law of motion of $\eta_t = N_t/(q_t K_t)$.

Step 2: The Law of Motion of η_t. It is convenient to express the laws of motion of the numerator and denominator of η_t by focusing on risks and risk premia. Specifically, the experts' net worth follows

$$\frac{dN_t}{N_t} = r_t\,dt + \underbrace{\frac{\chi_t\psi_t}{\eta_t}(\sigma+\sigma_t^q)}_{\text{risk}} \underbrace{(-\sigma_t^\theta)}_{\text{risk premium}} \; dt + \frac{\chi_t\psi_t}{\eta_t}(\sigma+\sigma_t^q)\,dZ_t - \frac{C_t}{N_t}\,dt.$$

[6] We can rule out the case that $-\sigma_t^\theta < -\underline{\sigma}_t^\theta$ and $\chi_t = 1$: experts cannot face lower risk premia than households if households hold zero risk.

For the capital gains of all capital in the economy, note that $\chi_t\psi_t$ of capital risk is held by experts and $1-\chi_t\psi_t$, households. Therefore, subtracting the dividend yield from the world portfolio of capital, we find that

$$\frac{d(q_tK_t)}{q_tK_t} = r_t\,dt + (\sigma+\sigma_t^q)\left(\psi_t\chi_t(-\sigma_t^\theta) + (1-\psi_t\chi_t)(-\sigma_t^\theta)\right)\,dt + (\sigma+\sigma_t^q)\,dZ_t$$

$$-\frac{\psi_t a + (1-\psi_t)\underline{a} - \iota_t}{q_t}\,dt.$$

Using the already familiar formula (9) for a ratio of two stochastic processes, we have

$$\frac{d\eta_t}{\eta_t} = \mu_t^\eta dt + \sigma_t^\eta dt = \left(\frac{\psi_t a + (1-\psi_t)\underline{a} - \iota_t}{q_t} - \frac{C_t}{N_t} + (\sigma+\sigma_t^q)(1-\psi_t\chi_t)(\sigma_t^{\underline{\theta}} - \sigma_t^\theta)\right)dt$$

$$+\left(1-\frac{\chi_t\psi_t}{\eta_t}\right)(\sigma+\sigma_t^q)(\sigma+\sigma_t^q+\sigma_t^\theta)\,dt + \left(\frac{\chi_t\psi_t}{\eta_t}-1\right)(\sigma+\sigma_t^q)\,dZ_t. \quad (21)$$

Step 3: Converting the equilibrium conditions and the laws of motion (21) into equations for $q(\eta)$, $\theta(\eta)$ and $\psi(\eta)$, etc. The steps we need to go through to convert the equilibrium conditions and the law of motion of η_t into numerically solvable equations for $q(\eta)$, $\psi(\eta)$, etc, can depend on the underlying assumptions on the agents' preferences (with the logarithmic case being the easiest to solve). In each case, we have to use Ito's lemma, which allows us to replace terms such as σ_t^q, σ_t^θ, μ_t^q etc. with expressions containing the derivatives of q and θ, in order to arrive at solvable differential equations for these functions in the end. Before we dive into this, let me give a very simple and well-known example to illustrate the gist of what we have to do.

Example 3. This example is from the well-known endogenous default model of Leland (1994). Equity holders are sitting on assets whose value follows a geometric Brownian motion

$$\frac{dV_t}{V_t} = r\,dt + \sigma\,dZ_t \quad (22)$$

under the risk-neutral measure. Default happens when the value of assets falls to some value of V_B (which is later endogenized). Before default, equity holders must be paying coupons to debt holders at rate C. In the event of default, equity holders abandon the assets, and debt holders receive the liquidating value of assets of αV_B, where $\alpha \in (0,1)$.

Under the risk-neutral measure, the expected return of any security must be r. Thus, if equity E_t follows $dE_t = \mu_t^E E_t\, dt + \sigma_t^E E_t\, dZ_t$, then we must have[7]

$$r = \mu_t^E - C/E_t. \tag{23}$$

That is, after paying coupons equity holders must receive an expected return of r.

