

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

September 6, 2012

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

This lecture covers the techniques of solving heterogeneous-agent models with financial frictions in continuous time. Financial frictions place restrictions on the agents' asset holdings and risk exposures. Because of these restrictions, the wealth distribution across agents matters. Macro shocks affect the wealth distribution in the economy, and thus asset allocation, economic activity, prices and risk premia. An equilibrium is 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. Such a map is an equilibrium if (1) all agents behave to maximize utility and (2) markets clear. This lecture covers the techniques to translate these two sets of conditions into an equilibrium characterization.

From this exercise, we can observe that

1. Shocks redistribute wealth and have persistent effect on economic activity
2. Price movements amplify wealth redistribution when assets are illiquid, i.e. have higher value in first-best use than in second-best use.

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

3. Amplification is not the same at all times: it is low when agents holding illiquid assets have adequate wealth, but large in crisis times when agents are constrained
4. Agents will tend to accumulate wealth towards stability, but occasional large shocks plunge the system into a crisis

Unstable crisis episodes can redistribute wealth away from productive agents, and caused prolonged inefficiencies in the economy. The frequency of crisis episodes depends on model parameters (such as asset liquidity), the structure of markets (such as the availability of instruments to hedge risks) and macro-prudential and monetary policies.

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 borrowed 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 market microstructure models with noise traders, agents have no price impact.

In the aggregate, the risky asset pays dividend

$$\frac{dD_t}{D_t} = g dt + \sigma dZ_t,$$

where g is the dividend growth rate, and Z is a standard Brownian motion. The price of the risky asset is also determined endogenously, and q_t denotes the price-to-dividend ratio. Thus, the aggregate value of all assets in the economy is $q_t D_t$. If N_t is the aggregate net worth of experts, then the aggregate net worth of households is $q_t D_t - N_t$.

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

$$\text{consumption} = \rho \cdot \text{net worth} \tag{1}$$

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

$$\text{volatility of wealth} = \text{Sharpe ratio of risky investment}, \tag{2}$$

where the volatility of wealth is measured in percent.²

We use equations (1) and (2) 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*

1. *agents choose consumption and portfolio allocation to maximize utility*
2. *and 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

²For 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.

evolve with the realizations of shocks Z . Usually, it is convenient to express this relationship using a state variable, which describes the distribution of wealth. A good state variable to use is the fraction of wealth owned by the experts,

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

which takes values between 0 and 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 maximization and market clearing, to write down equations that 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 , as a function of the history of macro shocks $\{Z_s, s \leq t\}$. After these two steps, we'll know how macro shocks map to η_t , and how η_t maps to q_t and r_t .

Step 1: The Equilibrium Conditions. First, from condition (1), the aggregate consumption of all agents is $\rho q_t D_t$, and aggregate output is D_t . From the market-clearing for consumption goods, these must be equal, and so

$$q_t = \frac{1}{\rho}. \tag{3}$$

Of the total output, experts consume $\rho N_t = \eta_t D_t$ and households, $(1 - \eta_t) D_t$.

Second, we can use condition (2) 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 (2) to get the risk-free rate.

Because q_t is constant, the risky asset earns the return of

$$dr_t^D = \underbrace{1/q_t dt}_{\rho, \text{ dividend yield}} + \underbrace{g dt + \sigma dZ_t}_{\text{capital gains rate}},$$

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

$$\frac{\rho + g - r_t}{\sigma}.$$

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

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

Using (2),

$$\frac{\sigma}{\eta_t} = \frac{\rho + g - r_t}{\sigma} \quad \Rightarrow \quad r_t = \rho + g - \frac{\sigma^2}{\eta_t}. \quad (4)$$

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 $\eta_t = N_t/(q_t D_t)$. We have, using Ito's lemma

$$dN_t = \underbrace{q_t D_t dr_t^D}_{\text{risky investment}} + \underbrace{(N_t - q_t D_t)r_t dt}_{\text{risk-free investment}} - \underbrace{\rho N_t dt}_{\text{consumption}}, \quad (5)$$

$$dD_t = gD_t dt + \sigma D_t dZ_t \quad \Rightarrow \quad d\frac{1}{q_t D_t} = (-g + \sigma^2)\frac{1}{q_t D_t} dt - \sigma\frac{1}{q_t D_t} dZ_t. \quad (6)$$

and so

$$d\eta_t = \frac{1}{q_t D_t} dN_t + N_t d\left(\frac{1}{q_t D_t}\right) + \text{Cov}\left(N_t, \frac{1}{q_t D_t}\right) = \frac{(1 - \eta_t)^2}{\eta_t} \sigma^2 dt + (1 - \eta_t) \sigma dZ_t. \quad (7)$$

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

$$\frac{\sigma}{\eta_t} = \frac{\rho + g - r_t}{\sigma},$$

which goes to ∞ as η_t goes to 0. Strangely, this is achieved due to the risk-free rate $r_t = \rho + g - \sigma^2/\eta_t$ 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

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.

