Uncertainty Shocks and Balance Sheet Recessions

Sebastian Di Tella*
Stanford GSB

First version November 2012
This version August, 2013

Abstract

This paper investigates the origin and propagation of balance sheet recessions in a general equilibrium model with financial frictions. I first show that in standard models driven by TFP shocks, the balance sheet channel completely disappears when agents are allowed to write contracts on the aggregate state of the economy. Optimal contracts sever the link between leverage and aggregate risk sharing, eliminating the concentration of aggregate risk that drives balance sheet recessions. I then show how the type of aggregate shock that hits the economy can help explain the concentration of aggregate risk. In particular, I show that uncertainty shocks can drive balance sheet recessions and "flight to quality" events, even when contracts can be written on the aggregate state of the economy. Finally, I explore implications for financial regulation.
JEL Codes: E32, E44, G1, G12

*I am grateful to my advisors Daron Acemoglu, Guido Lorenzoni, and Ivan Werning for their invaluable guidance. I also thank George-Marios Angeletos, Vladimir Asriyan, Ricardo Caballero, Marco Di Maggio, David Gamarnik, Veronica Guerrieri, Leonid Kogan, Andrey Malenko, William Mullins, Juan Passadore, Mercedes Politi, Felipe Severino, Rob Townsend, Victoria Vanasco, Xiao Yu Wang, Juan Pablo Xandri, Luis Zermeño, and seminar participants at the MIT Macroeconomics Lunch for useful suggestions and comments. Any remaining errors are my own.
Correspondence: Stanford GSB, 655 Knight Way, Stanford, CA 94305, email: sditella@stanford.edu

1 Introduction

The recent financial crisis has underscored the importance of the financial system in the transmission and amplification of aggregate shocks. During normal times, the financial system helps allocate resources to their most productive use, and provides liquidity and risk sharing services to the economy. During crises, however, excessive exposure to aggregate risk by leveraged agents can lead to *balance sheet recessions*.[1] Small shocks will be amplified when these leveraged agents lose net worth and become less willing or able to hold assets, depressing asset prices and growth. And since it takes time for balance sheets to be rebuilt, transitory shocks can become persistent slumps. While we have a good understanding of why balance sheets matter in an economy with financial frictions, we don't have a good explanation for why agents are so exposed to aggregate risk in the first place.[2] The answer to this question is important not only for understanding the balance sheet channel, but also for the design of effective financial regulation. In this paper I show that uncertainty shocks can help explain the apparently excessive exposure to aggregate risk that drives balance sheet recessions.

In order to understand agents' aggregate risk-sharing decisions, I derive financial frictions from a standard moral hazard problem. I allow them to write contracts on all observable variables, and I find that the type of aggregate shock hitting the economy takes on a prominent role. The first contribution of this paper is to show that in standard models of balance sheet recessions driven by TFP shocks, the balance sheet channel completely disappears when agents are allowed to write contracts contingent on the observable aggregate state of the economy. Optimal contracts break the link between leverage and aggregate risk sharing, and eliminate the excessive exposure to aggregate risk that drives balance sheet recessions. As a result, balance sheets play no role in the transmission and amplification of aggregate shocks. Furthermore, these contracts are simple to implement using standard financial instruments such as equity and a market index. In fact, the balance sheet channel vanishes as long as agents can trade a simple market index. The intuition behind this result goes beyond the particular environment in this model.

The second contribution is to show that, in contrast to standard TFP shocks, uncertainty shocks can create balance sheet recessions, even when contracts can be written on the aggregate state of the economy. I introduce an aggregate uncertainty shock that increases idiosyncratic risk in the economy. With financial frictions, an increase in idiosyncratic risk depresses asset prices and growth, and generates an endogenous hedging motive that induces more productive (leveraged) agents to take on aggregate risk ex-ante. Weak balance sheets therefore amplify the effects of the uncertainty shock, further depressing asset prices and growth. This balance sheet channel in turn amplifies the hedging motive, inducing agents to take even more aggregate risk ex-ante, in a two-way feedback

[1] The idea of balance sheet recessions goes back to Fisher (1933) and, more recently, Bernanke and Gertler (1989) and Kiyotaki and Moore (1997). Several papers make the empirical case for balance sheet effects, such as Sraer et al. (2011), Adrian et al. (2011) and Gabaix et al. (2007).

[2] In standard models of balance sheet recessions such as Bernanke and Gertler (1989) and Kiyotaki and Moore (1997), or more recently Brunnermeier and Sannikov (2012), He and Krishnamurthy (2011), or Kiyotaki et al. (2011) agents face ad-hoc constraints on their ability to share aggregate risk. However, Begenau et al. (2013) show banks, for example, have large trading positions on derivatives that allow them to insure against the aggregate risk in their traditional business, but use them instead to *amplify* their exposure to aggregate risk.

loop. In addition, an increase in *idiosyncratic* risk leads to an endogenous increase in *aggregate* risk, and triggers a "flight to quality" event with low interest rates and high risk premia.[3]

I use a continuous-time growth model similar to the Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011) models of financial crises (BS and HK respectively). I derive financial frictions from a moral hazard problem, and allow agents to write contracts on all observable variables.[4] There are two types of agents: experts who can trade and use capital to produce, and consumers who finance them. Experts can continuously trade capital, which is exposed to both aggregate and (expert-specific) idiosyncratic Brownian TFP shocks. They want to raise funds from consumers and share risk with them, but they face a moral hazard problem that imposes a "skin in the game" constraint: experts must keep a fraction of their equity to deter them from diverting funds to a private account. This limits their ability to share idiosyncratic risk, and makes leverage costly. The more capital an expert buys, the more idiosyncratic risk he must carry on his balance sheet. Experts will therefore require a higher excess return on capital when idiosyncratic risk is high and their balance sheets are weak.

When contracts cannot be written on the aggregate state of the economy, experts are mechanically exposed to aggregate risk through the capital they hold, and any aggregate shock that depresses the value of assets will have a large impact on their leveraged balance sheets. In contrast, when contracts can be written on the aggregate state of the economy, the decision of how much capital to buy (leverage) is separated from aggregate risk sharing, and optimal contracts hedge the (endogenously) stochastic investment possibility sets provided by the market. In equilibrium, aggregate risk sharing is governed by the hedging motive of experts relative to consumers. Brownian TFP shocks don't affect the investment possibility sets of experts and consumers, so they share this aggregate risk proportionally to their wealth. In equilibrium, TFP shocks have only a direct impact on output, but are not amplified through balance sheets and do not affect the price of capital, growth rate of the economy, or the financial market.

In contrast to Brownian TFP shocks, aggregate uncertainty shocks that increase idiosyncratic risk for all experts create an endogenous hedging motive that induces experts to choose a large exposure to aggregate risk. The intuition is as follows. Downturns are periods of high uncertainty, with endogenously depressed asset prices and high risk premia. Experts who invest in these assets and receive the risk premia have relatively better investment opportunities during downturns, and get more utility per dollar compared to consumers. On the one hand, this creates a *substitution effect*: if experts are risk-neutral, they will prefer to have more net worth during downturns in order to get more "bang for the buck". This effect works against the balance sheet channel, since it induces experts to insure against aggregate risk. On the other hand, experts require more net

[3]Empirically, idiosyncratic risk rises sharply during downturns as Bloom et al. (2012) document: during the financial crisis in 2008-2009, plant level TFP shocks increased in variance by 76%, while output growth dispersion increased by 152%. An increase in idiosyncratic risk could also reflect greater disagreement over the value of assets (Simsek (2013)) or an increased interest in acquiring information about assets (Gorton and Ordoñez (2013)).

[4]Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011) also derive financial frictions from a similar contracting problem, but they impose constraints on the contract space that limit agents ability to share aggregate risk.

worth during booms in order to achieve any given level of utility. This creates a *wealth effect*: risk averse experts will prefer to have more net worth during booms. I argue the empirically relevant case is the one in which the wealth effect dominates[5] the substitution effect, and drives the balance sheet amplification channel.

