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

September 12, 2015

Notes prepared for Yuliy's lectures at "Princeton Initiative: Macro, Money and Finance," based on work with Markus as well as other recent literature.

In previous Princeton Initiatives, I have focused this lecture on methods, building towards solving numerically the model from the paper with Markus "A Marcoeconomic Model with a Financial Sector." Since the first Princeton Initiative (this one is 5th), a lot of interesting work has been done in the area of financial frictions, and continuous time has proved a useful modeling technique. So, this time I would like to take a step forward and adopt a broader perspective. I would like to talk more generally about classic methodologies that are at the core of models developed recently, including optimal portfolio choice and general equilibrium with more than one types of agents. Also, instead of concentrating on particular model, I would like to give examples of various constructions - bringing in ideas about how to model labor market, housing, money, etc., - to open a broader perspective for things that can be done.

The motivation for a broader perspective is to be of most use to the attending Ph.D. students, to provide with building blocks that can be used immediately to start exploring ideas. This lecture will cover some theory, and then provide examples to illustrate how theory can be used.

Optimal Portfolio Choice with Logarithmic Utility.

Consider an agent who has logarithmic utility with discount rate ρ, of the form

$$E\left[\int_0^\infty e^{-\rho t} \log c_t \, dt\right], \tag{1}$$

where c_t is consumption at time t. Logarithmic utility has two convenient properties, which greatly streamline analysis in many models. First, for agents with log utility

$$\text{consumption} = \rho \cdot \text{net worth} \tag{2}$$

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

$$\text{volatility of wealth} = \text{Sharpe ratio.} \tag{3}$$

The Sharpe ratio on the right-hand side corresponds to the optimal portfolio, which maximizes expected return for a given level of risk. The higher the Sharpe ratio, the more risk the agent is willing to take on.

With multiple assets, equation (3) implies that the difference in expected returns of any two assets corresponds to the correlation of their relative risk to the risk of wealth. That is,

$$E[dr_t^1 - dr_t^2] = \text{Cov}(dr_t^1 - dr_t^2, dn_t/n_t), \tag{4}$$

where n_t on the right-hand side reflects the volatility of net worth.

The meaning of these equations may not be completely clear until we see how they can be applied.

Example 1: Basak and Cuoco (1998).

This example illustrates how optimal portfolio choice conditions, together with market clearing, can be used to derive endogenous prices and equilibrium dynamics in an economy with financial frictions.

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. What happens in equilibrium?

Production Technology and Returns. Capital evolves according to

$$\frac{dk_t}{k_t} = (\Phi(\iota_t) - \delta)\, dt + \sigma\, dZ_t, \tag{5}$$

and produces output

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

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.

If (endogenous in equilibrium) price of capital per unit q_t follows

$$\frac{dq_t}{q_t} = \mu_t^q\, dt + \sigma_t^q\, dZ_t, \tag{6}$$

then using Ito's lemma, we get capital gains of

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

Then, the return on capital is

$$dr_t^k = \underbrace{\frac{a - \iota_t}{q_t}\, dt}_{\text{dividend yield}} + \underbrace{(\Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q)\, dt + (\sigma + \sigma_t^q)\, dZ_t}_{\frac{d(k_t q_t)}{k_t q_t},\ \text{the capital gains rate}}. \tag{7}$$

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

The optimal investment rate $\iota(q)$ that maximizes return must satisfy

$$\Phi'(\iota_t) = \frac{1}{q_t}. \tag{8}$$

One convenient functional form is

$$\Phi(\iota) = \frac{\log(\kappa\iota + 1)}{\kappa}, \quad \text{then} \quad \iota(q) = \frac{q-1}{\kappa}. \tag{9}$$

Here κ is the adjustment cost parameter. We have $\Phi'(0) = 1$. Higher κ makes function Φ more concave, and as $\kappa \to 0$, $\Phi(\iota) \to \iota$, a fully elastic investment function with no adjustment costs.

Equilibrium. To find an equilibrium, we use optimal portfolio choice and market clearing conditions to ultimately write down a map from any initial allocation and any history of shocks Z to prices q_t, r_t, etc.

