FIN 3005 Asset Pricing/Macrofinans

Fall/Høst 2017

Final exam

Answer Question 1 and two of the remaining four questions.
All questions count equally

All the questions can be answered with a lot of mathematical detail. Although knowledge and understanding of the math should be rewarded, the main question is whether the candidate shows understanding of the main issue in each question. Below, the main points are italicized.

Text in English

1. Define the equity premium and the equity premium puzzle. Discuss the extent to which this puzzle can be explained within the Lucas Tree Model under the assumption of expected power utility. Use mathematical formulations as appropriate.

The equity premium is the difference between the expected return on equity and the return on a risk-free asset. Basic finance theory predicts that this premium is positive because, under risk aversion, equity returns tend to be low when the marginal utility is high. Investors will thus be reluctant to invest in equity unless they can expect a premium:

$$E_t R_{e,t+1} - R_{f,t+1} = -\frac{cov_t\left[R_{e,t+1}, U'(c_{t+1})\right]}{E_t U'(c_{t+1})}.$$

When estimated with real data, the equity premium usually turns out to be positive, as predicted. However, it usually is estimated as much larger than the theory would predict with plausible values of risk aversion, as estimated in studies on micro data. The puzzle is thus a quantitative one, not qualitative.

The Lucas Tree model allows analysis of the equity premium within a simple macroequilibrium framework. In this model, owning trees and harvesting the perishable fruits is the only source of income and consumption. Trees are equity, and the model allows computation of the rate of return on this asset.

The expected power utility is

$$U_t = E_t \sum_{s=0}^{\infty} \beta^s \frac{c_{t+s}^{1-\gamma}}{1-\gamma}, 0 < \beta < 1, \gamma > 0.$$

The Euler equation becomes

$$U'(c_t) = \beta E_t\left[R_{e,t+1} U'(c_{t+1})\right].$$

Those who can reproduce the following derivation should be rewarded.

Letting y_t denote the dividend, i.e. the fruit harvest at time t, the rate of return on equity is defined as

$$R_{e,t+1} = \frac{y_{t+1} + p_{t+1}}{p_t}.$$

Under the conjecture that the asset price is proportional to the dividend, a conjecture that turns out to be correct, we find

$$R_{e,t+1} = \frac{z_{t+1}}{\beta E_t[(z_{t+1})(x_{t+1})^{-\gamma}]},$$

where z_t and x_t are the gross growth rates of dividends and consumption, respectively. Similarly,

$$R_{f,t+1} = \frac{1}{\beta E_t x_{t+1}^{-\gamma}}.$$

Under the assumption that z_t and x_t are i.i.d. log-normally distributed, we now find

$$\ln E_t R_{e,t+1} - \ln R_{f,t+1} = \gamma \sigma_{xz}.$$

In equilibrium, consumption must equal dividends, so that

$$\ln E_t R_{e,t+1} - \ln R_{f,t+1} = \gamma \sigma_x^2.$$

Unfortunately, when empirical data are plugged in for the two rates of return and the variance of the log of the gross growth rate of consumption, satisfaction of this equation requires an unrealistically large value of γ*.*

Thus, the Lucas Tree Model does not solve the equity premium puzzle. However, it points towards possible avenues for further exploration. Research can, on the one hand look for alternative specification of consumer preferences that could imply more risk aversion, i.e. a larger γ. This is what is done in habit formation models. On the other hand, it can look for a higher σ_x^2, i.e. arguments that would suggest that risk is higher than in a model of simple i.i.d. distributions. Examples of such models include rare disasters (Barro) or long-term risk (Bansal and Yaron).

2. Define Epstein-Zin preferences. How do they differ from preferences defined by expected power utility? Do Epstein-Zin preferences make it easier to explain the equity premium puzzle and/or the risk-free rate puzzle? If so, how?

Epstein-Zin preferences distinguish between risk aversion and aversion to planned variations in consumption as expressed by the reciprocal of the elasticity of intertemporal substitution. They are defined recursively by the following kind of value function:

$$V_t(A_t) = \max_{c_t,\omega_t} \left\{ (1-\beta) c_t^{\rho} + \beta \left[E_t V_{t+1}(A_{t+1})^{\alpha} \right]^{\rho/\alpha} \right\}^{1/\rho}.$$

Technically, it is a two-level CES combination. The first level specifies the elasticity of intertemporal substitution, $\sigma \equiv 1/(1-\rho)$*. The second level describes risk aversion with the RRA parameter* $\gamma \equiv 1-\alpha$*.* Using the transformed parameters ρ and α makes the model easier to manipulate than with the underlying parameters σ and γ. Power expected utility is a special case of Epstein-Zin preferences, where $\gamma = 1/\sigma$, i.e. $\alpha = \rho$.