Investment opportunities, however, are endogenous and depend on the health of experts' balance sheets. After an uncertainty shock, experts' balance sheets are weak, which reduces their willingness to hold capital and further depresses asset prices and growth. This amplifies the hedging motive, inducing experts to take even more aggregate risk ex-ante. The equilibrium is a fixed point of this two-way feedback between aggregate risk sharing and endogenous hedging motives. The continuous-time setup allows me to characterize it as the solution to a system of partial differential equations, and analyze the full equilibrium dynamics instead of linearizing around a steady state. It also makes results comparable to the asset pricing literature.

These results suggest that the type of aggregate shock hitting the economy can play an important role explaining the concentration of aggregate risk that drives balance sheet recessions. When the wealth effect dominates, experts will choose to face large loses after an aggregate shock that (endogenously) widens the gap in investment opportunities between them and consumers. The same tools presented here can be used to study the effects of other aggregate shocks. In particular, I show that uncertainty shocks are equivalent to an exogenous shock to the degree of moral hazard (how efficient experts are at stealing capital) that translates into an exogenous tightening of financial constraints. The intuition for this result is as follows. In an economy without financial frictions idiosyncratic risk shouldn't matter, since it can be aggregated away. Moral hazard, however, forces agents to keep a fraction of the idiosyncratic risk in their capital. It is immaterial to them whether they must keep a constant fraction of more idiosyncratic risk, or a larger fraction of a constant idiosyncratic risk.

A possible concern with an optimal contracts approach is that they might require very complex and unrealistic financial arrangements. I show this is not the case. Optimal contracts can be implemented in a complete financial market with minimal informational requirements. Experts can be allowed to invest, consume, and manage their portfolios, subject only to an equity constraint. In fact, the TFP-neutrality result does not require the financial market to be complete. It is enough that it spans the aggregate return to capital. A market index of experts' equity accomplishes this. Empirically, Begenau et al. (2013) show banks have access to and actively trade derivatives (interest rate swaps) that allow them to hedge the aggregate risk in their traditional business. Instead of offloading this risk on the market, however, they use them to *increase* their exposure to aggregate risk. This is difficult to reconcile with a theory of incomplete markets, but is consistent with the mechanism in this paper.

Understanding why aggregate risk is concentrated on some agents' balance sheets is important for the design of financial regulation. If markets are incomplete and agents are not able to share

[5]The wealth effect dominates when the coefficient of relative risk aversion is larger than 1 (agents are more risk averse than log).

aggregate risk, it is optimal to facilitate this risk-sharing and eliminate the balance sheet channel, for example through fiscal policy. This is the case in the setting in Brunnermeier and Sannikov (2012) for example. But if agents are able but choose not to share aggregate risk, two issues arise. First, agents might undo policy interventions by taking more aggregate risk.[6] Second, even if it is possible to control their exposure to aggregate risk, they may actually have good reasons to take on so much risk. I show that, while the competitive equilibrium is not constrained efficient, a policy that aims to eliminate the balance sheet channel is not optimal either, and can even be worse than the competitive equilibrium. Optimal financial regulation must take into account the underlying reasons for the concentration of aggregate risk.

Literature Review. This paper fits within the literature on the balance sheet channel going back to the seminal works of Bernanke and Gertler (1989), Kiyotaki and Moore (1997), and Bernanke et al. (1999). It is most closely related to the more recent Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011). The main difference with these papers is that I allow agents to write contracts on all observable variables, including the aggregate state of the economy.

Krishnamurthy (2003) was the first to explore the concentration of aggregate risk and its role in balance sheet recessions when contracts can be written on the aggregate state of the economy. He finds that when agents are able to trade state-contingent assets, the feedback from asset prices to balance sheets disappears. He then shows this feedback reappears when limited commitment on consumers' side is introduced: if consumers also need collateral to credibly promise to make payments during downturns, they might be constrained in their ability to share aggregate risk with experts. This mechanism also appears in Holmstrom and Tirole (1996). The limited commitment on the consumers' side is only binding, however, when experts as a whole need fresh cash infusions from consumers. Typically, debt reductions are enough to provide the necessary aggregate risk sharing, and evade consumers' limited commitment (experts' debt can play the role of collateral for consumers). Rampini and Viswanathan (2010) and Rampini and Viswanathan (2013) also study the concentration of aggregate risk, focusing on the tradeoff between financing and risk-sharing. They show that firms that are severely collateral constrained might forgo insurance in order to have more funds up front for investment. Cooley et al. (2004) show how limited contract enforceability can prevent full aggregate risk sharing. After a positive shock raises entrepreneurs' outside option, their continuation utility must also go up to keep them from walking away. This relaxes the contractual problem going forward and propagates even transitory aggregate shocks. In contrast to these papers, in the setting here agents are able to leverage and share aggregate risk freely, which highlights their incentives for sharing different types of aggregate shocks.

Kiyotaki et al. (2011) also tackle the question of why banks' balance sheets are so highly exposed to aggregate risk, and focus on the debt vs. equity tradeoff for banks, while Adrian and Boyarchenko (2012) build a model of financial crises where experts use long-term debt and face a time-varying leverage constraint. Here, instead, I don't impose an asset structure on agents. Geanakoplos (2009)

[6]This is not due to moral hazard, but rather an attempt to obtain their desired exposure to aggregate risk, similar to the ineffectiveness of mandatory savings on unconstrained agents.

emphasizes the role of heterogeneous beliefs. More optimistic agents place a higher value on assets and are naturally more exposed to aggregate risk. The balance sheet channel in my model, in contrast, does not rely on heterogenous beliefs.[7] Experts take on more aggregate risk in order to take advantage of endogenous investment possibility sets. Myerson (2012), on the other hand, builds a model of credit cycles allowing long-term contracts with a similar moral hazard problem to the one in this paper. In his model the interaction of different generations of bankers can generate endogenous credit cycles, even without aggregate shocks. Shleifer and Vishny (1992) and Diamond and Rajan (2011) look at the liquidation value of assets during fire sales, and Brunnermeier and Pedersen (2009) focus on the endogenous determination of margin constraints.

Several papers make the empirical case for the balance sheet amplification channel. Sraer et al. (2011), for example, use local variation in real estate prices to identify the impact of firm collateral on investment. They find each extra dollar of collateral increases investment by $0.06. Gabaix et al. (2007) provide evidence for balance sheet effects in asset pricing. They show that the marginal investor in mortgage-backed securities is a specialized intermediary, instead of a diversified representative agent. Adrian et al. (2011) use shocks to the leverage of securities broker-dealers to construct an "intermediary SDF" and use it to explain asset returns.

The role of uncertainty in business cycles is explored in Bloom (2009) and, more recently, Bloom et al. (2012), who build a model where higher volatility leads to the postponement of investment and hiring decisions.[8] More closely related to the model in this paper, Christiano et al. (2012) introduce shocks to idiosyncratic risk in a model with financial frictions and incomplete contracts. They fit the model to U.S. data and find this uncertainty shock to be the most important factor driving business cycles[9]. Angeletos (2006) studies the effects of uninsurable idiosyncratic capital risk on aggregate savings. In the asset pricing literature, Campbell et al. (2012) introduce a volatility factor into an ICAPM asset pricing model. They find this volatility factor can help explain the growth-value spread in expected returns. Bansal and Yaron (2004) study aggregate shocks to the growth rate and volatility of the economy, and Bansal et al. (2012) study a dynamic asset pricing model with cash flow, discount rate and volatility shocks. Idiosyncratic risk, in particular, is studied by Campbell et al. (2001). Eggertsson and Krugman (2010), Guerrieri and Lorenzoni (2011) and Buera and Moll (2012) also consider exogenous shocks to financial frictions.

Layout. The rest of the paper is organized as follows. In Section 2 I introduce the setup of the model and the contractual environment. In Section 3 I characterize the equilibrium using a recursive formulation, and study the effects of different types of structural shocks. Section 4 looks at financial regulation. Section 5 concludes.