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*

1. *agents behave to maximize utility and*
2. *markets clear*

Let us use the equilibrium conditions. First, from (2), the aggregate consumption of all agents is $\rho q_t K_t$. Aggregate output is $(a - \iota(q))K_t$, so the market clearing condition for consumption goods is

$$\rho q_t = a - \iota(q_t). \tag{10}$$

Condition (10) implies that the price q_t in this model is *constant.* For function $\Phi(\iota)$ of the form (9), we have

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

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 (3) 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 (3) 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

$$\frac{\rho + \Phi(\iota) - \delta - r_t}{\sigma}. \tag{11}$$

Because experts must hold all the risky capital in the economy, if N_t is the aggregate net worth of experts, then the volatility of their net worth is

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

Equating this with (11), we can express the risk-free rate as

Using (3),

$$r_t = \rho + \Phi(\iota) - \delta - \sigma^2 \frac{q_t K_t}{N_t}. \tag{12}$$

To finish the description of equilibrium, we have to express how N_t evolves with shocks Z. N_0 depends on the initial allocation. After that, it is convenient to derive the law of motion of N from the amount of risk experts take and the Sharpe ratio they earn, directly, i.e.

$$\frac{dN_t}{N_t} = r_t\,dt + \frac{\rho + \Phi(\iota) - \delta - r_t}{\sigma}\frac{q_t K_t}{N_t}\sigma\,dt + \frac{q_t K_t}{N_t}\sigma\,dZ_t - \underbrace{\rho\,dt}_{\text{consumption}}.$$

We could stop here, but it is even more convenient to summarize the history of shocks in terms of a state variable, the experts' wealth share

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

Then, we can rewrite

$$\frac{dN_t}{N_t} = \left(r_t + \frac{\sigma^2}{\eta_t^2} - \rho\right)\,dt + \frac{\sigma}{\eta_t}\,dZ_t.$$

Since world capital earns the same Sharpe ratio,

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

Consequently,[1]

$$\frac{d\eta_t}{\eta_t} = (r_t + \sigma^2/\eta_t^2 - \rho - r_t - \sigma^2/\eta_t + \rho + \sigma^2 - \sigma^2/\eta_t)\, dt + (\sigma/\eta_t - \sigma)\, dZ_t$$
$$= \frac{(1-\eta_t)^2}{\eta_t^2}\sigma^2\, dt + \frac{1-\eta_t}{\eta_t}\sigma\, dZ_t. \tag{14}$$

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. This is achieved (somewhat strangely) 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. A richer model would include all these effects.[2] I am going to save the richer model "for

[1] 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

$$\frac{d(X_t/Y_t)}{X_t/Y_t} = (\mu_t^X - \mu_t^Y + (\sigma_t^Y)^2 - \sigma_t^X \sigma_t^Y)\, dt + (\sigma_t^X - \sigma_t^Y)\, dZ_t. \tag{13}$$

[2] Another feature of this model - as any model in which different agents have the same preferences but different investment opportunities - is that in the 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. 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 *retirement* 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.

dessert" at the end of this lecture.

One important goal of this example is seeing how portfolio choice conditions can be used to derive the evolution of wealth distribution, and endogenous evolution of prices. The next example also uses logarithmic utility - it introduces idiosyncratic shocks and presents a model of employment.

Example 2: Quadrini (2014).

This is a nice model, which is very tractable already in discrete time. Let me illustrate some interesting elements using a continuous-time framework. There are two types of agents - entrepreneurs who hire labor and a unit mass of workers. Entrepreneurs have to face idiosyncratic risk, and they save to self-insure. The model is about the availability of assets they can hold - these assets are ultimately tied to borrowing of the household sector, sometimes through banks (but I will not talk about banks here). Quadrini likes to contrast this model to a class of other models, related to Bernanke-Gertler-Gilchrist and Kiyotaki-Moore, in which entrepreneurs are borrowers rather than savers.

Entrepreneurs with logarithmic utility (1) choose how much labor to hire at wage w_t. Employment generates expected output $\bar{z}$ but exposes entrepreneurs to idiosyncratic Brownian risk $\tilde{Z}$. Specifically, entrepreneur net worth evolves according to

$$\frac{dn_t}{n_t} = (r_t - c_t)\, dt + h_t \left((\bar{z} - w_t)\, dt + \sigma\, d\tilde{Z}_t \right), \tag{15}$$

where r_t is the risk-free rate and c_t is the consumption ratio. This problem is identical to optimal portfolio choice with risky investment with Sharpe ratio equal to $(\bar{z} - w_t)/\sigma$. For log utility, the optimal hiring rate must satisfy