Returns with Investment and Endogenous Risk.

Consider a productive asset (capital) in the amount k_t , which produces gross output $ak_t dt$ and evolves according to

$$\frac{dk_t}{k_t} = (\Phi(\iota_t) - \delta) dt + \sigma dZ_t, \quad (8)$$

where ι_t is the investment rate per unit of capital and $\Phi(\iota_t)$ is a standard 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 $(a - \iota_t)k_t dt$.

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

$$\frac{dq_t}{q_t} = \mu_t^q dt + \sigma_t^q dZ_t, \quad (9)$$

³Besides the absence of price effects, another problem with this model is that in the long run expert sector becomes so large that it overwhelms the whole economy. To see this, note 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.

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

$$dr_t^k = \underbrace{\frac{a - \iota_t}{q_t} dt}_{\text{dividend yield}} + \underbrace{(\Phi(\iota_t) - \delta + \mu_t^q + \sigma \sigma_t^q) dt + (\sigma + \sigma_t^q) dZ_t}_{\frac{d(k_t q_t)}{k_t q_t}, \text{ the capital gains rate}}. \quad (10)$$

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

Note that the rate of 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

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

Optimal Portfolio Choice.

Consider an agent, whose marginal utility of wealth θ_t follows

$$\frac{d\theta_t}{\theta_t} = \mu_t^\theta dt + \sigma_t^\theta dZ_t. \quad (11)$$

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

$$0 = \mu_t^\theta - \rho + \mu_t^A + \sigma_t^A \sigma_t^\theta. \quad (12)$$

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 important and used often in analyses of continuous-time heterogeneous-agent models.

Example 1. Let us see how equation (2) for a log utility agent follows from a more general relationship (12). Note that the agent's marginal utility is $\theta_t = 1/c_t$, where consumption c_t is proportional to net worth according to (1). Therefore, if the volatility of net worth is σ_t^n , then $\sigma_t^\theta = -\sigma_t^n$. For a risky asset with return r_t^A , (12) implies

$$0 = \mu_t^\theta - \rho + \mu_t^A - \sigma_t^A \sigma_t^n. \quad (13)$$

For the risk-free asset, whose volatility is 0,

$$0 = \mu_t^\theta - \rho + r_t. \quad (14)$$

Subtracting (14) from (13), we get

$$\mu_t^A - r_t - \sigma_t^A \sigma_t^n = 0 \quad \Rightarrow \quad \frac{\mu_t^A - r_t}{\sigma_t^A} = \sigma_t^n,$$

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

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

Any risky investment with return dr_t^A , accessible to this agent, must satisfy the pricing equation

$$0 = -\gamma\mu_t^c + \frac{\gamma(\gamma+1)}{2}(\sigma_t^c)^2 - \rho + \mu_t^A - \gamma\sigma_t^c \sigma_t^A.$$

We will not use this relationship in this lecture, but it can be used to solve heterogeneous-agent models with risk aversion, such as that of He and Krishnamurthy (2010).

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 from Brunnermeier and Sannikov (2012). We will be able to get a number of important takeaways from the model:

1. *Normal times vs. crises:* 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 they channel excess profits to payouts. They can easily absorb usual macro shocks by adjusting payouts, and 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, feed into asset prices. This creates feedback effects, which cause high endogenous risks.
2. *Stationary distribution:* High volatility during crisis times may push the system in a 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 is 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.

In addition, we can do experiments to see how various policies affect equilibrium.

While the model is a close to those of BGG and KM, there are two key differences. First, unlike in KM and BGG, we will be able to conveniently describe equilibrium dynamics completely, not just near the steady state. In particular, we uncover the difference between dynamics in normal times and in crises.

Second, in this model the wealth distribution endogenous: i.e. agents choose capital cushions/payouts endogenously given the amount of risk in the system. In contrast, in KM the steady state is pinned down by the exogenous leverage constraint, and in BGG, by an exogenous parameter that determines the exit rate of experts. Endogenous wealth distribution delivers results such as the volatility paradox and the instability due to financial innovations.

The model is as follows. There are two types of agents - experts and households. There are two assets: capital in positive net supply, and the risk-free asset in zero net supply. The financial friction is that neither experts nor households can issue equity backed by their asset holdings - they can only borrow through risk-free debt.⁴

Experts are more productive at managing capital than households. The experts' production technology is characterized by (8). Capital held by households produces a lower dividend stream of $\underline{a}k_t$ instead of ak_t , where $\underline{a} \leq a$, and depreciates at a faster rate $\underline{\delta} \geq \delta$. Under their management,

$$\frac{dk_t}{k_t} = (\Phi(\iota_t) - \underline{\delta}) dt + \sigma dZ_t.$$

Thus, households earn the return of

$$dR_t^k = \underbrace{\frac{\underline{a} - \iota_t}{q_t} dt}_{\text{dividend yield}} + \underbrace{(\Phi(\iota_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q) dt + (\sigma + \sigma_t^q) dZ_t}_{\frac{d(q_t k_t)}{q_t k_t}, \text{ the capital gains rate}} \quad (15)$$

when they manage capital.