[7] A related explanation could be built on heterogenous preferences for risk. Less risk averse agents value risky assets more, and also take on more aggregate risk. The mechanism in this paper does not depend on heterogenous preferences either.

[8] On the other hand, Bachmann and Moscarini (2011) argue that causation may run in the opposite direction, with downturns inducing higher risk.

[9] Fernández-Villaverde and Rubio-Ramírez (2010) study the impact of uncertainty shocks in standard macroeconomic models, and Fernández-Villaverde et al. (2011) look at the impact of volatility of international interest rates on small open economies.

2 The model

The model purposefully builds on Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011), adding idiosyncratic risk and general EZ preferences to their framework. As in those papers, I derive financial frictions endogenously from a moral hazard problem. In contrast to those papers, however, contracts can be written on all observable variables.

Technology. Consider an economy populated by two types of agents: "experts" and "consumers", identical in every respect except that experts are able to use capital.[10] There are two goods, consumption and capital. Denote by k_t the aggregate "efficiency units" of capital in the economy, and by $k_{i,t}$ the individual holdings of an expert $i \in [0,1]$, where $t \in [0,\infty)$ is time. An expert can use capital to produce a flow of consumption goods

$$y_{i,t} = (a - \iota(g_{i,t}))\, k_{i,t}$$

The function ι with $\iota' > 0$, $\iota'' > 0$ represents a standard investment technology with adjustment costs: in order to achieve a growth rate g for his capital stock, the expert must invest a flow of $\iota(g)$ consumption goods. The capital he holds evolves[11]

$$\frac{dk_{i,t}}{k_{i,t}} = g_{i,t}dt + \sigma dZ_t + \nu_t dW_{i,t} \tag{1}$$

where $Z = \left\{Z_t \in \mathbb{R}^d; \mathcal{F}_t, t \geq 0\right\}$ is an aggregate brownian motion, and $W_i = \{W_{i,t}; \mathcal{F}_t, t \geq 0\}$ an idiosyncratic brownian motion for expert i,[12] in a probability space $(\Omega, P, \mathcal{F})$ equipped with a filtration $\{\mathcal{F}_t\}$ with the usual conditions.[13] The aggregate shock can be multidimensional, $d \geq 1$, so the economy could be hit by many aggregate shocks. For most results, however, there is no loss from taking $d = 1$ and focusing on a single aggregate shock.[14] While the exposure of capital to aggregate risk $\sigma \geq 0 \in \mathbb{R}^d$ is constant,[15] its exposure to idiosyncratic risk $\nu_t > 0$ follows an

[10] We could allow consumers to use capital less productively, as in Brunnermeier and Sannikov (2012) or Kiyotaki and Moore (1997). This doesn't change the main results.

[11] This formulation where capital is exposed to aggregate risk is equivalent to a standard growth model where TFP a_t follows a geometric Brownian Motion. Then if $\kappa_{i,t}$ is physical capital, $k_{i,t} = a_t \kappa_{i,t}$ is "effective capital" in the hands of expert i, so aggregate shocks to $k_{i,t}$ can be interpreted as persistent shocks to TFP a_t, i.e. $da_t = a_t \sigma dZ_t$. To preserve scale invariance we must also have investment costs proportional to a_t, which makes sense if we think investment requires diverting capital from consumption to investment (or in a richer model with labor).

[12] Idiosyncratic shocks $W_{i,t}$ represent shocks to the capital held by expert i over a short period, not to the productivity of the expert i. All experts are always equally good at using all capital. An increase in idiosyncratic risk could also reflect greater disagreement over the value of assets (Simsek (2013)) or an increased interest in acquiring information about assets (Gorton and Ordoñez (2013)).

[13] I will use an exact law of large numbers, which requires that we actually work with an extension of the Lebesgue interval $([0,1], \mathcal{I}, M)$ and a Fubini extension of the product space, $([0,1] \times \Omega, \mathcal{I} \boxtimes \mathcal{F}, M \boxtimes P)$, such that the $\{W_i\}_{i\in[0,1]}$ and Z are essentially pairwise independent, and such that for any i, $\int_{[0,1]} W_{i,t} dM = \int_{[0,1]} W_{i,t} dP = 0$ P-almost surely. I will abuse notation however and write $\int_{[0,1]} W_{i,t} di$ instead of $\int_{[0,1]} W_{i,t} dM$ to keep notation simple. See Sun and Zhang (2009) for details.

[14] This is in fact the approach I take when computing numerical solutions.

[15] I will use the convention that σ is a row vector, while Z_t a column vector. Throughout the paper I will not point this out unless it's necessary for clarity.

exogenous stochastic process

$$d\nu_t = \lambda\left(\bar{\nu} - \nu_t\right)dt + \sigma_\nu \sqrt{\nu_t} dZ_t \tag{2}$$

where $\bar{\nu}$ is the long-run mean and λ the mean reversion parameter.[16] The loading of the idiosyncratic volatility of capital on aggregate risk $\sigma_\nu \leq 0$ so that we may think of Z as a "good" aggregate shock that increases the effective capital stock and reduces idiosyncratic risk. This is just a naming convention. With multiple aggregate shocks, $d > 1$, we may take some shocks to be pure TFP shocks with $\sigma_\nu^{(i)} = 0$, other pure uncertainty shocks with $\sigma^{(i)} = 0$, and yet other mixed shocks.

The law of motion for the aggregate capital stock $k_t = \int_{[0,1]} k_{i,t} di$ is not affected by the idiosyncratic shocks $W_{i,t}$

$$dk_t = \left(\int_{[0,1]} g_{i,t} k_{i,t} di\right) dt + \sigma k_t dZ_t$$

Preferences. Both experts and consumers have Epstein-Zin preferences with the same discount rate ρ, risk aversion γ and elasticity of intertemporal substitution (EIS) ψ^{-1}. If we let $\gamma = \psi$ we get the standard CRRA utility case as a special case.

$$U_t = \mathbb{E}_t\left[\int_t^\infty e^{-\rho u} \frac{c_u^{1-\gamma}}{1-\gamma} du\right]$$

Epstein-Zin preferences separate risk-aversion from the EIS, which play different roles in the balance sheet amplification channel. They are defined recursively (see Duffie and Epstein (1992)):

$$U_t = \mathbb{E}_t\left[\int_t^\infty f\left(c_u, U_u\right) du\right] \tag{3}$$

where

$$f\left(c, U\right) = \frac{1}{1-\psi}\left\{\frac{\rho c^{1-\psi}}{\left[\left(1-\gamma\right)U\right]^{\frac{\gamma-\psi}{1-\gamma}}} - \rho\left(1-\gamma\right)U\right\}$$

I will later also introduce turnover among experts in order to obtain a non-degenerate stationary distribution for the economy. Experts will retire with independent Poisson arrival rate τ and become consumers. There is no loss in intuition from taking $\tau = 0$ for most of the results, however.

Markets. Experts can trade capital continuously at a competitive price $p_t > 0$, which we conjecture follows an Ito process:

$$\frac{dp_t}{p_t} = \mu_{p,t} dt + \sigma_{p,t} dZ_t$$

The price of capital depends on the aggregate shock Z but not on the idiosyncratic shocks $\{W_i\}_{i \in [0,1]}$, and is determined endogenously in equilibrium. The total value of the aggregate capital stock is $p_t k_t$ and it constitutes the total wealth of the economy.

[16] If $2\lambda\bar{\nu} \geq \sigma_\nu^2$, this Cox-Ingersoll-Ross process is always strictly positive and has a long-run distribution with mean $\bar{\nu}$.

There is also a complete financial market[17] with SDF η_t:

$$\frac{d\eta_t}{\eta_t} = -r_t dt - \pi_t dZ_t$$

Here r_t is the risk-free interest rate and π_t the price of aggregate risk Z. Both are determined endogenously in equilibrium. I am already using the fact that idiosyncratic risks $\{W_i\}_{i\in[0,1]}$ have price zero in equilibrium because they can be aggregated away.