$$\sigma h_t = \frac{\bar{z} - w_t}{\sigma}.$$

Entrepreneur net worth/savings are liabilities of households, so let us discuss the households. There is a unit mass of households that provide labor. They are risk-neutral with the same discount rate ρ as entrepreneurs and have utility function

$$C_t - \alpha \frac{H_t^{1+1/\nu}}{1 + 1/\nu}.$$

Optimal labor supply is given by

$$H_t = \left(\frac{w_t}{\alpha}\right)^{\nu}.$$

Household also hold housing in fixed total amount of $\bar{K}$, which generates goods at rate χ per unit of the asset, and can be traded at price p_t. Households can also borrow an amount $L_t \leq \lambda p_t \bar{K}$ from entrepreneurs at the risk-free rate of r_t, where $\lambda \in (0, 1)$ determines the collateral value of housing.

The aggregate *financial* wealth of households (they also have human capital) follows

$$d(p_t \bar{K} - L_t) = (w_t H_t + \chi \bar{K} + \bar{K} dp_t - C_t - r_t L_t)\, dt.$$

The aggregate net worth of entrepreneurs is L_t. Let us solve this model.

Equating labor supply and demand, equilibrium wage is

$$\frac{\bar{z} - w_t}{\sigma^2} L_t = \left(\frac{w_t}{\alpha}\right)^{\nu}. \tag{16}$$

Wage w_t is increasing in the savings L_t of the entrepreneur sector.

Steady State. At the steady state, L_t is constant, so aggregating (15) we must have

$$0 = r_t - \rho + h_t(\bar{z} - w_t) \quad \Rightarrow \quad r_t = \rho - \frac{(\bar{z} - w_t)^2}{\sigma^2}. \tag{17}$$

The risk-free rate is always less than ρ, so households want to borrow to the limit at this rate.

The return on a portfolio of housing short of cash, levered to the maximum, must be ρ. Hence

$$\frac{\chi \bar{K} - r_t \lambda p_t \bar{K}}{(1 - \lambda) p_t \bar{K}} = \frac{\chi}{(1 - \lambda) p_t} - \frac{r_t \lambda}{1 - \lambda} = \rho. \tag{18}$$

As p_t increases, $\chi/((1 - \lambda)p_t)$ decreases. At the same time, $L_t = \lambda p_t \bar{K}$ increases, w_t increases and r_t given by (17) increases. Hence, the left-hand side of (18) is decreases, and there is a unique solution p^*.

This model illustrates how the presence of idiosyncratic shocks can lead to a demand for financial assets. Idiosyncratic shocks are important - in my opinion - because many macro models focus on representative agents and ignore idiosyncratic uncertainty. At the same time, idiosyncratic uncertainty is key - arguably individuals face a lot more idiosyncratic risks than aggregate risk. Bewley models include idiosyncratic risk, but unlike those models, this model can be solved in closed form. Another example that can be solved in closed form is in the "I-Theory of Money," where exposure to idiosyncratic risk is also a choice similar to portfolio choice (i.e. agents earn a risk premium for choosing to expose themselves to idiosyncratic risk). In the I-Theory, idiosyncratic risk leads to a demand for money.

At this point, let me move on, and talk more about methodology. Ultimately, conditions (2), (3) and (4) come from dynamic optimization - he did not derive them, we just used them. Let me derive conditions for optimal portfolio choice with general utility functions.

Optimal Portfolio Choice.

Here the key result is that any asset available to the agent to invest in can be priced from the agent's marginal utility of consumption.[3]

Specifically, if the agent has discount rate ρ and consumption utility $u(c_t)$, then $\xi_t = e^{-\rho t} u'(c_t)$ is the stochastic discount factor (SDF) to price assets. We can write

$$\frac{d\xi_t}{\xi_t} = -r_t\, dt - \pi_t\, dZ_t, \tag{19}$$

where r_t is the risk-free rate and π_t is the price of risk dZ. For any asset A that the agent can invest in, with return

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

we must have

$$\mu_t^A = r_t + \pi_t \sigma_t^A. \tag{20}$$

If wealth n_t is invested in asset A, so that $dn_t/n_t = dr_t^A$, then (19) and (20) imply that $n_t \xi_t$ is a martingale. Indeed, the drift of $n_t \xi_t$ by Ito's lemma is

$$\mu_t^A - r_t - \sigma_t^A \pi_t = 0.$$

Equations (19) and (20) are simple, yet extremely powerful. As we shall see later, these are simply first-order conditions for optimal portfolio choice. The

[3] If the agent is risk-neutral, then have to use the marginal utility of wealth rather than the marginal utility of consumption when the agent is "at the corner," i.e. consumes 0.

first-order condition for optimal consumption is that the marginal utility of consumption equals the marginal utility of wealth.