After considerable manipulation, the Euler equation for Epstein-Zin preferences turns out to be

$$1 = E_t\left(k_{t+1} R_{e,t+1}\right),$$

where

$$k_{t+1} = \beta^{\alpha/\rho} (c_{t+1}/c_t)^{\alpha(1-1/\rho)} R_{p,t+1}^{\alpha/\rho - 1}.$$

Students who have memorized this formula should be rewarded.

Also,

$$\ln R_{f,t+1} = -\ln E_t k_{t+1}.$$

Implementing this into the Lucas Tree Model gives the same equation for the equity premium as with expected power utility:

$$\ln E_t R_{e,t+1} - R_{f,t+1} = \gamma \sigma_x^2.$$

Thus, within this model, Epstein-Zin preferences provides no help in explaining the equity premium puzzle. *However, it does make it clear that the parameter γ in the above formula reflects risk aversion only and has nothing to do with intertemporal substitution. Furthermore, it helps explain the risk-free rate puzzl*e because the formula for the risk-free rate now becomes

$$\ln R_{f,t+1} = -\ln \beta + (1/\sigma)\mu_x - \frac{1}{2}[\gamma - (1-\gamma)/\sigma)]\sigma_x^2.$$

In particular, the trend consumption growth rate μ_x is now multiplied by the reciprocal of the elasticity of intertemporal substitution, $1/\sigma$, rather than the risk aversion parameter γ. Thus, the more complex preference ordering helps explain the observed low values of the risk-free rate provided the aversion to planned changes in consumption, $1/\sigma$, is lower than the aversion γ to stochastic changes.

However, if we go beyond the i.i.d assumptions of the Lucas Tree Model, Epstein-Zin preferences can even help explain the equity premium puzzle. This is discussed further in the answer to the next question.

3. Rare disasters and long-term risks have been advanced as explanations of the equity premium puzzle. Describe and explain the similarities and differences between these two approaches and how successful they have been in explaining the equity premium puzzle.

Barro's model of rare disasters and Bansal and Yaron's model of long-term risk both seek to explain the equity premium puzzle by claiming that real-world risk is greater than what is captured by the i.i.d. log-normal specification of consumption growth or equity return. They thus focus on the severity of market and macro risk rather than on the specification of consumer or investor preferences. That said, however, the model in Bansal and Yaron depends crucially on their use of Epstein-Zin preferences, which allows them to assume that risk aversion exceeds the aversion to planned variations in consumption. This is important because the main term in their formula for the equity premium is multiplied by the factor $\gamma - 1/\sigma$.

The two approaches also differ in the way that they specify deviations from the i.i.d. log-normal specification of consumption growth. *Barro keeps the i.i.d. log-normal specification for normal times, but adds a probability that, sometimes, a rare disaster may strike. The effect of this assumption is to fatten the tails of the unconditional probability distribution of consumption growth.* With the fattened tail, the unconditional distribution is not normal. It is, however, i.i.d, which means that it can be implemented within the Lucas Tree Model. Using then the same formulas for the equity return and the riskless return, the equity premium becomes (in terms of net returns in continuous time), approximately

$$r_e - r_f = \gamma\sigma^2 + \ pEb[(1-b)^{-\gamma} - 1].$$

Here, the first term on the right is the same as in the lognormal model. The second term reflects the rare disasters in three ways: (i) the probability of a rare disaster, p, (ii) the likely *magnitude of a rare disaster, once it occurs,* Eb*, and finally (iii) the difference in marginal utility between a disaster state and a normal state,* $(1-b)^{-\gamma}-1$.

Bansal and Yaron take a different approach by allowing for changes in consumption growth to display persistence and thus possibly cast long shadows. Specifically, consumption growth takes the form

$$g_{t+1}=x_t+\eta_{t+1},$$

where x_t is the persistent part, specified as an AR(1), and η_{t+1} is white noise. Then the unconditional variance of g_t has the order of magnitude of $1/(1-\rho^2)$, which may be large if the process is highly persistent. However, estimation of the persistence parameter ρ by standard econometric techniques tend to provide downward biased estimates, and the bias is larger the greater the persistence. Thus, Bansal and Yaron argue that the observed magnitude of the equity premium may provide more reliable information of the amount of persistence than direct estimates.