Consumers' problem. Consumers face a standard portfolio problem. They cannot hold capital but they have access to a complete financial market. They start with wealth w_0 derived from ownership of a fraction of aggregate capital (which they immediately sell to experts). Taking the aggregate process η as given, they solve the following problem.

$$U_0 = \max_{(c,\sigma_w)} \mathbb{E}\left[\int_0^\infty f(c_t, U_t)\, dt\right]$$

$$st: \quad \frac{dw_t}{w_t} = (r_t + \sigma_{w,t}\pi_t - \hat{c}_t)\, dt + \sigma_{w,t} dZ_t$$

and a solvency constraint $w_t \geq 0$, where U_t is defined recursively as in (3), and the hat on $\hat{c}$ denotes the variable is normalized by wealth. I use w for the wealth of consumers, and reserve n for experts', which I will call "net worth". Consumers get the risk free interest rate on their wealth, plus a premium π_t for the exposure to aggregate risk $\sigma_{w,t}$ they chose to take. Since the price of expert-specific idiosyncratic risks $\{W_i\}$ is zero in equilibrium, consumers will never buy idiosyncratic risk. This is already baked into consumers' dynamic budget constraint.

Experts' problem. Experts face a more complex problem. An expert can continuously trade and use capital for production, as well as participate in the financial market. The cumulative return from investing a dollar in capital for expert i is $R_i^k = \{R_{i,t}^k; t \geq 0\}$ with

$$dR_{i,t}^k = \underbrace{\left(\frac{a - \iota(g_{i,t})}{p_t} + g_{i,t} + \mu_{p,t} + \sigma\sigma_{p,t}'\right)}_{\mathbb{E}_t\left[dR_{i,t}^k\right]} dt + \left(\sigma + \sigma_{p,t}\right) dZ_t + \nu_t dW_{i,t}$$

He would like to share risk with the market, but he faces a "skin in the game" constraint that forces him to keep an exposure to his own return $\tilde{\phi} \geq \phi \in (0,1)$. In Appendix A I derive this financial friction from a moral hazard problem, similar to Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011). The expert can secretly divert capital to a private account, but can only keep a fraction $\phi \in (0,1)$ of what he steals. In order to provide incentives to not steal, the expert must keep an exposure $\tilde{\phi}_t \geq \phi$ to the return of his capital, so that he loses more from stealing than

[17] A complete financial market could be implemented with different asset structures. For example, a natural asset structure would include risk-free debt, equity in each expert's investments and d market indices to span Z. If $d = 1$ we can do with only one market index.

what he wins. Importantly, I allow contracts to be written on the aggregate state of the economy. The expert's net worth therefore follows

$$\frac{dn_{i,t}}{n_{i,t}} = \tilde{\mu}_{i,n,t}dt + \tilde{\phi}_{i,t}p_t\hat{k}_{i,t}dR^k_{i,t} + \theta_{i,t}dZ_t \tag{4}$$

where $\tilde{\mu}_{i,n,t} = r_t(1 - p_t\hat{k}_{i,t}) + p_t\hat{k}_{i,t}(1 - \tilde{\phi}_{i,t})\Big(\mathbb{E}_t\left[dR^k_{i,t}\right] - (\sigma + \sigma_{p,t})\pi_t\Big) + \theta_{i,t}\pi_t - \hat{e}_{i,t}$. As before, the hatted variables denote they are divided by the net worth $n_{i,t}$.

The expert keeps an exposure $\tilde{\phi}_{i,t} \geq \phi$ to his own return and sells the rest $1 - \tilde{\phi}_{i,t}$ on the market.[18] This "skin in the game" constraint limits the expert's ability to share the idiosyncratic risk. Crucially, however, it does not limit his ability to share aggregate risk, which does not interact with the moral hazard problem. This is captured in (4) by the term $\theta_{i,t}$. This is the main difference with the contractual setup in Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011), where the additional constraint $\theta_{i,t} = 0$ is imposed (contracts cannot be written on the aggregate state of the economy). We can think of $\theta_{i,t}$ as the fraction of the expert's wealth invested in a set of aggregate securities that span Z (normalized to have an identity loading on Z). In the special case with only one aggregate shock, $d = 1$, we can think of this security as a normalized market index.

Given that the expert can use $\theta_{i,t}$ to adjust his exposure to aggregate risk, the "skin in the game" constraint will always be binding, i.e. $\tilde{\phi} = \phi$.[19] Re-writing the dynamic budget constraint (4) in terms of the structural shocks Z and W_i, the expert's problem is to maximize his expected utility

$$V_0 = \max_{(\hat{e},g,\hat{k},\theta)} \mathbb{E}\left[\int_0^\infty f\left(e_t, V_t\right)dt\right]$$

subject to a solvency constraint $n_t \geq 0$ and the dynamic budget constraint

$$\frac{dn_{i,t}}{n_{i,t}} = \left[\mu_{i,n,t} - \hat{e}_{i,t}\right]dt + \sigma_{i,n,t}dZ_t + \tilde{\sigma}_{i,n,t}dW_{i,t} \tag{5}$$

$$\begin{aligned}
\mu_{i,n,t} &= r_t + p_t\hat{k}_{i,t}\left(\mathbb{E}_t\left[dR^k_{i,t}\right] - r_t\right) - (1-\phi)\,p_t\hat{k}_{i,t}\left(\sigma + \sigma_{p,t}\right)\pi_t + \theta_{i,t}\pi_t \\
\sigma_{i,n,t} &= \phi p_t\hat{k}_{i,t}\left(\sigma + \sigma_{p,t}\right) + \theta_{i,t} \\
\tilde{\sigma}_{i,n,t} &= \phi p_t\hat{k}_{i,t}\nu_t
\end{aligned}$$

where V_t is defined recursively as in (3).

Notice that $\theta_{i,t}$ separates the decision of how much capital to buy, $\hat{k}_{i,t}$, from the decision of how much aggregate risk to carry, $\sigma_{i,n,t}$. When contracts cannot be written on the aggregate state of the economy we are restricted to $\theta_{i,t} = 0$, and the two decisions become entangled. The separation

[18] The market doesn't price the idiosyncratic risk $\nu_t dW_{i,t}$ contained in $dR^k_{i,t}$, but it does price the aggregate risk $(\sigma + \sigma_{p,t})dZ_t$ with π_t.

[19] Any desired exposure to aggregate risk can be handled with $\theta_{i,t}$, so $\tilde{\phi}_{i,t}$ should be minimized to reduce exposure to idiosyncratic risk, which is not rewarded by the market.

between investment in capital (or leverage) and aggregate risk-sharing is at the core of the TFP neutrality result. More generally, we can consider the intermediate case where contracts may only be written on a linear combination of aggregate shocks $\tilde{Z}_t = BZ_t$ for some full rank matrix $B \in \mathbb{R}^{d' \times d}$ with $d' < d$.[20] In this case we will be restricted to choosing $\theta_{i,t} = \tilde{\theta}_{i,t} B$.[21,22]

The optimal contract is easy to implement. The expert creates a firm with $p_t k_t$ assets, keeps a fraction ϕ of the equity, and sells the rest and borrows to raise funds (if $n_{i,t} > \phi p_t k_{i,t}$ he doesn't need to borrow, and he invests $n_{i,t} - \phi p_t k_{i,t}$ outside the firm). In addition, he trades aggregate securities (possibly indices of other firms' equity), and he receives a payment as CEO of the firm, which compensates him for the idiosyncratic risk he takes by keeping a fraction ϕ of his firm's equity).

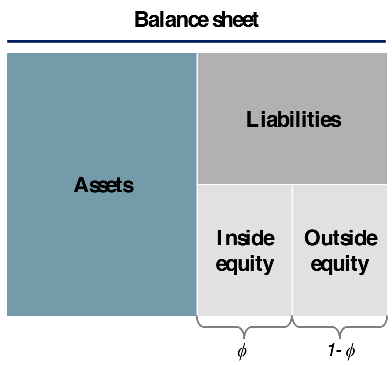

Equilibrium Denote the set of experts $\mathbb{I} = [0,1]$ and the set of consumers $\mathbb{J} = (1,2]$. We take the initial capital stock k_0 and its distribution among agents $\{k_i^0\}_{i\in\mathbb{I}}$, $\{k_j^0\}_{j\in\mathbb{J}}$ as given[23], with $\int_{\mathbb{I}} k_i^0 di + \int_{\mathbb{J}} k_j^0 dj = k_0$. Let $k_i^0 > 0$ and $k_j^0 > 0$ so that all agents start with strictly positive net worth.