Remark 1. With CRRA utility, the agent's value function takes a power form

$$\frac{u(\omega_t n_t)}{\rho}.$$

This form comes from the fact that if the agent's wealth changes by a factor of x, then his optimal consumption at all future states changes by the same factor - hence ω_t is determined so that $u(\omega_t)/\rho$ is the value function at unit wealth. Marginal utility of consumption and marginal utility of wealth are equated if $c_t^{-\gamma} = \omega_t^{1-\gamma} n_t^{-\gamma}/\rho$, or

$$\frac{c_t}{n_t} = \rho^{1/\gamma} \omega_t^{1-1/\gamma}. \tag{21}$$

For log utility, $\gamma = 1$ and this equation implies that $c_t/n_t = \rho$ as we claimed in (2).

For any utility function other than log, we can also solve for ω_t from the agent's consumption rate c_t/n_t, and hence compute the agent's welfare.

Remark 2. Let us justify our asset-pricing relationships for logarithmic utility. Equation (2) follows directly from (21). Next, since the SDF is $\xi_t = e^{-\rho t}/c_t = e^{-\rho t}/n_t$ (for any ω_t) it follows that $\sigma_t^n = \sigma_t^c = \pi_t$ (i.e. minus the volatility of ξ_t). Hence, (20) implies that

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

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

Remark 3. Back to general CRRA utility, given the agent's consumption process of

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

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

$$\frac{d(c_t^{-\gamma})}{c_t^{-\gamma}} = \left(-\gamma \mu_t^c + \frac{\gamma(\gamma+1)}{2}(\sigma_t^c)^2\right) dt - \gamma \sigma_t^c \, dZ_t. \tag{22}$$

Substituting this into (20), we obtain the following relationship for the pricing of any risky asset relative to the risk-free asset:

$$\frac{\mu_t^A - r_t}{\sigma_t^A} = \gamma \sigma_t^c = \pi_t. \tag{23}$$

Minus the drift of the SDF is the risk-free rate, i.e.

$$r_t = \rho + \gamma \mu_t^c - \frac{\gamma(\gamma+1)}{2}(\sigma_t^c)^2. \tag{24}$$

Derivation. Let me derive the pricing equations using the stochastic maximum principle (for a change). Consider an agent whose net worth follows

$$dn_t = n_t \left(r_t\, dt + \sum_A x_t^A ((\mu_t^A - r_t)\, dt + \sigma_t^A\, dZ_t) \right) - c_t\, dt,$$

where x_t^A are portfolio weights on various assets A.

The agent would like to maximize

$$E\left[\int_0^\infty e^{-\rho s} u(c_s)\, ds\right]$$

and has initial wealth $n_0 > 0$. Investment opportunities are stochastic and exogenous, i.e. they do not depend on the agent's strategy.

Stochastic maximum principle allows us to derive first-order conditions for maximization from the Hamiltonian. Introducing a multiplier ξ_t on n_t (i.e. marginal utility of wealth) and denoting the volatility of ξ_t by $-\pi_t \xi_t$, the Hamiltonian is written as

$$H = e^{-\rho t} u(c) + \xi_t \underbrace{(r_t + \sum_A x^A (\mu_t^A - r_t)) n_t - c}_{\text{drift of } n_t} - \pi_t \xi_t \underbrace{\sum_A x^A \sigma_t^A}_{\text{volatility of } n_t}.$$

By differentiating the Hamiltonian with respect to controls, we get the first-order conditions, and by differentiating it with respect to the state n_t, we get the law of motion of the multiplier ξ_t.

The first-order condition with respect to c is

$$e^{-\rho t} c_t^{-\gamma} = \xi_t,$$

implies that the multiplier on the agent's wealth is his discounted marginal utility of consumption. The first-order condition with respect to the portfolio weight x^A implies

$$\xi_t(\mu_t^A - r_t) - \pi_t \xi_t \sigma_t^A = 0,$$

which implies (20).