This specification of consumption growth is obviously not i.i.d. and so cannot be implemented in the Lucas Tree Model. Bansal and Yaron instead rely on the linear approximation of Campbell and Shller for the log of the gross equity return:

$$r_{e,t+1}=\kappa_0+\kappa_1 z_{t+1}-z_t+g_{t+1}.$$

where z_t is the log of the stock price – dividend ratio, g_t the dividend growth rate, and κ_1 positive and just a little less than one. For the simplest case, dividends are assumed to equal consumption. In an extension, the two are modelled separately.

An additional assumption is that the log stock price – consumption ratio (or, the log stock price – dividend ratio) can be approximated as a linear function of the persistent component of consumption growth. When that is substituted into the Euler equation for the Epstein-Zin preferences, we obtain a function that depends on the current value of the persistent component of consumption/dividend growth. Because the Euler equation must hold for all values of this component, its coefficient must be zero. This condition can then be used to solve for the slope coefficient of the persistent component in the stock price – consumption/dividend ratio. Given this result, the next trick is to be able to see that *the log of the equity premium must equal the negative of the conditional covariance of the log gross equity return and the log stochastic discount factor:*

$$\ln E_t R_{e,t+1} - \ln R_{f,t+1} = -cov_t(m_{t+1}, r_{e,t+1}).$$

For the case where consumption equals dividends, we then find

$$\ln E_t R_{e,t+1} - \ln R_{f,t+1} = (\gamma - 1/\sigma)(1 - 1/\sigma)\left(\frac{\kappa_1}{1 - \rho\kappa_1}\right)^2 \sigma_\varepsilon^2 + \gamma\sigma_\eta^2.$$

The last term on the right, $\gamma\sigma_\eta^2$, is the answer we would have gotten with i.i.d. consumption growth and expected power utility. Here, we see that the log equity premium exceeds this term if $\gamma > 1/\sigma$ and $\sigma > 1$. Thus, the elasticity of intertemporal substitution must be large enough for the substitution effect to dominate the income effect for comparative-static changes in the risk-free rate of return. It must also be large enough so that the aversion to stochastic variations in consumption exceeds the aversion to planned variations. Given these conditions, the equity premium can be quite large because the denominator $1 - \rho\kappa_1$ is likely to be small when consumption growth is persistent.

For the case where consumption and dividends are modeled separately, this result is modified to

$$\ln E_t R_{m,t+1} - \ln R_{f,t+1} = (\gamma - 1/\sigma)\,(\phi - 1/\sigma)\left(\frac{\kappa_1}{1 - \rho\kappa_1}\right)\left(\frac{\nu_1}{1 - \nu_1\rho}\right)\sigma_\varepsilon^2 + \gamma\psi\sigma_\eta^2.$$

Here, the parameter ϕ is the response of dividends to an innovation in the persistent part of consumption growth, assumed greater than unity. Thus, with this extension, it is not necessary for the elasticity of intertemporal substitution to exceed unity.

Quantitatively, Bansal and Yaron seem to have come closest to explaining the equity premium puzzle. However, the assumptions needed in terms of parameter values for γ and σ make their theory somewhat vulnerable to criticism.

4. Physical assets can be used as inputs to production as well as collateral to secure debt. Describe and discuss how this dual role may exacerbate macroeconomic fluctuations. Use mathematical formulations as needed to support your narrative. Why is moral hazard important in this context?

This question refers to Kiyotaki and Moore's model of credit cycles. They present a simple model of fruit-growing farmers and gatherers. They all use capital, best interpreted as land. Farmers can produce with constant returns, but their output depends on their continued effort during the growing season. That creates a temptation for them to "stiff" their creditors by diverting the borrowed funds for personal consumption. For the creditors, this is a problem of moral hazard. As a result, they will ration credit to farmers, limiting it to the resale value of the farmer's land in case he or she diverts. Gathers also grow fruit. Their production does not depend on their effort during the growing season, so for them, the moral hazard problem does not arise. Their credit is not constrained. However, their production is subject to decreasing returns, and their marginal product is lower than that of farmers for low levels of production. In equilibrium, they thus choose to be creditors, whereas the farmers leverage their capital to buy and hold as much land as possible given their credit constraint.