Definition 1. *An equilibrium is a set of aggregate stochastic processes adapted to the filtration generated by Z: the price of capital $\{p_t\}$, the state price density $\{\eta_t\}$, and the aggregate capital stock $\{k_t\}$, and a set of stochastic processes*[24] *for each expert $i \in \mathbb{I}$ and each consumer $j \in \mathbb{J}$: net worth and wealth $\{n_{i,t}, w_{j,t}\}$, consumption $\{e_{i,t} \geq 0, c_{j,t} \geq 0\}$, capital holdings $\{k_{i,t}\}$, investment $\{g_{i,t}\}$, and aggregate risk sharing, $\{\sigma_{i,n,t}, \sigma_{j,w,t}\}$, such that:*

i. *Initial net worth satisfies $n_{i,0} = p_0 k_i^0$ and wealth $w_{j,0} = p_0 k_j^0$.*

ii. *Each expert and consumer solves his problem taking aggregate conditions as given.*

iii. *Market Clearing:*

$$\int_{\mathbb{I}} e_{i,t} di + \int_{\mathbb{J}} c_{i,t} dj + \int_{\mathbb{I}} \iota\left(g_{i,t}\right) k_{i,t} di = \int_{\mathbb{I}} a k_{i,t} di$$

[20] In terms of $\theta_{i,t}$ as a set of aggregate securities, this corresponds to an incomplete financial market.

[21] In particular, with $B = 0$ contracts cannot be written on Z.

[22] In this case, the "skin in the game" constraint may not be always binding.

[23] Consumers start holding capital and will immediately sell it to experts.

[24] *(each adapted to the filtation generated by Z and the $\{W_i\}_{i\in\mathbb{I}}$)*

$$\int_{\mathbb{I}} k_{i,t} di = k_t$$

$$\int_{\mathbb{I}} \sigma_{i,n,t} n_{i,t} di + \int_{\mathbb{J}} \sigma_{j,w,t} w_{j,t} dj = \int_{\mathbb{I}} p_t k_{i,t} \left(\sigma + \sigma_{p,t}\right) di$$

iv. *Law of motion of aggregate capital:*

$$dk_t = \left(\int_{\mathbb{I}} g_{i,t} k_{i,t} di \right) dt + k_t \sigma dZ_t$$

The market clearing conditions for the consumer goods and capital market are standard. The condition for market clearing in the financial market is derived as follows: we already know each expert keeps a fraction ϕ of his own equity. If we aggregate the equity sold on the market into indices with identity loading on Z, there is a total supply of these indices $(1 - \phi) p_t k_t (\sigma + \sigma_{t,p})$. Consumers absorb $\int_{\mathbb{J}} \sigma_{j,w,t} w_{j,t} dj$ and experts $\int_{\mathbb{I}} \theta_{i,t} n_{i,t} di$ of these indices. Rearranging we obtain the expression above. By Walras' law, the market for risk-free debt clears automatically.

3 Solving the model

Experts and consumers face a dynamic problem, where their optimal decisions depend on the stochastic investment possibility sets they face, captured by the price of capital p and the SDF η. The equilibrium is driven by the exogenous stochastic process for ν_t and by the endogenous distribution of wealth between experts and consumers. The recursive EZ preferences generate optimal strategies that are linear in net worth, and allow us to simplify the state-space: we only need to keep track of the net worth of experts relative to the total value of assets that they must hold in equilibrium, $x_t = \frac{n_t}{p_t k_t}$. The distribution of net worth across experts, and of wealth across consumers, is not important. The strategy is to use a recursive formulation of the problem and look for a Markov equilibrium in (ν_t, x_t), taking advantage of the scale invariance property of the economy which allows us to abstract from the level of the capital stock.

The layout of this section is as follows. First I solve a first best benchmark without moral hazard, and show the economy follows a stable growth path. Then back to the moral hazard case, I recast the equilibrium in recursive form and characterize agents' optimal plans. I study the effect of Brownian TFP shocks under different contractual environments. I then show how uncertainty shocks can create balance sheet recessions as a result of agents' optimal aggregate risk sharing decisions. Finally, I consider general aggregate shocks.

3.1 Benchmark without moral hazard

Without any financial frictions this is a standard AK growth model where balance sheets don't play any role. Because there is no moral hazard, experts share all of their idiosyncratic risk, so the dynamics of idiosyncratic shocks ν_t are irrelevant. Without financial frictions, the price of capital and the growth rate of the economy do not depend on experts' net worth: balance sheets are only relevant to determine consumption of experts and consumers. The economy follows a stable growth path.

Proposition 1 (First best benchmark)**.** *If $\rho - (1-\psi)g^* + (1-\psi)\frac{\gamma}{2}\sigma^2 > 0$ and without any financial frictions, there is a stable growth equilibrium, where the price of capital is p^* and the growth rate g^*, given by:*

$$\iota'(g^*) = p^* \tag{6}$$

$$p^* = \frac{a - \iota(g^*)}{\rho - (1-\psi)g^* + (1-\psi)\frac{\gamma}{2}\sigma^2} \tag{7}$$

3.2 Back to moral hazard

First, from homothetic preferences we know that the value function for an expert with net worth n takes the following power form:

$$V(\xi_t, n) = \frac{(\xi_t n)^{1-\gamma}}{1-\gamma}$$

for some stochastic process $\xi = \{\xi_t > 0; t \geq 0\}$. I call ξ the "net worth multiplier". It captures the stochastic, general equilibrium investment possibility set the expert faces (it does not depend on his own net worth n_t). When ξ_t is high the expert is able to obtain a large amount of utility from a given net worth n_t, as if his actual net worth was $\xi_t n_t$. Conjecture that it follows an Ito process

$$\frac{d\xi_t}{\xi_t} = \mu_{\xi,t}dt + \sigma_{\xi,t}dZ_t$$

where $\mu_{\xi,t}$ and $\sigma_{\xi,t}$ must be determined in equilibrium. For consumers, the utility function takes the same form, $U(\zeta_t, n) = \frac{(\zeta_t n)^{1-\gamma}}{1-\gamma}$ but instead of ξ_t, we have ζ_t as the "wealth multiplier" which follows $\frac{d\zeta_t}{\zeta_t} = \mu_{\zeta,t}dt + \sigma_{\zeta,t}dZ_t$, also determined in equilibrium.

I use a dynamic programming approach to solve agents' problem. For experts, we have the Hamilton-Jacobi-Bellman equation after some algebra:

$$\frac{\rho}{1-\psi} = \max_{\hat{e},g,\hat{k},\theta} \left\{ \frac{\hat{e}^{1-\psi}}{1-\psi}\rho\xi_t^{\psi-1} + \mu_n - \hat{e} + \mu_\xi - \frac{\gamma}{2}\left(\sigma_n^2 + \sigma_\xi^2 - 2(1-\gamma)\sigma_n\sigma_\xi + \tilde{\sigma}_n^2\right) \right\} \tag{8}$$

subject to the dynamic budget constraint (5), and a transversality condition. Consumers have an analogous HJB equation.

Proposition 2. *[Linearity] All experts chose the same $\hat{e}_t$, g_t, $\hat{k}_t$ and θ_t, and all consumers the same $\hat{c}_t$ and $\sigma_{w,t}$. In addition, growth is determined by a static FOC*

$$\iota'(g_t) = p_t$$

Proposition 2 tells us two things. The first is that the growth rate of the economy is linked to asset prices in a straightforward way. Anything that depresses asset prices will have a real effect on the growth rate of the economy. For example, with a quadratic adjustment cost function $\iota(g) = Ag^2$, the growth rate of the economy is simply $g_t = \frac{p_t}{2A}$.