In addition, the drift of ξ_t is

$$-H_n = -\xi_t r_t,$$

where we already used the first-order conditions with respect to x^A to perform cancellations. It follows that the law of motion of ξ_t is

$$d\xi_t = -\xi_t r_t \, dt - \pi_t \xi_t \, dZ_t,$$

which corresponds to (19).

Example 3: CRRA and Constant Investment Opportunities.

With constant investment opportunities, then ω_t is a constant, hence (21) implies that $\sigma_t^c = \sigma_t^n$, just like in the logarithmic case. Hence, (23) implies that

$$\underbrace{\frac{\mu_t^A - r}{\sigma_t^A}}_{\pi} = \gamma \sigma_t^n.$$

i.e. the volatility of net worth is Sharpe ratio divided by the risk aversion coefficient γ.[4]

Now, the agent's net worth follows

$$\frac{dn_t}{n_t} = r \, dt + \frac{\pi^2}{\gamma} \, dt + \frac{\pi}{\gamma} \, dZ_t - c_t/n_t \, dt$$

and, since consumption is proportional to net worth, (24) implies that

$$r = \rho + \gamma \left(r + \frac{\pi^2}{\gamma} - \frac{c_t}{n_t} \right) - \frac{\gamma(\gamma+1)}{2} \frac{\pi^2}{\gamma^2} \quad \Rightarrow \quad \frac{c_t}{n_t} = \rho + \frac{\gamma - 1}{\gamma} \left(r - \rho + \frac{\pi^2}{2\gamma} \right).$$

Hence, consumption ratio increases with better investment opportunities when $\gamma > 1$ and falls otherwise.

[4]This property also holds more generally when ω_t has no volatility, the agent's investment opportunities evolve deterministically.

A Model with Price Effects and Instabilities.

Let me know solve a more complex model, borrowed from from Brunnermeier and Sannikov (2014). Relative to the Basak-Cuoco model, 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.

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.

I am going to modify the Basak-Cuoco model described earlier to allow for firesales, i.e. households can also hold capital. This will let us talk about capital illiquidity and endogenous risk.

Assume that households are less productive than experts - when they hold capital, their productivity parameter $\underline{a} < a$ is lower than that of experts.

Thus, households earn the return of

$$d\underline{r}_t^k = \underbrace{\frac{\underline{a} - \iota_t}{q_t} dt}_{\text{dividend yield}} + \underbrace{(\Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q)\, dt + (\sigma + \sigma_t^q)\, dZ_t}_{\frac{d(q_t k_t)}{q_t k_t},\ \text{the capital gains rate}} \qquad (25)$$

when they manage capital. The households' return differs from that of experts, (7), only in the dividend yield that they earn.

Assume that experts have discount rate ρ while households have a lower discount rate r - this assumption prevents experts from "saving their way out" away from the constraints. I am going to start by deriving equilibrium conditions in a general form, and later on use logarithmic utility to "fast forward" to the solution.

Denote by $\psi_t \leq 1$ the fraction of world capital held by experts at any time t. We want to characterize how any history of shocks $\{Z_s, s \leq t\}$ maps to equilibrium prices q_t and r_t, and allocations of capital ψ_t and consumption so that (1) all agents maximize utility through optimal consumption and portfolio choice and (2) markets clear.

Let me break down analysis by discussing first, the asset-pricing conditions and then the evolution of the wealth distribution. At this stage, I am going to allow the agent's utility to be of any general CRRA form. After that, we could transform the conditions to a system of ordinary differential equations that can be solved numerically. However, for the sake of time, I am going to come back to logarithmic utility at that point and solve the model for this special case.

Step 1: The Equilibrium Conditions. Experts and household have different investment opportunities, but both of them can trade the risk-free asset. Therefore, they have different stochastic discount factors, which follow

$$\frac{d\xi_t}{\xi_t} = -r_t\, dt - \pi_t\, dZ_t, \quad \text{and} \quad \frac{d\underline{\xi}_t}{\underline{\xi}_t} = -r_t\, dt - \underline{\pi}_t\, dZ_t$$

respectively.

Then (20) implies the following asset-pricing relationship for capital held by experts:

$$\frac{\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - r_t}{\sigma + \sigma_t^q} = \pi_t. \qquad (26)$$

An analogous relationship for households is

$$\frac{\frac{\underline{a}-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - r_t}{\sigma + \sigma_t^q} \leq \underline{\pi}_t, \tag{27}$$

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.