In equilibrium, farmers will hold capital in the amount

$$K_t = \frac{1}{u_t}[(a + q_t)K_{t-1} - RB_{t-1}],$$

where $u_t = q_t - q_{t+1}/R$ is the user cost of capital (also the down payment needed for buying a unit of capita), q_t the price of capital, and R the gross interest rate. a is the (constant) amount of marketable fruit output per unit of capital, and B_{t-1} is the farmer's one-period debt. The latter is limited to

$$B_t = \frac{1}{R} q_{t+1} K_t.$$

Finally, the user cost is determined as the gatherers' marginal product when they hold all the capital that the farmers are not able to buy:

$$u_t = \frac{1}{R} G'(\bar{K} - K).$$

This equilibrium is inefficient because the farmers' marginal product exceeds that of the gatherers, and yet they aren't allow to take over more of the fruit growing because of the capital constraint. The graph in the article illustrates this nicely.

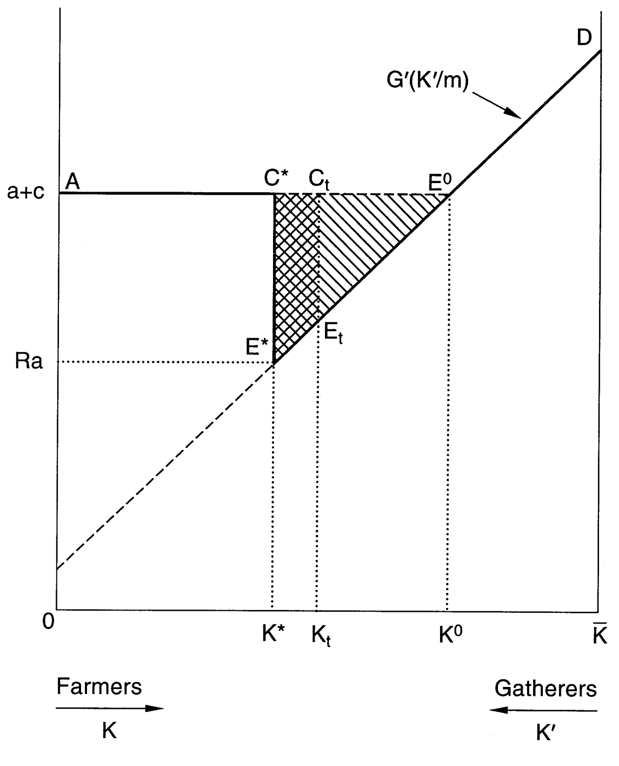

The interesting point in this analysis is of the effects of a temporary productivity shock for farmers. A positive such shock will raise the value of capital by making it more productive. T*his raises farmers' net worth, so they can buy more land, even at the higher price. This, in turn, has two effects. On the one hand, aggregate productivity is increased even more because more of production is now carried out by farmers. On the other hand, even though the productivity improvement is temporary, farmers will want to keep some of their new capital in future periods as well. Thus, future prices of capital are expected to rise, which in turn makes the current value of capital to rise even further.* The graph in the article illustrates this process for the opposite case of a temporary drop in productivity.

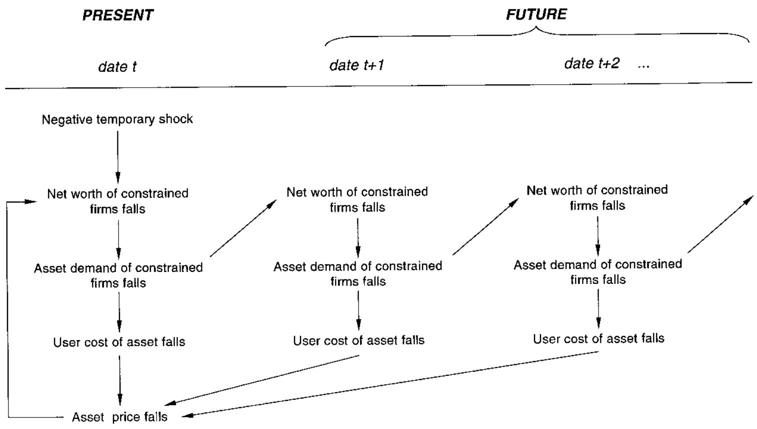

The double use of capital as an input to production and as collateral for debt thus creates *two multipliers, one static and one dynamic.* The dynamic part is much more powerful. The full effects, in terms of log differences of changes because of the shock Δ, are:

$$\hat{q}_t = \frac{1}{\eta}\Delta$$

$$\hat{K}_t = \frac{1}{1 + 1/\eta}\left(1 + \frac{R}{R-1}\frac{1}{\eta}\right)\Delta.$$

Students don't necessarily need to memorize these formulae exactly, but should definitely explain the relative orders of magnitude of the static and the dynamic components.