Suppose we would like to calculate how the value of equity E_t depends on the value of assets V_t. Then we are face a problem that is completely analogous to that in the model of Brunnermeier and Sannikov (2014). We have a law of motion of the state variable V_t and a relationship (23) that the stochastic motion of E_t has to satisfy, and we would like to characterize E_t as a function of V_t.

How can we do this? Easy. Using Ito's lemma

$$\mu_t^E E_t = rV_t E'(V_t) + \frac{1}{2}\sigma^2 V_t^2 E''(V_t),$$

and so (23) becomes

$$r = \frac{rVE'(V) + \frac{1}{2}\sigma^2 V^2 E''(V)}{E(V)} - \frac{C}{E(V)}. \tag{24}$$

If function $E(V)$ satisfies this equation, then the process $E_t = E(V_t)$ will satisfy (23). We are able to go from an equation like (23) to a differential equation (24) by assuming that the value of equity is a *function* of the value of assets.

We can solve the second-order ordinary differential equation (ODE) (24) if we have two boundary conditions. The relevant boundary conditions in the context of the Leland (1994) model are $E(V_B) = 0$ and that $V - E(V) \to C/r$ as $V \to \infty$.

Our problem is similar to that of Leland (1994): we have an equation for the stochastic law of motion of the state variable (21), as well as the equilibrium conditions that processes q_t, θ_t and ψ_t must satisfy. Certainly, the equations are more complicated than those of Leland (1994), and the law of motion of η_t is endogenous (i.e. it depends on q_t, θ_t and ψ_t). However, the mechanics of solving these equations is the same - use Ito's lemma. For example, using Ito's lemma we can tie the volatility of q_t with the first derivative of $q(\eta)$,

$$\sigma_t^q q(\eta) = q'(\eta)\, \underbrace{(\chi_t\psi_t - \eta_t)(\sigma + \sigma_t^q)}_{\eta\sigma_t^\eta}. \tag{25}$$

[7]Unlike in Leland (1994), I assumed here that there are no taxes, so equity holders do not get any tax shield benefits by paying coupons.

Likewise, using Ito's lemma,

$$\mu_t^q q(\eta) = \mu_t^\eta \eta q'(\eta) + \frac{1}{2}(\sigma_t^\eta)^2 \eta^2 q''(\eta) \tag{26}$$

Below, we demonstrate how the model can be solved for particular cases.

The Risk-Neutral Model. Assume that experts and households are risk neutral, and while experts must consume non-negatively, households can have both positive and negative consumption. This set of assumptions corresponds to those of the baseline model of Brunnermeier and Sannikov (2014). We would like to construct differential equations to solve for the functions $q(\eta)$, $\theta(\eta)$ and $\psi(\eta)$. The equations will be of second order in $q(\eta)$ and $\theta(\eta)$, i.e. we'll design a procedure to compute $q''(\eta)$ and $\theta''(\eta)$, as well as $\psi(\eta)$, from η, $q(\eta)$, $q'(\eta)$ and $\theta(\eta)$, $\theta'(\eta)$. Note also that, since households are risk-neutral and financially unconstrained, $\underline{\theta}_t = 1$ at all times, so $\sigma_t^{\underline{\theta}} = 0$ and experts issue the maximum allowed fraction of equity to households, so $\chi_t = \bar{\chi}$ at all times.

In this case $q(\eta)$ is an increasing function that satisfies the boundary condition

$$q(0) = \max_{\iota} \frac{\underline{a} - \iota}{r - \Phi(\iota) + \delta},$$

the Gordon growth formula for the value of capital when it is permanently managed by households. Any expert can get infinite utility if he can buy capital at the price of $q(0)$, so

$$\lim_{\eta \to 0} \theta(\eta) = \infty. \tag{27}$$

Function $\theta(\eta)$ is decreasing: the marginal value of the experts' net worth is declining as η rises, and investment opportunities become less valuable. Experts refrain from consumption whenever $\theta(\eta) > 1$, and consume only at point η^* where $\theta(\eta^*) = 1$, i.e. the marginal value of the experts' net worth is exactly 1. That point becomes the reflecting boundary of the system. In addition, at the reflecting boundary η^* functions $q(\eta)$ and $\theta(\eta)$ must satisfy

$$q'(\eta^*) = \theta'(\eta^*) = 1.$$

Now to the differential equations. Equation (25) implies that

$$\sigma + \sigma_t^q = \frac{\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(\bar{\chi}\psi_t - \eta_t)}, \tag{28}$$

and by Ito's lemma,

$$\sigma_t^\theta = \frac{\theta'(\eta)}{\theta(\eta)} \frac{(\bar{\chi}\psi_t - \eta_t)\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(\bar{\chi}\psi_t - \eta_t)}. \tag{29}$$