It is useful to combine (26) and (27), eliminating μ_t^q and r_t, to obtain

$$\frac{(a - \underline{a})/q_t}{\sigma + \sigma_t^q} \geq \pi_t - \underline{\pi}_t, \tag{28}$$

with equality if $\psi_t < 1$.

The required risk premia can be tied to the agent's consumption processes via (22) in the CRRA case and to the agent's net worth processes in the special logarithmic case.

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

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

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

$$\frac{d(q_t K_t)}{q_t K_t} = r_t\,dt + (\sigma + \sigma_t^q)\left((\psi_t\pi_t + (1-\psi_t)\underline{\pi}_t)\;dt + dZ_t\right)$$

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

Using the already familiar formula (13) for a ratio of two stochastic processes, we have

$$\frac{d\eta_t}{\eta_t} = \left(\frac{\psi_t a + (1-\psi_t)\underline{a} - \iota_t}{q_t} - \frac{C_t}{N_t} + (1-\psi_t)(\pi_t - \underline{\pi}_t)(\sigma + \sigma_t^q)\right)\,dt$$

$$-\frac{\psi_t - \eta_t}{\eta_t}(\sigma + \sigma_t^q)(\sigma + \sigma_t^q - \pi_t)\, dt + \frac{\psi_t - \eta_t}{\eta_t}(\sigma + \sigma_t^q)\, dZ_t. \tag{29}$$

Model with Logarithmic Utility. Logarithmic utility significantly simplifies the computation of the solution in this (and many other) environments, and leads to a characterization that allows for analytic proofs of many interesting properties.

First, from (2), the market-clearing condition for output is

$$(r(1-\eta) + \rho\eta)q = \psi a + (1-\psi)\underline{a} - \iota(q). \tag{30}$$

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.

Since

$$\pi_t = \frac{\psi_t}{\eta_t}(\sigma + \sigma_t^q)$$

$$\text{and} \quad (\pi_t - \underline{\pi}_t)(\sigma + \sigma_t^q) = \frac{a - \underline{a}}{q_t} \quad \text{when } \psi_t < 1,$$

the law of motion of η_t can be simplified to

$$\frac{d\eta_t}{\eta_t} = \left((r-\rho)(1-\eta_t) + (1-\psi_t)\frac{a - \underline{a}}{q_t} + (\sigma_t^\eta)^2 \right) dt + \underbrace{\frac{\psi_t - \eta_t}{\eta_t}(\sigma + \sigma_t^q)}_{\sigma_t^\eta}\, dZ_t. \tag{31}$$

The relationship between the function $q(\eta)$ and the volatility of η is as follows. Using Ito's lemma,

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

Hence,

$$\sigma_t^q = \frac{q'(\eta)}{q(\eta)} \frac{(\psi_t - \eta_t)\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(\psi_t - \eta_t)}, \quad \sigma_t^\eta = \frac{\psi_t - \eta_t}{\eta_t} \frac{\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(\psi_t - \eta_t)}. \tag{32}$$

Normal regime. From the price of capital in the normal regime, given by (30), we can compute directly the equilibrium dynamics there. From (32),

$$\sigma_t^\eta = \frac{1-\eta_t}{\eta_t} \frac{\sigma}{1 - \frac{q'(\eta)}{q(\eta)}(1-\eta_t)}, \quad \mu_t^\eta = (r-\rho)(1-\eta_t) + (\sigma_t^\eta)^2.$$

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 $q(\eta)$, (30) allows us to compute ψ. Furthermore, we can use the (28),

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

to determine σ_t^q and (32),

$$\sigma_t^q q(\eta) = q'(\eta)(\psi_t - \eta_t)(\sigma + \sigma_t^q)$$

to determine $q'(\eta)$.

Figure 1 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 $\Phi(\iota) = \frac{1}{\kappa}(\sqrt{1 + 2\kappa\iota} - 1)$, $\kappa = 10$, for 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,