Proposition 2 also tells us that policy functions are linear in net worth. This is a useful property of homothetic preferences and allows us to abstract from the distribution of wealth across experts and across consumers, and simplifies the state space of the equilibrium. We only need to keep track of the fraction of aggregate wealth that belongs to experts: $x_t = \frac{n_t}{p_t k_t} \in [0,1]$. I look for a Markov equilibrium with two state variables: the volatility of idiosyncratic shocks ν_t, and x_t:

$$p_t = p(\nu_t, x_t), \quad \xi_t = \xi(\nu_t, x_t), \quad \zeta_t = \zeta(\nu_t, x_t), \quad r_t = r(\nu_t, x_t), \quad \pi_t = \pi(\nu_t, x_t)$$

where p, ξ and ζ are conjectured to be twice continuously differentiable. The first state variable ν_t evolves exogenously according to (2). The state variable x_t is endogenous, and has an interpretation in terms of experts' balance sheets. Since experts must hold all the capital in the economy, the denominator captures their assets while the numerator is the net worth of the expert sector as a whole. I will sometimes refer to x_t as "experts' balance sheets".

We know from Proposition 1 that without moral hazard, experts would be able to offload all of their idiosyncratic risk onto the market and hence neither ν_t nor x_t would play any role in equilibrium. In contrast, in an economy with financial frictions, $\phi > 0$, experts' balance sheets will play an important role. We say *balance sheets matter* if equilibrium objects depend on x_t. In order for balance sheets to play a role in the transmission and amplification of aggregate shocks, we also need them to be exposed to aggregate shocks. In principle, x_t could be exposed to aggregate risk Z_t through its volatility term $\sigma_{x,t}$, or through a stochastic drift $\mu_{x,t}$. In practice, what we usually mean when we talk about a balance sheet channel is that experts' balance sheets are disproportionally hit by aggregate shocks, so we want to focus on $\sigma_{x,t} > 0$.[25] We say there is a *balance sheet amplification channel* if balance sheets matter and, in addition, $\sigma_{x,t} > 0$.

We can now give a definition for a Markov equilibrium..

Definition 2. *A Markov Equilibrium in (ν, x) is a set of aggregate functions p, ξ, ζ, r, π and policy functions $\hat{e}, g, \hat{k}, \theta$ for experts and $\hat{c}, \sigma_{w,t}$ for consumers, and a law of motion for the endogenous aggregate state variable $dx_t = \mu_x(\nu, x)\,dt + \sigma_x(\nu, x)\,dZ_t$ such that:*

i. *ξ and ζ solve the experts' and consumers' HJB equations (8), and $\hat{e}, g, \hat{k}, \theta$ and $\hat{c}, \sigma_{w,t}$ are the corresponding policy functions, taking p, r, π and the laws of motion of ν_t and x_t as given.*

ii. *Market clearing:*

$$\hat{e}px + \hat{c}p(1-x) = a - \iota(g)$$

$$p\hat{k}x = 1$$

$$\sigma_n x + \sigma_w(1-x) = \sigma + \sigma_p$$

[25] As it turns out, this distinction won't be important, since in the TFP shocks case both $\sigma_{x,t} = 0$ and the drift $\mu_{x,t}$ is non-stochastic, while with uncertainty shocks both $\sigma_{x,t} > 0$ and $\mu_{x,t}$ is stochastic.

iii. *x follows the law of motion (9) derived using Ito's lemma:*

$$\mu_x(\nu, x) = x\Big(\mu_n - \hat{e} - g - \mu_p - \sigma\sigma_p' + \big(\sigma + \sigma_p\big)^2 - \sigma_n\big(\sigma + \sigma_p\big)'\Big) \tag{9}$$

$$\sigma_x(\nu, x) = x\left(\sigma_n - \sigma - \sigma_p\right)$$

This recursive definition abstracts from the absolute level of the aggregate capital stock, which we can recover using $\frac{dk_t}{k_t} = g_t dt + \sigma dZ_t$.

Capital holdings. Experts demand for capital is pinned down by the FOC from the HJB equation. After some algebra we obtain an expression that pins down the demand for capital $\hat{k}$:

$$\frac{a - \iota_t}{p_t} + g_t + \mu_{p,t} + \sigma\sigma_{p,t}' - r_t \leq \left(\sigma + \sigma_{p,t}\right)\pi_t + \gamma p_t \hat{k}_t \left(\phi\nu_t\right)^2$$

Idiosyncratic risk is not priced in the financial market, because it can be aggregated away. However, because experts face an equity constraint that forces them to keep an exposure ϕ to the return of their capital, they know that the more capital they hold, the more idiosyncratic risk they must bear on their balance sheets $\tilde{\sigma}_{n,t} = \phi p_t \hat{k}_t \nu_t$. They consequently demand a premium on capital for that idiosyncratic risk. Using the equilibrium condition $p\hat{k}x = 1$ we obtain an equilibrium pricing equation for capital:

$$\underbrace{\frac{a - \iota_t}{p_t} + g_t + \mu_{p,t} + \sigma\sigma_{p,t}' - r_t}_{\text{excess return}} = \underbrace{\left(\sigma + \sigma_{p,t}\right)\pi_t}_{\text{agg. risk premium}} + \underbrace{\gamma\frac{1}{x_t}\left(\phi\nu_t\right)^2}_{\text{id. risk premium}} \tag{10}$$

The left hand side is the excess return of capital. The right hand side is made up of the risk premium corresponding to the aggregate risk capital carries, and a risk premium for the idiosyncratic risk it carries. When experts balance sheets are weak (low x_t) and idiosyncratic risk ν_t high, experts demand a high premium on capital. This is how x_t and ν_t affect the economy, and we can see that without moral hazard, $\phi = 0$, neither x_t nor ν_t would play any role, and experts would be indifferent about how much capital to hold as long as it was properly priced. With moral hazard, instead, they have a well defined demand for capital, proportional to their net worth.

It is useful to reformulate experts' problem with a "fictitious" price of idiosyncratic risk

$$\alpha_t = \gamma\frac{\phi\nu_t}{x_t}$$

Under this formulation, each expert faces a complete financial market without the equity constraint, but where his own idiosyncratic risk W_i pays a premium α_t. Capital is priced as an asset with exposure $\phi\nu_t$ to this idiosyncratic risk, and can be abstracted from.[26] We can verify that the expert

[26]We can use (10) to rewrite experts' dynamic budget constraint

$$\frac{dn_t}{n_t} = \left(r_t + \pi_t\sigma_{n,t} + \alpha_t\tilde{\sigma}_{n,t}\right)dt + \sigma_{n,t}dZ_t + \tilde{\sigma}_{n,t}dW_{i,t}$$

will chose an exposure to his own idiosyncratic risk $\tilde{\sigma}_{n,t} = \frac{\alpha_t}{\gamma} = \phi \frac{1}{x_t} \nu_t$ as required in equilibrium. In this sense the fictitious price of idiosyncratic risk α_t is "right". An advantage of the fictitious price formulation is that the only difference between experts' and consumers' problem is that experts perceive a positive price for idiosyncratic risk $\alpha_t > 0$, while consumers perceive a price of zero.