η is the elasticity of user cost to a change in farmers' capital holdings, whose value is of the order of magnitude of 1. In the first equation, we then see that the value of land raises by about the same amount as the productivity shock, even though the latter is temporary. And in the second, we see that the term $\frac{R}{R-1}$ multiplies up considerably the amount of land that farmers then buy from the gatherers and thus cultivate more efficiently. Had only the static effect been realized, the results would have been the much more modest:

$$\hat{q}_t = \frac{R}{R-1}\frac{1}{\eta}\Delta, \qquad \hat{K}_t = \Delta.$$

5. Securitization has been advanced as a financial innovation that improves the efficiency of risk allocation. Is this claim justified? Why or why not? And what are the implications of securitization for aggregate risk?

This is the issue analyzed in the article by Gennaioli, Shleifer, and Vishny. Their setup consists of financial intermediaries, i.e. investment banks, or banks for short, and investors, i.e. households, mutual funds, pension funds, etc., households for short. Whereas households are infinitely risk averse (i.e. minimax optimizers), the banks are completely risk neutral. Banks have access to investment in projects, that can be thought of as mortgages. Of these, a set, limited to a size of 1, are high quality (H) with a high and riskless gross return of 1. In additions, banks can also invest in low-quality (L) projects. These are subject to idiosyncratic as well as aggregate risk. The idiosyncratic risk is modeled such that a share π_ω og such projects will succeed and yield a gross return of A if the aggregate state is ω. The rest yield zero. The aggregate state can be either one of growth, a downturn, or a recession, such that $\pi_g > \pi_d > \pi_r$.

In addition, *banks can trade their projects amongst each other. This is the form of securitization in the model. By buying a diverse set of other banks' projects, a bank can diversify away the idiosyncratic risk on that part of its portfolio. The bank does not do this to reduce its own risk, for it is risk neutral. However, trading projects against a diversified portfolio of other banks' projects enables it to promise risk-free returns to its clients, the households, whose only investment consists of bank deposits.*

Perhaps needless to say, there will be no securitization of high-quality projects, only of low-quality ones.

The key constraint in this model is the limit on the amount of deposits from infinitely risk-averse clients that a bank can credibly promise to repay at maturity:

$$rD \leq RI_H + \pi_r AT_L.$$

Here, the first term is the bank's gross earnings from high-quality projects and the second its return on its diversified portfolio of purchased assets, a share π_r of which will earn the return A and the rest zero. Because all banks are symmetric, we will also have $T_L = S_L$, where S_L denotes projects sold.

The graph illustrates the equilibrium rates of return as a function of demand, indicated by household wealth w. If this demand is low enough so that the sum of it and banks' own capital w_{int} is less than or equal to 1, the available supply of high-quality projects, deposits well earn the return of those projects. Then there is no securitization. If demand is greater than this, the return on deposits falls to the expected return on the marginal, low-quality project, i.e. $E_\omega(\pi_\omega)A$. If demand is less than R divided by this return, banks will still have enough return on their high-quality assets to be able to repay clients even in a recession, which is the worst case of aggregate risk. If beyond rises above this level, they will, however, need to do some securitization in order to be able to honor their obligations without risk for the households. According to the funding constraint, they can do this an still credibly promise the same rate or return as long as

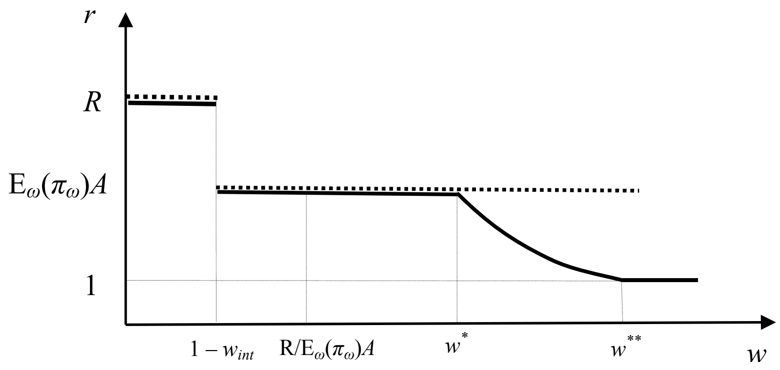

$$rD = rw = E_\omega(\pi_\omega)A \leq RI_H + \pi_r AT_L = R + \pi_r AS_L = R + \pi_r A(w_{int} + w - 1).$$