Regarding *preferences*, we assume that both experts and households are risk-neutral, but (1) the experts' discount rate ρ is higher than that of households, r , and (2) experts cannot have negative consumption, but households

⁴Brunnermeier and Sannikov (2012) allows for some equity issuance (see Appendix A), but here we restrict attention to debt only to simplify exposition.

can consume negatively. The second assumption simplifies analysis - it implies that households are always financially unconstrained, and that they are willing to lend and borrow arbitrary amounts at the risk-free rate of r . To summarize, experts and household maximize, respectively

$$E \left[\int_0^\infty e^{-\rho t} dc_t \right], \quad dc_t \geq 0, \quad \text{and} \quad E \left[\int_0^\infty e^{-rt} dc_t \right].$$

We denote the fraction of capital allocated to experts by $\psi_t \leq 1$, and look for an equilibrium. That is, we want to characterize how any history of shocks $\{Z_s, s \leq t\}$ maps to the price of capital q_t , asset allocation ψ_t and consumption so that (1) all agents maximize utility and (2) markets clear.

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

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

fraction of wealth in the economy that belongs to experts, 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 , θ_t and ψ_t as functions of η_t . We solve these equations numerically.

Step 1: The Equilibrium Conditions. From the asset pricing equation (12), experts price the risk-free asset according to

$$0 = \mu_t^\theta - \rho + r, \tag{16}$$

and capital, with return given by (10), according to

$$0 = \mu_t^\theta - \rho + \underbrace{\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma \sigma_t^q + (\sigma + \sigma_t^q) \sigma_t^\theta}_{E[dr_t^k]/dt}. \tag{17}$$

In equilibrium, the experts' marginal utility of wealth must always satisfy $\theta_t \geq 1$, and experts consume only when $\theta_t = 1$. We will see that in equilibrium experts consume only at one point, when η_t reaches a critical level η^* .

Because households can consume both positive and negative amounts, their marginal utility of wealth is always 1. If the expected return on the

risky asset is r according to (15), households are willing to hold some of it, i.e. ψ_t can be less than 1. The expected household return from risky capital cannot exceed r (otherwise they demand an infinite amount of the risky asset, and markets will not clear), but it can be less than r if $\psi_t = 1$. Thus, we have

$$\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q \leq r, \quad \text{with equality if } \psi_t < 1. \quad (18)$$

We will use three conditions (16), (17) and (18) to characterize q_t , θ_t and ψ_t 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 . The law of motion of N_t in this model is analogous to (5) except that experts invest only wealth $\psi_t q_t K_t$ in capital, and they consume sporadically (only when $\theta_t = 1$). We have

$$dN_t = \psi_t q_t K_t dr_t^k - (\psi_t q_t K_t - N_t)r dt - dC_t,$$

where dr_t^k is given by (10). Furthermore,

$$\frac{d(q_t K_t)}{q_t K_t} = \underbrace{dr_t^k - \frac{a - \iota_t}{q_t} dt}_{\text{capital gains rate}} - \underbrace{(1 - \psi_t)(\underline{\delta} - \delta) dt}_{\text{adjustment for households}} \Rightarrow$$

$$\frac{d(1/(q_t K_t))}{1/(q_t K_t)} = -dr_t^k + \frac{a - \iota_t}{q_t} dt + (1 - \psi_t)(\underline{\delta} - \delta) dt + (\sigma + \sigma_t^q)^2 dt.$$

Using Ito's lemma again,

$$d\eta_t = (dN_t) \frac{1}{q_t K_t} + N_t d\left(\frac{1}{q_t K_t}\right) + \psi_t q_t K_t (\sigma + \sigma_t^q) \frac{-1}{q_t K_t} (\sigma + \sigma_t^q) dt =$$

$$(\psi_t - \eta_t)(dr_t^k - r dt - (\sigma + \sigma_t^q)^2 dt) + \eta_t \frac{a - \iota_t}{q_t} dt + \eta_t (1 - \psi_t)(\underline{\delta} - \delta) - \eta_t d\zeta_t, \quad (19)$$

where $d\zeta_t = dC_t/N_t$ is the experts consumption rate.

Step 3: Converting the equilibrium conditions (16), (17) and (18) and the law of motion (19) into equations for $q(\eta)$, $\theta(\eta)$ and $\psi(\eta)$. This step boils down to multiple applications of Ito's lemma to convert equilibrium conditions (16), (17) and (18) as well as the law of motion (19)

into differential equations for $q(\eta)$, $\theta(\eta)$ and $\psi(\eta)$ through multiple applications of Ito's lemma. Ito's lemma allows us to replace terms such as σ_t^q , μ_t^θ , etc. with expressions containing the derivatives of q and θ .

The derivations are pretty mechanical, but somewhat lengthy. They can take a very long time, and blow into unmanageably long expressions if you do them in the wrong order. I'll show you a relatively quick, optimized derivation route.