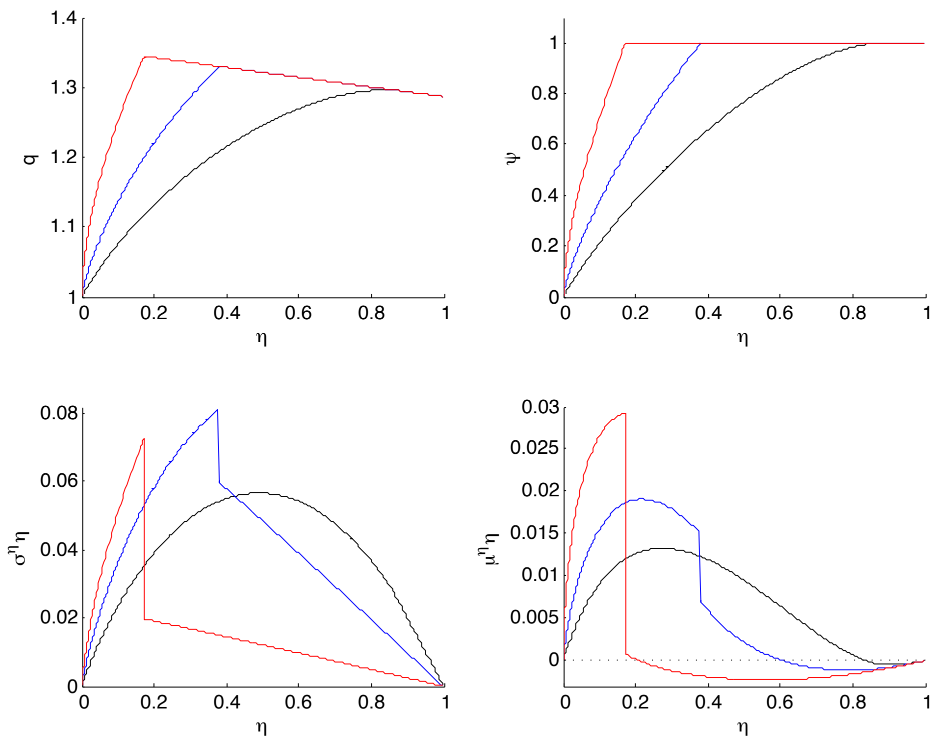

Figure 1: Equilibrium with $\sigma = 20\%$ (black), 10% (blue) and 2.5% (red).

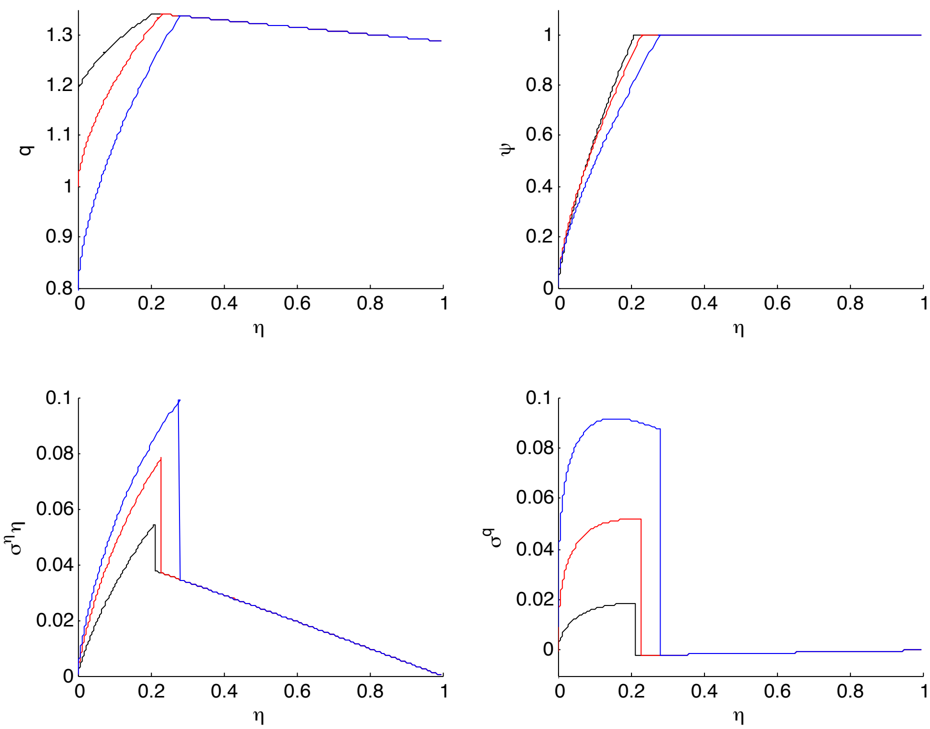

Figure 2: Equilibrium with $\sigma = 5\%$ and $\underline{a} = .8$ (black), .05 (red) and .02 (blue).

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 - 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 2 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.