Aggregate risk sharing. Optimal contracts allow experts to share aggregate risk freely. The optimal contract effectively separates the decision of how much capital to hold $k_{i,t}$ from the decision of how much aggregate risk to hold $\sigma_{n,t}$. The FOC for aggregate risk sharing for experts are:

$$\sigma_{n,t} = \underbrace{\frac{\pi'_t}{\gamma}}_{\text{myopic}} - \underbrace{\frac{\gamma - 1}{\gamma} \sigma_{\xi,t}}_{\text{hedging motive}} \tag{11}$$

Experts' optimal aggregate risk exposure depends on a myopic risk-taking motive given by the price of risk[27] and the risk-aversion parameter, $\frac{\pi'_t}{\gamma}$, and a hedging motive driven by the stochastic investment possibility sets, $\frac{\gamma-1}{\gamma}\sigma_{\xi,t}$. This hedging motive is standard in Intertemporal CAPM models, going back to Merton (1973), and it will play a crucial role in the amplification and propagation of aggregate shocks through experts' balance sheets. Recall the "net worth multiplier" ξ_t captures the stochastic general equilibrium conditions the expert faces

$$V_t(n) = \frac{(\xi_t n)^{1-\gamma}}{1-\gamma}$$

If the expert is risk neutral, he will prefer to have more net worth when ξ_t is high, since he can obtain a lot of long-term utility out of each unit of net worth. This is a "substitution effect". On the other hand, when ξ_t is low he requires more net worth to achieve any given level of utility. If the expert is risk averse, he will prefer to have more net worth when ξ_t is low. This is a "wealth effect". Which effect dominates depends on the risk aversion parameter.[28] When $\gamma < 1$, equation (11) tells us the expert wants his net worth to be positively correlated with ξ_t: the substitution effect dominates. When $\gamma > 1$, instead, the wealth effect dominates. I focus on the case where the wealth effect dominates, $\gamma > 1$.

---

where the expert can freely choose $\sigma_{n,t}$ and $\tilde{\sigma}_{n,t}$. Experts problem then is to maximize their objective function, subject to an intertemporal budget constraint

$$\mathbb{E}\left[\int_0^\infty \tilde{\eta}_u e_u du\right] = n_0$$

where the fictitious SPD $\tilde{\eta}$ follows: $\frac{d\tilde{\eta}_t}{\tilde{\eta}_t} = -r_t dt - \pi_t dZ_t - \alpha_t dW_{i,t}$ for expert i.

[27] π_t is a column vector and must be transposed, hence π'_t.

[28] EZ preferences separate risk aversion γ from the EIS ψ^{-1}. Below I explore the role of each parameter in the model.

Consumers have analogous FOC conditions for aggregate risk sharing

$$\sigma_{w,t} = \underbrace{\frac{\pi'_t}{\gamma}}_{\text{myopic}} - \underbrace{\frac{\gamma - 1}{\gamma}\sigma_{\zeta,t}}_{\text{hedging motive}} \tag{12}$$

where the only difference is that consumers' investment possibility sets are captured by ζ_t instead of ξ_t. Since consumers cannot buy capital, its price and idiosyncratic risk-premium does not affect them, but they still face a stochastic investment possibility set from interest rates r_t and the price of aggregate risk π_t.

The volatility of balance sheets $\sigma_{x,t}$ arises from the interaction of experts' and consumers' risk-taking decisions. Using the equilibrium condition $\sigma_n x + \sigma_w (1 - x) = (\sigma + \sigma_p)$ we obtain the following aggregate risk-sharing equation

$$\sigma_{x,t} \quad = \quad (1 - x_t)x_t \frac{1 - \gamma}{\gamma} \underbrace{(\sigma_{\xi,t} - \sigma_{\zeta,t})}_{\substack{\text{relative} \\ \text{hedging motive}}} \tag{13}$$

The term $(1-x_t)x_t$ arises because experts are able to hedge their investment possibility sets only to the extent that consumers as a whole are willing to take the other side of the hedge. The $\frac{1-\gamma}{\gamma}$ term captures the "substitution" and "wealth" effects, while $\sigma_{\zeta,t} - \sigma_{\xi,t}$ captures experts' and consumers' relative hedging motive. Since experts and consumers cannot both hedge in the same direction in equilibrium, it is the difference in their hedging motives which will cause experts' balance sheets to be overexposed to aggregate risk.

To understand aggregate risk-sharing better, notice that because experts have the option of investing in capital, they always get more utility per dollar of net worth than consumers, i.e. $\xi_t > \zeta_t$. Call

$$\Omega_t = \log \xi_t - \log \zeta_t > 0$$

the *investment opportunity gap* between experts and consumers. This gap is not constant, however: it depends on the aggregate state of the economy. The relative hedging motive is the loading of the investment opportunity gap Ω_t on the aggregate shock dZ_t. Equation (13) says that if the wealth effect dominates ($\gamma > 1$) agents will share aggregate risk so that experts have a smaller share of aggregate wealth when the gap is large (because they are already relatively better off in this state, compared to consumers), and a larger share when the gap is small

$$\sigma_{x,t} = (1 - x_t)x_t \frac{1 - \gamma}{\gamma} vol(\Omega_t)$$

Experts' and consumers' investment possibility sets, and hence the log gap, depends on balance sheets x_t, and so are endogenously determined in equilibrium in a two way feedback: experts' balance sheets are exposed to aggregate risk to hedge stochastic investment possibility sets, but the volatility of investment possibility sets actually depends on the exposure of experts' balance sheets to aggregate risk. We can use Ito's lemma to obtain a simple expression for the volatility of the investment opportunity gap Ω_t:

$$vol(\Omega_t) = \underbrace{\Omega_\nu \sigma_\nu \sqrt{\nu_t}}_{\text{exogenous}} + \underbrace{\Omega_x \sigma_{x,t}}_{\text{endogenous}} \tag{14}$$

where the function Ω is evaluated at (ν_t, x_t). The locally linear representation allows a neat decomposition into an *exogenous* source, driven by the uncertainty shock to ν_t, and an *endogenous* source from optimal contracts' aggregate risk sharing $\sigma_{x,t}$. We can solve for the fixed point of this two-way feedback:

$$\sigma_{x,t} = \frac{(1-x_t)x_t \frac{1-\gamma}{\gamma}\Omega_\nu}{1-(1-x_t)x_t\frac{1-\gamma}{\gamma}\Omega_x}\sqrt{\nu_t}\sigma_\nu \tag{15}$$

Notice that even though the presence of moral hazard does not directly restrict experts' ability to share aggregate risk, it introduces hedging motives through the general equilibrium which would not be present without moral hazard, as shown by Proposition 1.

3.3 Brownian TFP shocks

When aggregate shocks come only in the form of Brownian TFP shocks ($\sigma_\nu = 0$) and we allow agents to write contracts on all observable variables, there is no balance sheet channel. After a negative TFP shock, the value of all assets $p_t k_t$ falls and everyone, experts and consumers alike, looses net worth proportionally, so $\sigma_{x,t} = 0$. Experts then have lower net worth, but the value of capital they must hold in equilibrium is also lower, so the idiosyncratic risk they must carry as a proportion of their net worth is not affected by TFP shocks. Investment possibility sets then are not affected by aggregate shocks, and consequently the investment opportunity gap Ω_t is not affected by aggregate shocks and there is no relative hedging motive, $vol(\Omega_t) = \sigma_{\zeta,t} - \sigma_{\xi,t} = 0$. Balance sheets x_t may still affect the economy, due to the presence of financial frictions derived from the moral hazard problem, but they won't be exposed to aggregate risk and hence won't play any role in the amplification of aggregate TFP shocks. In fact, the equilibrium is completely deterministic, up to the direct effect of TFP shocks on the aggregate capital stock.