Before that, 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

$$\frac{dV_t}{V_t} = r dt + \sigma dZ_t \quad (20)$$

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 have⁵

$$r = \mu_t^E - C/E_t. \quad (21)$$

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 of Brunnermeier and Sannikov (2012) model: We have a law of motion of the state variable V_t and a relationship (21) 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

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

and so (21) becomes

$$r = \frac{r V E'(V) + \frac{1}{2} \sigma^2 V^2 E''(V)}{E(V)} - \frac{C}{E(V)}. \quad (22)$$

⁵Unlike in Leland (1994), I assumed here that there are no taxes, so equity holders do not get any tax shield benefits by paying coupons.

If function $E(V)$ satisfies this equation, then the process $E_t = E(V_t)$ will satisfy (21). We are able to go from an equation like (21) to a differential equation (22) 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) (22) 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 $V - E(V) \rightarrow C/r$ as $V \rightarrow \infty$.

Our problem is similar to that of Leland (1994): we have an equation for the stochastic law of motion of the state variable (19), as well as conditions (16), (17) and (18) that processes q_t and θ_t must satisfy. Certainly, the equations are more complicated than those of Leland (1994), as

- we have two functions $q(\eta)$ and $\theta(\eta)$ whose derivatives get involved and
- the law of motion of η_t is endogenous (depends on q_t , θ_t and ψ_t)

However, the basic idea for solving these equations is the same: we must express highest-order derivatives $q''(\eta)$ and $\theta''(\eta)$, as functions of lower-order derivatives $q(\eta)$, $q'(\eta)$, $\theta(\eta)$, $\theta'(\eta)$ as well as η_t . In the process, we also determine ψ_t .

To get started, let us eye-ball the equations we got. First, Ito's lemma and (19) lead to

$$\sigma_t^q q(\eta) = q'(\eta) \underbrace{(\psi_t - \eta_t)(\sigma + \sigma_t^q)}_{\sigma_t^\eta}, \quad (23)$$

which has two unknowns, σ_t^q and ψ_t . From (23), we can get another relationship that ties σ_t^q and ψ_t by following

$$\sigma_t^\theta = \frac{\theta'(\eta)}{\theta(\eta)} (\psi_t - \eta_t) (\sigma + \sigma_t^q) \quad (24)$$

likewise from Ito's lemma,

$$\mu_t^q = r - \frac{a - \iota_t}{q_t} - \Phi(\iota_t) + \delta - \sigma \sigma_t^q - (\sigma + \sigma_t^q) \sigma_t^\theta \quad (25)$$

from (17), and checking whether (18) is satisfied, i.e.

$$HH \equiv \frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma \sigma_t^q \leq r \quad (26)$$

with equality if $\psi_t < 1$. Therefore, we can use these equations to solve for ψ_t , σ_t^q , σ_t^θ and μ_t^q through the following procedure.

Procedure 1. Note that the domain for ψ_t is between η_t (which corresponds to no expert leverage and no endogenous risk) and $\min(1, \eta + q(\eta)/q'(\eta))$ (as $\psi_t \rightarrow q(\eta)/q'(\eta)$, endogenous risk σ_t^q converges to ∞). Therefore, we guess $\psi_L = \eta$ and $\psi_H = \min(1, \eta + q(\eta)/q'(\eta))$ and repeat the following loop 30 times.

Guess $\psi_t = (\psi_L + \psi_H)/2$ and using (23), compute

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

Then do (24), (25) and find HH from (26). If $HH > r$ (i.e. households want to hold more capital) lower ψ_t by setting $\psi_H = \psi_t$. Otherwise, set $\psi_L = \psi_t$. Repeat.

After finding ψ_t , it is straightforward to find $q''(\theta)$ and $\theta''(\eta)$ using Ito's lemma.

Procedure 2. Using (16), find

$$\mu_t^\theta = \rho - r \tag{27}$$

and using (19), compute $\mu_t^\eta \eta$, the drift of η_t ,⁶

$$= (\psi_t - \eta_t) \left(\Phi(\iota_t) - \delta + \mu_t^q + \sigma \sigma_t^q - r - (\sigma + \sigma_t^q)^2 \right) + \psi_t \frac{a - \iota_t}{q_t} + \eta_t (1 - \psi_t) (\underline{\delta} - \delta).$$

Note that we left out the experts' consumption $d\zeta_t$ for reasons that will become clear later.