Therefore, plugging these expressions into the asset-pricing equation (20), we obtain

$$\frac{a - \underline{a}}{q(\eta)} \geq -\bar{\chi}\frac{\theta'(\eta)}{\theta(\eta)} \frac{(\bar{\chi}\psi - \eta)\sigma^2}{\left(1 - \frac{q'(\eta)}{q(\eta)}(\bar{\chi}\psi - \eta)\right)^2}. \tag{30}$$

Assuming that $q'(\eta) > 0$ and $\theta'(\eta) < 0$, the right-hand side is increasing from 0 to ∞ as $\bar{\chi}\psi - \eta$ rises from 0 to $q(\eta)/q'(\eta)$). Thus, we have to set $\psi = 1$ it is possible to do so (i.e. $\bar{\chi} - \eta < q(\eta)/q'(\eta)$) and if it allows the inequality in (30) to hold. Otherwise we determine ψ by solving the quadratic equation (30), in which we replace the $\geq$ sign with equality.

After that, we can find σ_t^q from (28), σ_t^θ from (29), μ_t^η and σ_t^η from (21) (where we set $C_t = 0$ since experts consume only at the boundary η^*), μ_t^q from the asset-pricing condition

$$\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - r = \bar{\chi}(\sigma + \sigma_t^q)(-\sigma_t^\theta),$$

μ_t^θ from the pricing condition for the risk-free asset

$$\mu_t^\theta = \rho - r,$$

$q''(\eta)$ from Ito's formula (26) and $\theta''(\eta)$ from Ito's formula

$$\mu_t^\theta \theta(\eta) = \mu_t^\eta \eta \theta'(\eta) + \frac{1}{2}(\sigma_t^\eta)^2 \eta^2 \theta''(\eta).$$

Matlab function fnct.m provided with these notes implements this algorithm: finding $q''(\eta)$ and $\theta''(\eta)$, as well as $\psi(\eta)$, from η, $q(\eta)$, $q'(\eta)$ and $\theta(\eta)$, $\theta'(\eta)$.

Solving the system of ODEs numerically. We can use function fnct.m together with an ODE solver in Matlab, such as ode45, to solve the system of equations. We need to perform a search, since our boundary conditions are defined at two endpoints of $[0, \eta^*]$, and we also need to deal with a singularity at $\eta = 0$. The following algorithm performs an appropriate search and deals with the singularity issue, effectively, by solving the system of equations with the boundary condition $\theta(0) = M$, for a large constant M, instead of (27):

Algorithm. *Set*

$$q(0) = \max_{\iota} \frac{\underline{a} - \iota}{r - \Phi(\iota) + \underline{\delta}}, \quad \theta(0) = 1 \quad \text{and} \quad \theta'(0) = -10^{10}.$$

Perform the following procedure to find an appropriate boundary condition $q'(0)$*. Set* $q_L = 0$ *and* $q_H = 10^{15}$*. Repeat the following loop 50 times. Guess* $q'(0) = (q_L + q_H)/2$*. Use Matlab function ode45 to solve for* $q(\eta)$ *and* $\theta(\eta)$ *on the interval* $[0, ?)$ *until one of the following events is triggered, either (1)* $q(\eta)$ *reaches the upper bound*

$$q_{\max} = \max_{\iota} \frac{a - \iota}{r - \Phi(\iota) + \delta},$$

(2) the slope $\theta'(\eta)$ *reaches 0 or (3) the slope* $q'(\eta)$ *reaches 0. If integration has terminated for reason (3), we need to increase the initial guess of* $q'(0)$ *by setting* $q_L = q'(0)$*. Otherwise, we decrease the initial guess of* $q'(0)$*, by setting* $q_H = q'(0)$*.*