This constraint holds as long as

$$w \leq \frac{R + \pi_r A(w_{int} - 1)}{E_\omega(\pi_\omega)A - \pi_r} \equiv w^*.$$

For $w > w^*$, there must be full securitization. And the same constraint implies

$$r = \frac{R + \pi_r A(w_{int} - 1)}{w} + \pi_r A < \frac{R + \pi_r A(w_{int} - 1)}{w^*} + \pi_r A.$$

At $w = w^{**}$, the gains from securitization are completely exhausted, and the net return on deposits is competed down to zero.

Provided some securitization occurs, the outcome for the lucky banks becomes the sum of the return on their high-quality investments, the return on their non-securitized low-quality investments, and the return on their purchased securities:

$$\begin{aligned} &RI_H + A(I_L - S_L) + \pi_\omega AS_L - rD \\ &= RI_H + A(I_L - S_L) + \pi_\omega AS_L - RI_H - \pi_r AS_L \\ &= A(I_L - S_L) + (\pi_\omega - \pi_r)AS_L. \end{aligned}$$

Both of the last two terms are obviously non-negative, so the lucky banks survive for sure, even in the worst aggregate case. For the unlucky banks, we similarly find

$$RI_H + 0 \cdot (I_L - S_L) + \pi_\omega AS_L - rD = (\pi_\omega - \pi_r)AS_L \geq 0.$$

So, *even the unlucky banks survive. In fact, they even make a positive profit except in the recession case, when they break even.*

These results demonstrate the microeconomic advantage of securitization. By promoting diversification, it makes sure that risk is allocated perfectly efficiently.

Unfortunately, this result is changed if agents ignore the worst aggregate risk, i.e. they act as if the recession case cannot occur. Then, things will look very good as long as this case is not realized, in the sense that securitization will be higher and the rate of return on deposits higher as well. This is illustrated in the graph.

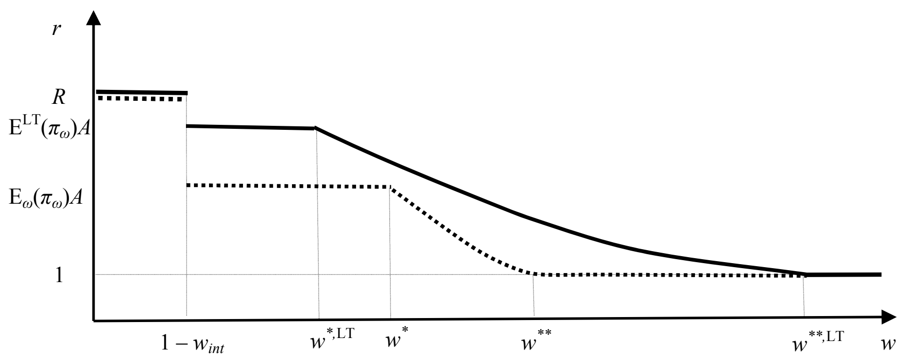

The main reason that things look better in this case is the fact that the banks' funding constraint now will be less strict because households ignore the risk of recession:

$$rD \leq RI_H + \pi_d AT_L.$$

However, the outcome will not look as good if the recession case in fact occurs. Then, the result for the unlucky banks will be $(\pi_r - \pi_d)AS_L < 0$*, i.e. they will go bankrupt, and their depositors won't get their promised return. The lucky banks will do better with* $A(I_L - S_L) + (\pi_r - \pi_d)AS_L$*. This expression is positive if*

$$\frac{I_L}{S_l} > 1 + \pi_d - \pi_r,$$

i.e. if securitization has not gone too far. However, with full securitization, even they go bankrupt because then $A(I_L - S_L) + (\pi_r - \pi_d)AS_L = A \cdot 0 + (\pi_r - \pi_d)AS_L < 0$*.*

This is the other main point of this analysis. *Securitization makes sure that all agents are equally exposed to aggregate risks. And if aggregate risk is ignored, the entire system may fail.*

The system becomes more efficient, but also more fragile.