Then using Ito's lemma,

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

We can use these expressions to express $q''(\eta)$ and $\theta''(\eta)$ as functions of what we already computed:

$$q''(\eta) = 2 \frac{\mu_t^q q(\eta) - \mu_t^\eta \eta q'(\eta)}{(\sigma_t^\eta)^2 \eta^2} \quad \text{and} \quad \theta''(\eta) = 2 \frac{\mu_t^\theta \theta(\eta) - \mu_t^\eta \eta \theta'(\eta)}{(\sigma_t^\eta)^2 \eta^2}. \tag{28}$$

⁶A bit of algebra gives a simpler expression

$$\mu_t^\eta \eta = -(\psi_t - \eta_t)(\sigma + \sigma_t^q)(\sigma + \sigma_t^q + \sigma_t^\theta) + \eta_t \left(\frac{a - \iota_t}{q_t} + (1 - \psi_t)(\underline{\delta} - \delta) \right).$$

To find $q''(\eta)$, $\theta''(\eta)$ as well as $\psi(\eta)$ from η , $q(\eta)$, $q'(\eta)$, $\theta(\eta)$ and $\theta'(\eta)$, we perform procedures 1 and 2.

Matlab function `funct.m` provided with these notes implements this.

Boundary Conditions. We have characterized an the equilibrium via a system of second-order differential equations for $q(\eta)$ and $\theta(\eta)$. In order to solve them, we need appropriate boundary conditions. We discuss now what these theoretically are, and how to implement them numerically.

The equilibrium domain is given by an interval $[0, \eta^*]$ over which $\theta(\eta)$ is decreasing towards

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

That is, the experts' marginal utility of wealth increases towards $\eta_t = 0$, when other experts become constrained and profitable opportunities arise. Consequently, the price of capital $q(\eta)$ increases from

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

over the interval $[0, \eta^*]$. Since $\theta(\eta) > 1$ for $\eta < \eta^*$, in equilibrium experts refrain from consumption until η_t reaches η^* .

In fact, η^* plays the role of a reflecting boundary, and consequently

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

We need a fifth boundary condition in order to solve the system of two second-order ODEs and also pin down η^* . That condition is

$$\lim_{\eta \rightarrow 0} \theta(\eta) = \infty, \tag{29}$$

i.e. there is a singularity of a particular kind at 0. We'll discuss next how to take care of it numerically.

Solving the system of ODE's numerically. We can use function `funct.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 (29):

Algorithm. Set

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

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 $\theta(\eta^*)$.⁷ 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.

Let me finish discussing the algorithm by providing remarks of practical matter.

Remark 1. The solution procedure works fine as long as η , $q(\eta)$, $q'(\eta)$, $\theta(\eta)$ and $\theta'(\eta)$ belong to the relevant domain. If these values “wander” off too far from the true solution (e.g. $q'(\eta)$ becomes negative or $q(\eta)$ becomes too large), then computation is no longer meaningful. This can be manifested

⁷We 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.

through complex numbers, problems with integration tolerances, etc. One has to be sensitive towards this problem. Sometimes figuring out the right domain is a matter of trial and error.

Remark 2. Finding the derivatives of relevant functions, such as $q''(\eta)$ and $\theta''(\eta)$, often involves solving nonlinear equations without good properties. In this model, this problem is not huge - it involves finding ψ_t that can be performed in multiple ways (e.g. through binary search). In general, however, one has to think about how to find the right solutions of these equations.

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 experts earn higher returns and refrain from consuming when $\eta_t < \eta^*$. Thus, the system is pushed towards η^* by the drift.

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 (23), 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. The stationary distribution is characterized by Kolmogorov forward equations. By characterizing asymptotic behavior of the system near $\eta_t = 0$, it is possible to prove analytically that stationary density converges to ∞ as $\eta_t \rightarrow 0$ (see Brunnermeier and Sannikov (2012)).

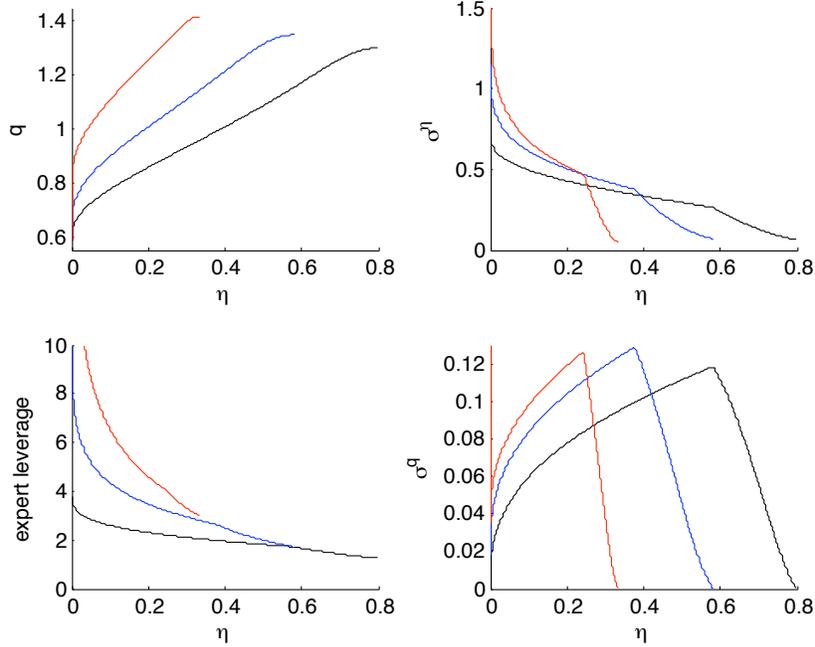


Figure 1: Equilibrium with $\sigma = 25\%$ (black), 10% (blue) and 2.5% (red).