At the end, $\theta'(0)$ *and* $q'(0)$ *reach 0 at about the same point, which we denote by* η^**. Divide the entire function* θ *by* $\theta(\eta^*)$*.*[8] *Then plot the solutions.*

Script solve_equilibrium.m provided with these notes implements this algorithm, and uses event function evntfct.m to terminate integration. The solution is economically meaningful even with the boundary condition $\theta(0) = M$: it corresponds to an assumption that, in the event all experts are wiped out, any measure-zero set of experts that still has wealth left gets utility M per dollar of net worth.

Properties of the Solution. Point η^* plays the role of the steady state of our system. The drift of η_t is positive everywhere on the interval $[0, \eta^*)$, because the expert sector, which is more productive than the household sector, is growing in expectation. Thus, the system is pushed towards η^* by the drift.

[8] We can do this because whenever functions θ and q satisfy our system of equation, so do functions $\Theta\theta$ and q for any constant Θ. Because of that, also, it is immaterial what we set $\theta(0)$ to.

It turns out that the steady state is relatively stable, because volatility is low near η^*. To see this, recall that the amount of endogenous risk in asset prices, from (25), is given by

$$\sigma_t^q = \frac{q'(\eta)}{q(\eta)} \frac{(\psi_t - \eta_t)\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(\psi_t - \eta_t)}.$$

From the boundary conditions, $q'(\eta^*) = 0$, so there is no endogenous risk near η^*.

However, below η^*, endogenous risk increases as $q'(\eta)$ becomes larger. As prices react to shocks, fundamental risk becomes amplified. As we see from the expression for σ_t^q, this amplification effect is nonlinear, since $q'(\eta)$ enters not only the numerator, but also denominator. This happens due to the feedback effect: an initial shock causes η_t to drop, which leads to a drop in q_t, which hurts experts who are holding capital and leads to a further decrease in η_t, and so on.

Of course, far in the depressed region the volatility of η_t, $\sigma_t^\eta \eta_t$, becomes low again in this model. This leads to a bimodal stationary distribution of η_t in equilibrium.[9]

Volatility paradox refers to the phenomenon that systemic risk can build up in quiet environments. We can illustrate this phenomenon through comparative statics on σ or the degree of the experts' equity constraint χ. One may guess that the system becomes a lot more stable as σ or χ decline.

This is not the case, as illustrated in Figure 1 for parameters $\rho = 6\%$, $r = 5\%$, $a = 11\%$, $\underline{a} = 5\%$, $\delta = 3\%$ and an investment function of the form $\Phi(\iota) = \frac{1}{\kappa}(\sqrt{1 + 2\kappa\iota} - 1)$, $\kappa = 10$, $\bar{\chi} = 1$ and various values of σ. (The investment technology in this example has quadratic adjustment costs: an investment of $\Phi + \kappa\Phi^2/2$ generates new capital at rate Φ.)

The volatility paradox shows itself in a number of metrics. As exogenous risk declines,

- maximal endogenous risk σ_t^q may increase (as σ drops from 25% to 10% in Figure 1)

[9] One can prove that the stationary distribution is bimodal analytically by analyzing the asymptotic properties of the solutions near $\eta = 0$ and using the Kolmogorov forward equations that characterize the stationary density - see Brunnermeier and Sannikov (2014) for details.

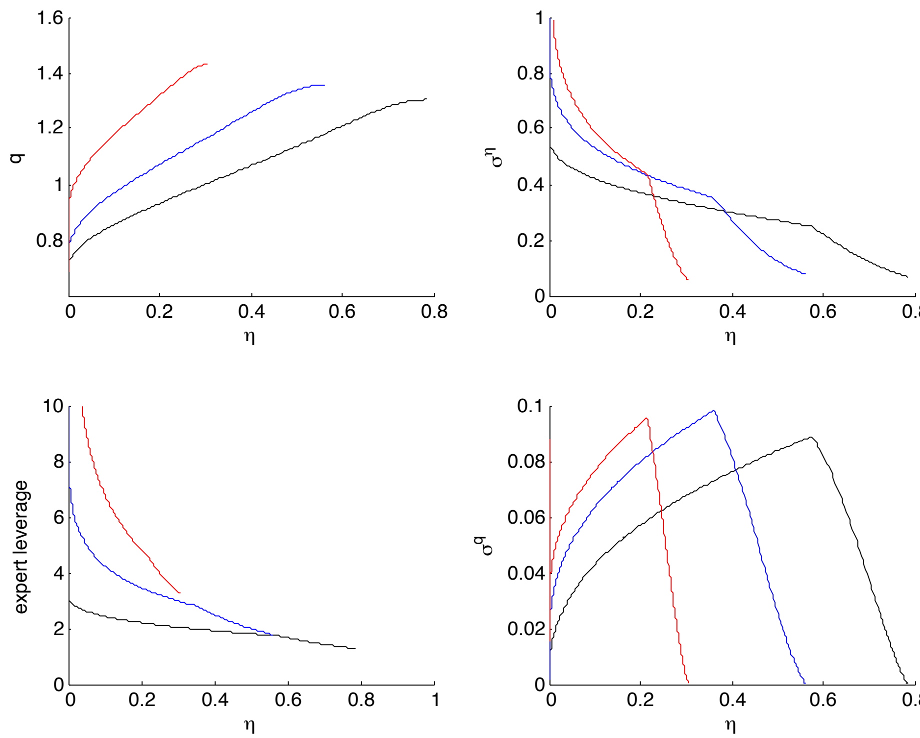

Figure 1: Equilibrium with $\sigma = 25\%$ (black), 10% (blue) and 2.5% (red).

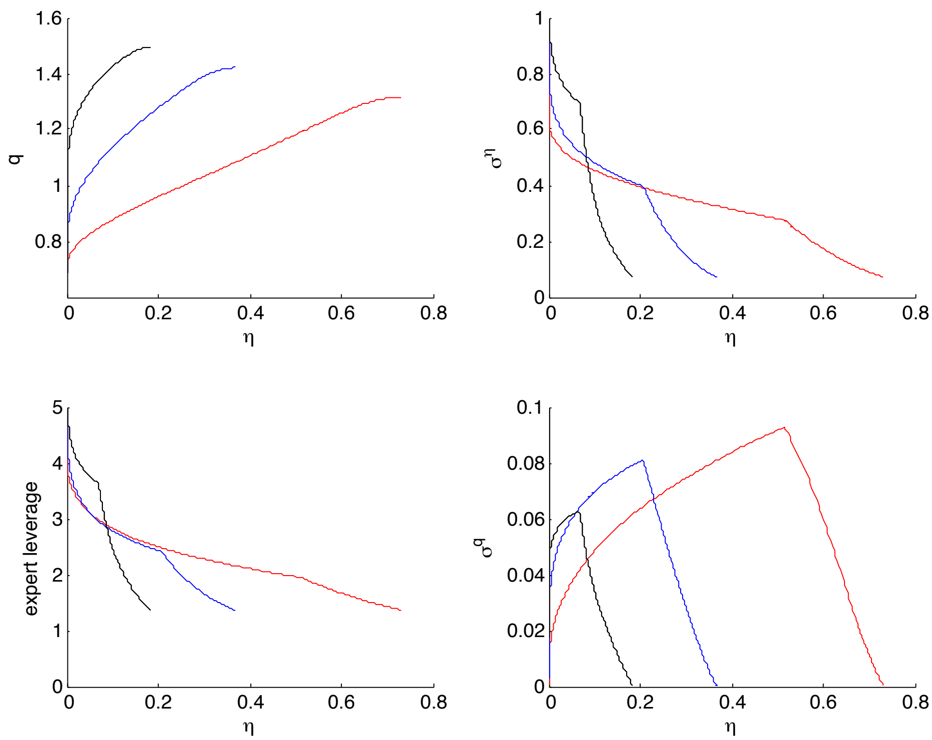

Figure 2: Equilibrium with $\bar{\chi} = 1$ (black), 0.5 (blue) and 0.25 (red).