Comparative Statics.

Volatility Paradox refers to the phenomenon that systemic risk can build up in quiet environments. We can illustrate this phenomenon through a comparative static on σ .

Figure 1 illustrates comparative static on σ for parameter values $\rho = 6\%$, $r = 5\%$, $a = 11\%$, $\underline{a} = 5\%$, $\delta = 3\%$, $\underline{\delta} = 5\%$, and $\Phi(\iota) = \frac{1}{10}(\sqrt{1 + 20\iota} - 1)$.⁸

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)
- the volatility σ_t^η near $\eta = 0$ rises (and this result can be proved analytically)

⁸The investment technology in this example has quadratic adjustment costs: an investment of $\Phi + 5\Phi^2$ generates new capital at rate Φ .

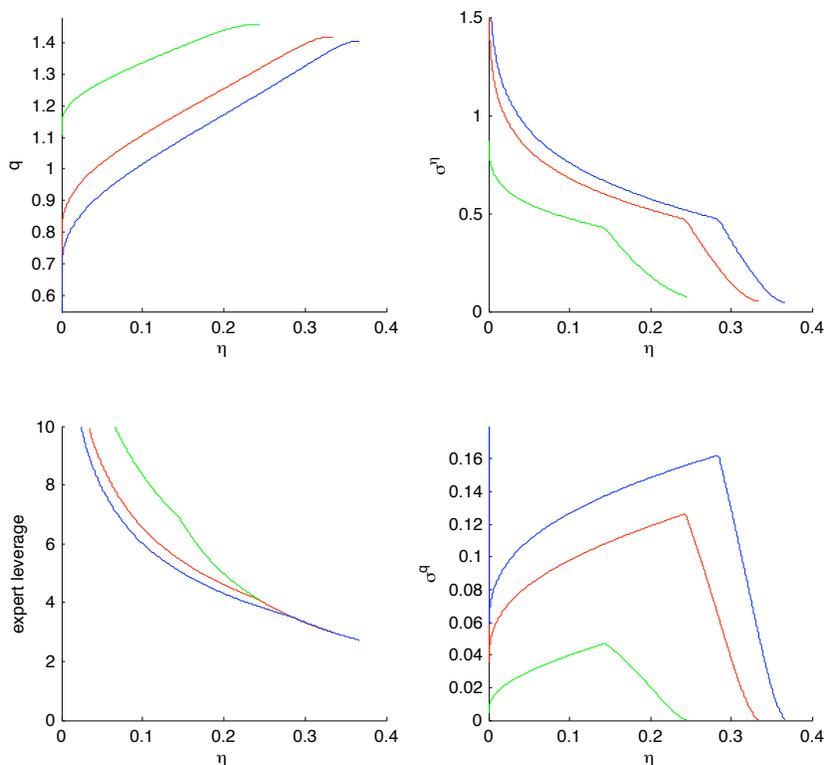


Figure 2: Market illiquidity: $\underline{a} = 2\%$ (blue) 5% (red) and 11% (green).

- 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 households⁹

The reason for the volatility paradox is that payouts, i.e. the location of η^* , and leverage are endogenous.

Illiquidity. Brunnermeier, Eisenbach and Sannikov (2012) distinguish three types of illiquidity. First, technological illiquidity, which is captured in this model by the concavity of function Φ . Second, market illiquidity, which is captured by the differences $a - \underline{a}$ and $\delta - \underline{\delta}$, which drive the difference

⁹However, as σ decreases, the system spends less time in the depressed region, so some measures of stability improve.

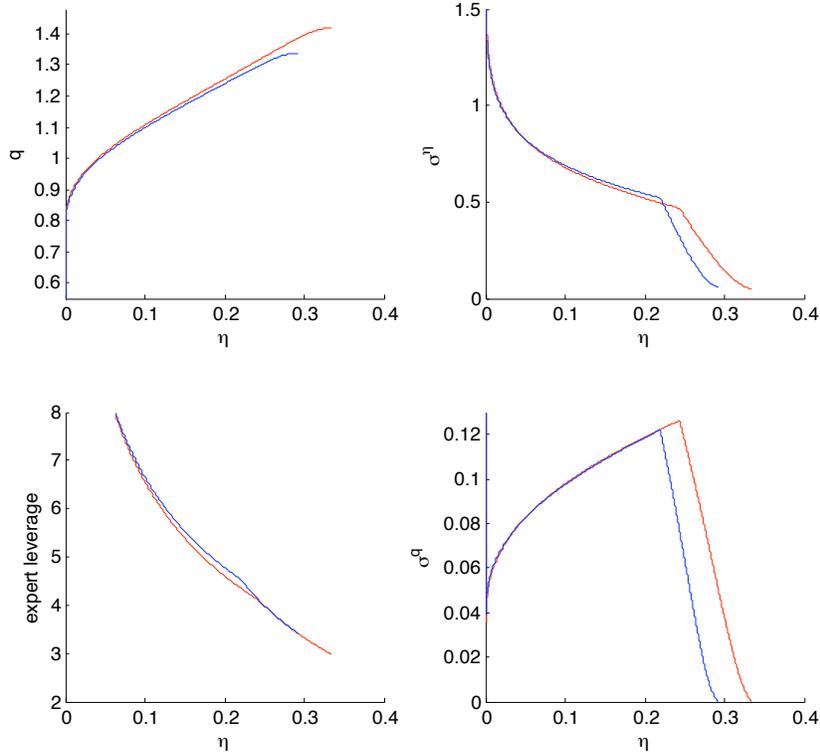


Figure 3: Technological illiquidity: $\iota = \Phi + 5\Phi^2$ (red) and $\Phi + 20\Phi^2$ (blue).

between the first-best (expert) and second-best valuations (household) of assets. There is also funding liquidity.

In this model, the main type of illiquidity that drives endogenous risk, and makes the system unstable, is the *market* illiquidity. Figure 2 presents a comparative static for our model for $\sigma = 2.5\%$ and the same other parameters as before (except \underline{a}).

The figure confirms our expectations: higher illiquidity results in higher endogenous risk, higher earnings retention and lower leverage. What is striking is the magnitude of effect: for the same level of exogenous risk of 2.5%, endogenous risk varies anywhere between 4% and 16% in this example.

In contrast, *technological* illiquidity matters a lot less. If we boost the adjustment cost parameter by a factor of 4, so that it takes investment of $\Phi + 20\Phi^2$ to generate new capital at rate Φ , the equilibrium changes much

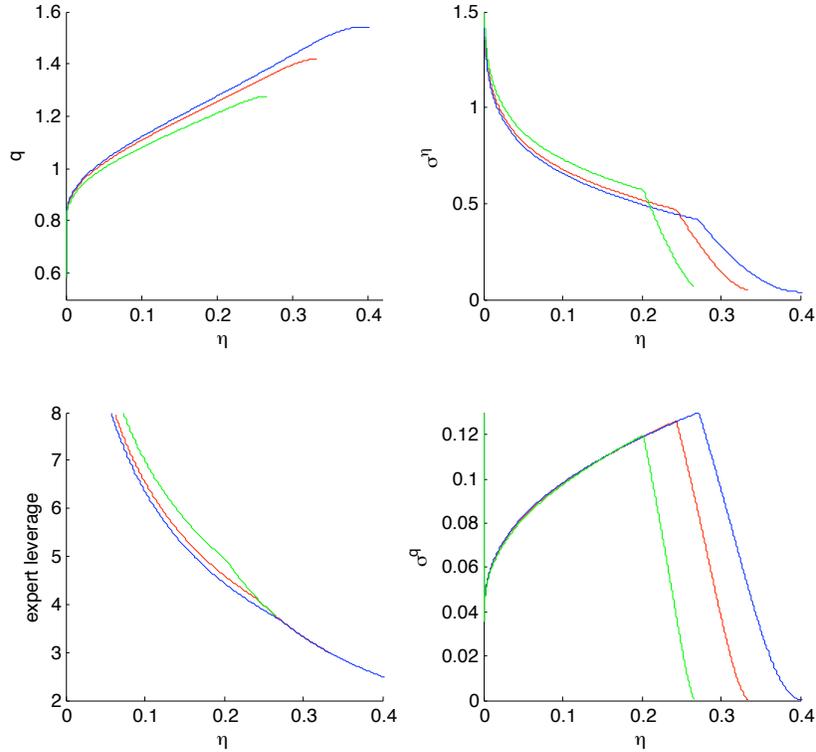


Figure 4: Experts' impatience: $\rho = 5.2\%$ (blue), 6% (red) and 8% (green).

less. In fact, endogenous risk *decreases* with higher adjustment costs instead of increasing. The reason is that lower adjustment costs lead to higher investment and higher prices in booms. The more room there is for prices to fall, the greater the endogenous risk.

Experts' Impatience. We assume that experts are less patient than households, i.e. $\rho > r$. If $\rho = r$, then in the long run η_t would converge to 1.

We do not have a perfect economic interpretation of the parameter ρ , but nevertheless it is interesting to do comparative statics on ρ . One would think that as ρ decreases towards r , the crises become less volatile and less frequent as experts accumulate more wealth.