- the volatility σ_t^η near $\eta = 0$ rises (and this result can be proved analytically)
- from the steady state η^* it takes less time for volatility $\sigma + \sigma_t^q$ to double
- from the steady state, it may take less time to reach the peak of the crisis η^ψ, where experts start selling capital to households[10]

Figure 2 takes the same parameters and $\sigma = 20\%$, but varies $\bar{\chi}$. As $\bar{\chi}$ falls, expert net worth at the steady state η^* drops significantly, and the volatility σ_t^η in the crisis regime rises.

Model with Logarithmic Utility. Logarithmic utility significantly simplifies the computation of the solution in this (and many other) environ-

[10] As σ declines, the system spends less time in the crisis region, so some measures of stability improve, but the amount of time spent in crisis does not converge to 0 as $\sigma \to 0$.

ments, and leads to a characterization that allows for analytic proofs of many interesting properties.

First, from (5), the market-clearing condition for output is

$$(r(1-\eta)+\rho\eta)q = \psi a + (1-\psi)\underline{a} - \iota(q). \tag{31}$$

This equation completely determines the price of output $q(\eta)$ in the normal regime when $\psi = 1$ and at the boundary $\eta = 0$, where $\psi = 0$. For example, in the special case that $\Phi(\iota) = \log(\kappa\iota + 1)/\kappa$, the price of output is given by

$$q(\eta) = \frac{a + 1/\kappa}{r(1-\eta)+\rho\eta + 1/\kappa}$$

in the normal regime and $q(0) = (\underline{a} + 1/\kappa)/(r + 1/\kappa)$ at the boundary.

The law of motion of η_t can be simplified to

$$\frac{d\eta_t}{\eta_t} = \left((r-\rho)(1-\eta_t) + \frac{1-\psi_t\chi_t}{1-\eta_t}(\sigma+\sigma_t^q)\sigma_t^\eta + (\sigma_t^\eta)^2\right)\,dt$$
$$+ \underbrace{\left(\frac{\chi_t\psi_t}{\eta_t} - 1\right)(\sigma+\sigma_t^q)}_{\sigma_t^\eta}\,dZ_t. \tag{32}$$

Normal regime. From the price of capital in the normal regime, given by (31), we can compute directly the equilibrium dynamics there. As long as

$$\frac{\bar{\chi}}{\eta}(\sigma+\sigma_t^q) \geq \sigma + \sigma_t^q,$$

i.e. $\eta_t \leq \bar{\chi}$, the volatility of the experts' net worth is greater than the volatility of capital, and so we have $\chi_t = \bar{\chi}$. The dynamics on the interval $[\eta^\psi, \bar{\chi}]$ then can be found by plugging $\psi_t = 1$ and $\chi_t = \bar{\chi}$ directly into (32) and using

$$\sigma + \sigma_t^q = \frac{\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(\bar{\chi}\psi_t - \eta)}.$$

On $[1/\bar{\chi}, 1]$, experts and households can share risk perfectly, i.e. $\chi_t = \eta_t$ and

$$\frac{d\eta_t}{\eta_t} = (r-\rho)(1-\eta_t)\,dt.$$

Crisis regime. While we know the dynamics in the normal regime, we have to solve a separate equation to determine where the normal regime ends and the crisis regime (with $\psi < 1$) starts. In order to do this, we need to derive an appropriate equation for the crisis regime - it will be a first-order ordinary differential equation for q with a boundary condition at $q(0)$. That is, we need a procedure to determine $q'(\eta)$ given η and $q(\eta)$.

Starting with η and $q(\eta)$, (31) allows us to compute ψ. We have $\chi_t = \bar{\chi}$, since experts face greater risk than households, and so they choose to issue the maximum allowed fraction of equity. Furthermore, we can use the (20),

$$\frac{a - \underline{a}}{q(\eta)} = \bar{\chi} \frac{\bar{\chi}\psi_t - \eta_t}{\eta_t(1 - \eta_t)} (\sigma + \sigma_t^q)^2$$

to determine σ_t^q and (25),

$$\sigma_t^q q(\eta) = q'(\eta)(\bar{\chi}\psi_t - \eta_t)(\sigma + \sigma_t^q)$$

to determine $q'(\eta)$.

Figure 3 illustrates a solution for the same parameters as before, $\rho = 6\%$, $r = 5\%$, $a = 11\%$, $\underline{a} = 5\%$, $\delta = 3\%$ and an investment function of the form $\Phi(\iota) = \frac{1}{\kappa}(\sqrt{1 + 2\kappa\iota} - 1)$, $\kappa = 10$, $\bar{\chi} = 1$ and various values of σ.

Here, point η^* where the drift of η_t becomes 0 plays the role of a steady state of the system. In the absence of shocks, the system stays still at the steady state due to be absence of drift there. It is a point at which risk premia decline sufficiently so that the experts' earnings are exactly offset by their slightly higher consumption rates. Note that the drift is positive below η^*, and negative above η^*.

As σ declines, η^* falls, i.e. experts become more levered at the steady state. Thus, leverage is endogenous. The point η^ψ where crisis starts also declines as σ falls. In fact, it is possible to prove both of these two facts analytically from the differential equation provided above.

But what happens as $\sigma \to 0$? Does endogenous risk disappear altogether, and does the solution converge to first best? It turns out that no: in the limit as $\sigma \to 0$, η^ψ converges not to 0 but a finite number. At the same time, η^* becomes equal to η^ψ for sufficiently low σ. Both of these facts, again, can be proved analytically and are left as an exercise.

If σ is not the crucial parameter that affects system stability - the economy is prone to crises even for very low σ - then what is? It turns out that the level of endogenous risk in crises depends primarily on the illiquidity of capital -

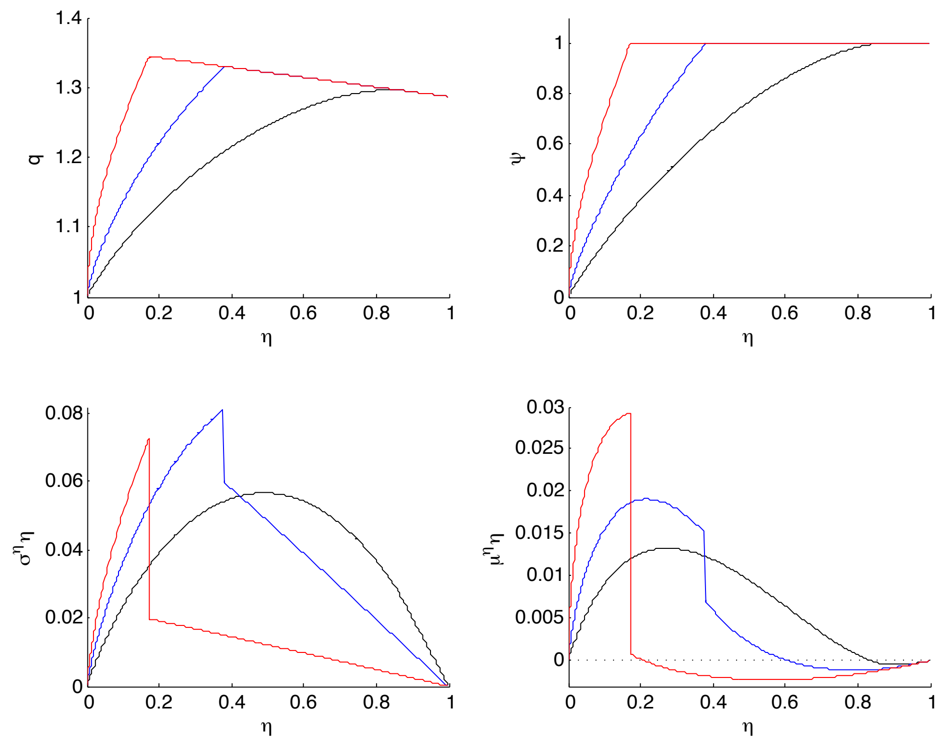

Figure 3: Equilibrium with $\sigma = 20\%$ (black), 10% (blue) and 2.5% (red).