That is only partially true! Crises can definitely become more volatile as ρ gets closer to r . As Figure 4 illustrates, lower ρ may lead to higher prices in booms as experts accumulate more wealth. As prices have more room to

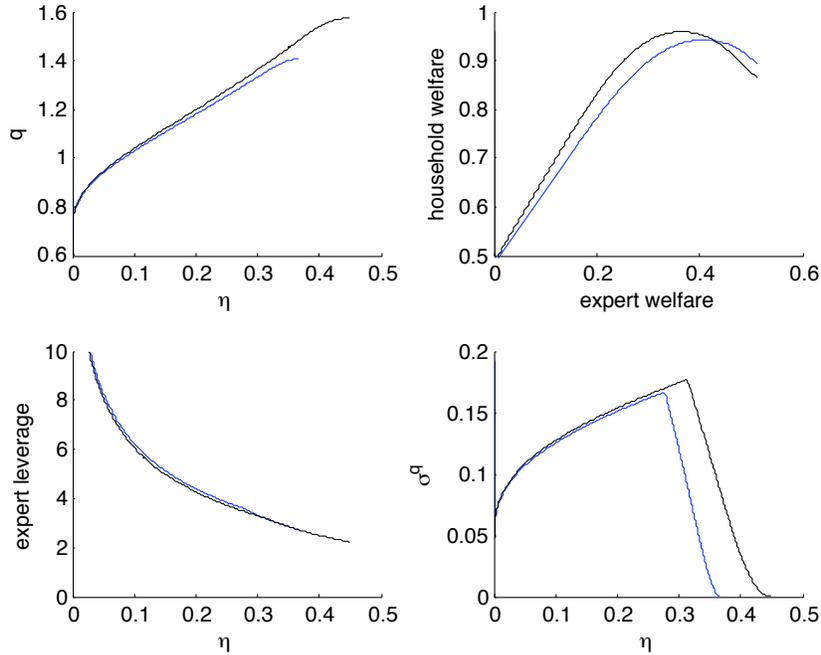


Figure 5: A constraint on payouts (black) vs. the baseline setting (blue).

fall, endogenous risk in crises rises. Financial frictions certainly get weaker as ρ gets closer to r , but that fuels asset price booms (or “bubbles”).

Policies.

When a policy is imposed, it affects some of the equilibrium equations, and thus equilibrium dynamics. To predict the effect of a policy, it is useful to ask the following questions. How does the policy affect the equilibrium payout rate and leverage? How does it affect asset allocation? How does it affect asset prices and endogenous risk?

While we can have good intuition about the effects of various policies, very often formal analysis reveals unintended consequences, which can be justified ex-post.

A Restriction on Payouts. Figure 6 shows what happens when experts are not allowed to consume wealth until η_t reaches the level of 0.45 (instead of $\eta^* = 0.365$ in equilibrium).

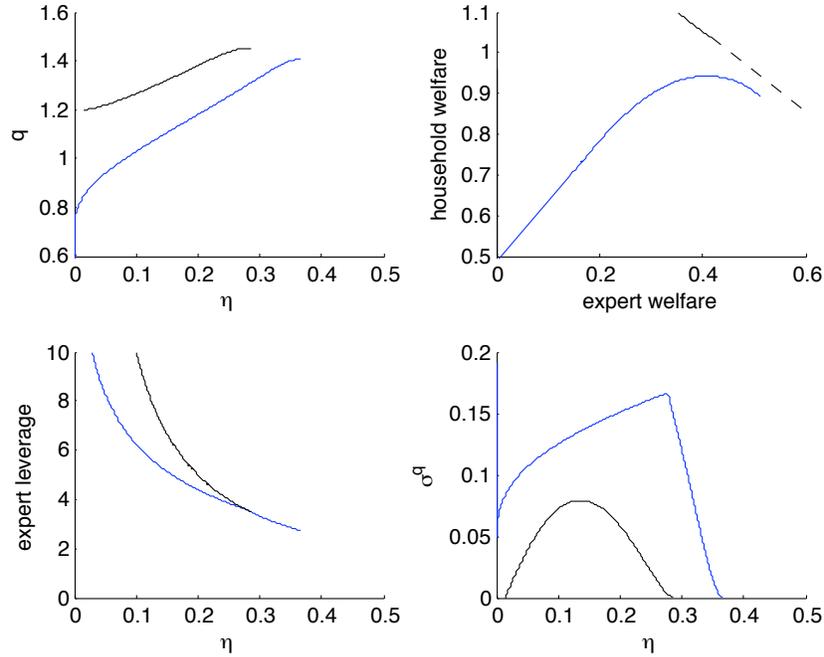


Figure 6: A price floor at $q = 1.2$ (black) vs. the baseline setting (blue).

This policy fuels a boom in prices, increases endogenous risk during crisis episodes and reduces welfare within the model. However, it also reduces the frequency of crises, and so could improve welfare for reasons outside the model, if there are spillovers from crisis episodes to the economy.

A Price Floor. If a government/central bank can support the price of illiquid assets, it can potentially increase welfare significantly in environments with high endogenous risk. Moreover, if exogenous risk is low, then the probability of having to apply the policy is low and thus the cost of the policy is low as well.

Note, however, that such a policy leads to higher leverage and earlier payouts in equilibrium. See Figure 7.