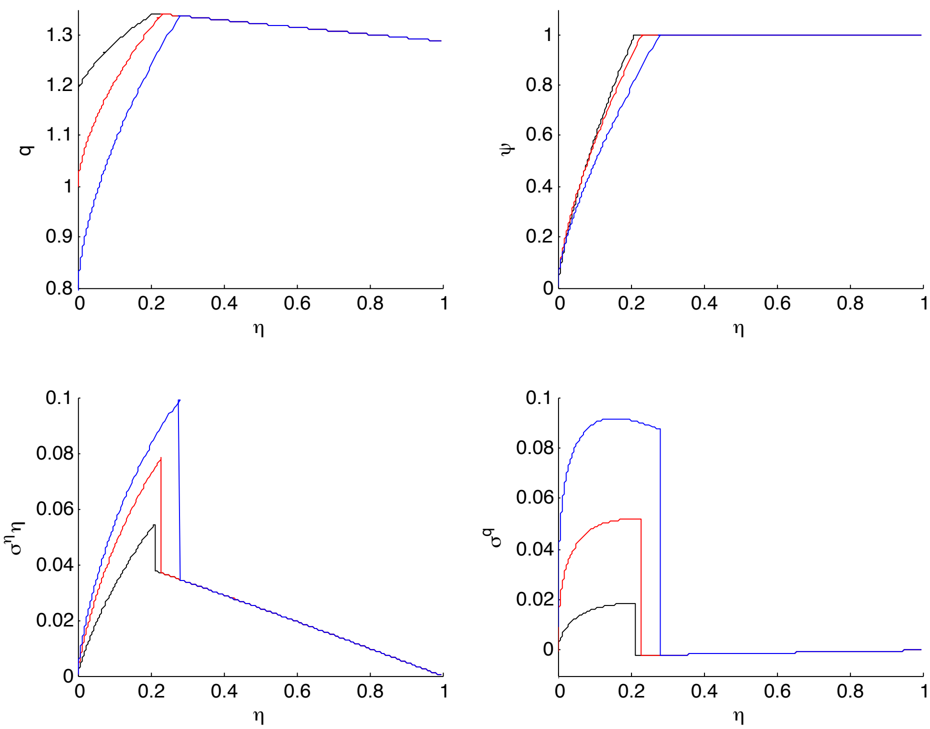

Figure 4: Equilibrium with $\sigma = 5\%$ and $\underline{a} = .8$ (black), .05 (red) and .02 (blue).

the difference between parameter a and $\underline{a}$ that determines how much less households value capital, in the event that they have to buy it, relative to experts. Figure 4 illustrates the equilibrium for several values of $\underline{a}$. Note that endogenous risk rises sharply as $\underline{a}$ drops. However, it is easy to see from the characterization above that neither the dynamics in the normal regime nor η^* when it falls in the interior of the normal regime depend on $\underline{a}$. Thus, while expert leverage responds endogenously to fundamental risk σ, it does not respond to endogenous tail risk when preferences are logarithmic.

Welfare. One of the most important characteristics of equilibria is welfare - the utility that a representative expert or household gets in equilibrium. Let us discuss how welfare can be computed.

In the risk-neutral model above, solution already provides welfare. Since the experts' total wealth is given by $\eta_t q_t K_t$, their utility is $\theta_t \eta_t q_t K_t$, while the households' total utility is $(1-\eta_t)q_t K_t$.

In the logarithmic model, we can find welfare by solving a differential equation. The value function of experts takes the form $\log(K_t)/\rho + H(\eta_t)$. Indeed, if K_t doubles while η_t stays the same, then the experts' consumption doubles in perpetuity and so their utility increases by $\log(2)/\rho$. The experts' value function must satisfy the HJB equation

$$\rho\left(\frac{\log(K)}{\rho} + H(\eta)\right) = \log(\rho\eta q(\eta)K) + \frac{1}{\rho}E\left[\frac{d\log(K)}{dt}\right] + E\left[\frac{dH(\eta)}{dt}\right].$$

Using Ito's lemma and simplifying, we obtain

$$\rho H(\eta) = \log(\rho\eta q(\eta)) + \Phi(\iota(q)) - \delta - \frac{\sigma^2}{2} + H'(\eta)\eta\mu_t^\eta + \frac{H''(\eta)}{2}\eta^2(\sigma_t^\eta)^2.$$