Proposition 3. *With only Brownian TFP shocks ($\sigma_\nu = 0$) if agents can write contracts on the aggregate state of the economy, the balance sheet channel disappears: the state variable x_t, the price of capital p_t, the growth rate of the economy g_t, the interest rate r_t, and the price of risk π_t all follow deterministic paths and are not affected by aggregate shocks.*

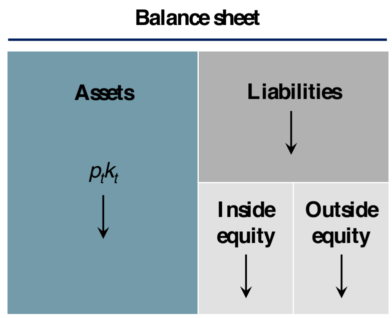

The negative result of Proposition 3 has two ingredients: 1) optimal contracts separate the decision of how much capital to buy (leverage) from the decision of how much aggregate risk to hold (risk sharing).[29] Risk sharing between experts and consumers will depend only on their relative hedging motives. The difference in hedging motives is ultimately traced to the fact that experts can trade and use capital, and the gap in investment opportunities Ω_t that this creates. This gives us expression (13):

$$\sigma_{x,t} = (1 - x_t)x_t \frac{1-\gamma}{\gamma} vol(\Omega_t)$$

And 2) aggregate Brownian TFP shocks don't affect investment possibility sets directly and so don't create a relative hedging motive by themselves. The exogenous source of relative hedging motive disappears, so we are left with only the endogenous component in expression (14):

$$vol(\Omega_t) = \underbrace{\Omega_\nu \sigma_\nu \sqrt{\nu_t}}_{\text{exogenous} = 0} + \underbrace{\Omega_x \sigma_{x,t}}_{\text{endogenous}}$$

With no exogenous source, however, the unique Markov equilibrium has deterministic investment possibility sets, no relative hedging motive, and hence no overexposure to aggregate risk which could endogenously affect investment possibility sets. The continuous-time setting provides a locally linear relationship that guarantees this is the unique Markov equilibrium, given by equation (15). Without any source of aggregate volatility, the economy then follows a deterministic path.

Implementation and constrained contracts. The optimal contract can be implemented with simple financial instruments: an expert buys capital $p_t k_t$ and sells a fraction $1 - \phi$ of his equity. He then buys a market index, or shorts it, to obtain the right exposure to aggregate risk. Even though the ability to short the market index is important for deriving Proposition 3 in general, experts might typically be going long on the market index. We can compute their investment in

[29] In Di Tella (2012) I explore under what conditions moral hazard can distort aggregate risk-sharing for incentive provision reasons. This can happen if the expert's private action affects the exposure to aggregate risk of his private benefit. Optimal contracts will overexpose experts to aggregate risk in order to deter them from taking a private action that further exposes them to it, creating a tradeoff between aggregate and idiosyncratic risk-sharing. However, this requires control over the expert's portfolio investments beyond his equity stake in the project he runs.

the normalized market index θ_t explicitly:

$$\theta_t = (\sigma + \sigma_{p,t}) \underbrace{\frac{x_t - \phi}{x_t}}_{\gtreqless 0} + \underbrace{\frac{\sigma_{x,t}}{x_t}}_{=0} = \underbrace{\frac{x_t - \phi}{x_t}}_{\gtreqless 0} \sigma$$

Their portfolio position on the market indices will be positive or negative depending on whether $x_t \gtreqless \phi$. With aggregate shocks don't affect the investment opportunity gap, $vol(\Omega_t) = 0$, experts and consumers will hold a fraction of aggregate wealth proportional to their net worth or wealth, respectively: $\sigma_{n,t} = \sigma_{w,t} = \sigma + \sigma_{p,t} = \sigma$. Experts are required to hold a fraction ϕ of their equity, which already exposes them to a fraction ϕ of aggregate risk. If their net worth represents more than fraction ϕ of aggregate wealth $x_t > \phi$, they will want to buy more aggregate risk by going long on a market index, or their competitors' equity. On the other hand, if $x_t < \phi$, they will short the market index to get rid of some of the aggregate risk contained in their equity.

In general, the economy may be hit by a large number of orthogonal aggregate shocks, i.e. $d > 1$. The negative result in Proposition 3 doesn't require complete markets, only that leverage and aggregate risk-sharing be separated. In terms of implementation in a financial market, we need the financial market to span the exposure to aggregate risk of the return to capital σdZ. In this case, an expert can buy capital and immediately get rid of the aggregate risk using financial instruments. He can then share aggregate risk with consumers using any available financial instruments. Without any endogenous hedging motives both experts and consumers will choose the same exposure to aggregate risk and eliminate the balance sheet channel.

Proposition 4. *Even if the financial market spans only a linear combination of aggregate shocks $\tilde{Z}_t = BZ_t$, for a matrix B with full rank $d' < d$, then as long as the exposure of capital to aggregate risk σ is in the row space of B, the result of Proposition 3 holds.*

If experts can short the equity of their competitors, who have a similar exposure to aggregate risk as they do, they can get rid of the aggregate risk in their capital. In a competitive market, there is a large number of competitors so their idiosyncratic risks can be aggregated away. In other words, an index made up competitors' equity is exactly the instrument required to separate leverage from risk sharing and obtain the negative result.

Corollary 1. *As long as a market index of experts' equity can be traded, the balance sheet channel disappears.*

In contrast to Proposition 3, when we rule out contracts on aggregate shocks, i.e. $B = 0$, experts' leverage and aggregate risk sharing become entangled. In the simplest case with $\phi = 1$ as in the baseline setting in BS, if experts are leveraged $p_t k_t > n_t$, then when a negative aggregate shock reduces the value of capital experts will lose net worth more than proportionally:

$$\sigma_{x,t} = x_t (\sigma_n - \sigma - \sigma_{p,t}) = x_t \underbrace{\left(p_t \hat{k}_t - 1\right)}_{>0} (\sigma + \sigma_{p,t}) > 0$$

This reduces their ability to hold capital and lowers asset prices, further hurting their balance sheets, and amplifying and propagating the initial shock. This is precisely the mechanism behind the balance sheet channel in Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2011).[30]

3.4 Uncertainty shocks

In this section I show that, in contrast to Brownian TFP shocks, uncertainty shocks that increase idiosyncratic risk depress asset prices and growth, and lead to balance sheet recessions. The "skin in the game" constraint forces experts to keep a fraction of their idiosyncratic risk, so during periods of high idiosyncratic risk and weak balance sheets, the price of capital and growth are low. Even though experts can share aggregate risk freely, they choose to be highly exposed to this aggregate risk ex-ante in order to take advantage of ex-post investment opportunity sets. Weak balance sheets further depress the price of capital and growth and create a balance sheet recession, which in turn amplifies experts' incentives to take even more aggregate risk ex-ante in a two-way feedback loop. In addition, an increase in *idiosyncratic* risk leads to an endogenous increase in *aggregate* risk, and can trigger a "flight to quality" event with low interest rate and high risk premiums.

The strategy to solve for the equilibrium with uncertainty shocks is to map it into a set of three partial differential equations for the price of capital $p(\nu, x)$ and the multipliers $\xi(\nu, x)$ and $\zeta(\nu, x)$. The pricing equation for capital (10), experts' HJB (8) and market clearing for consumption goods provide three functional equations.[31] In the interest of simplicity, I consider a one dimensional exogenous Brownian shock, i.e. $d = 1$. This shock will affect both capital directly through $\sigma > 0$ and the volatility of idiosyncratic risk ν_t through $\sqrt{\sigma_\nu}\nu_t < 0$.[32] I also add turnover among experts as described in Section 2 to keep them from asymptotically accumulating all the wealth in the economy.[33] I show the procedure for solving the equilibrium in detail in Appendix C.

Balance sheet channel. An uncertainty shock increases idiosyncratic risk ν_t in the economy and endogenously reduces the fraction of aggregate wealth that belongs to experts x_t. Both effects increase the idiosyncratic risk premium on capital and drive its price p_t down[34]

$$\underbrace{g_t + \mu_{p,t} + \sigma\sigma'_{p,t} + \frac{a - \iota_t}{p_t} - r_t}_{\text{excess return}} = \underbrace{(\sigma + \sigma_{p,t})\,\pi_t}_{\text{agg. risk premium}} + \underbrace{\gamma\frac{1}{x_t}(\phi\nu_t)^2 \uparrow}_{\text{id. risk premium}}$$

In addition, the fall in the price of capital drives growth down, given by the FOC $\iota'(g_t) = p_t$. An uncertainty shock therefore produces a downturn in the economy, with depressed asset values and

[30] In He and Krishnamurthy (2011) a similar mechanism underlies the volatility of experts' net worth (specialists in their model), but the price of capital falls because consumers are more impatient and interest rates must rise for consumption-goods markets to clear.

[31] Two second order partial differential equations and an algebraic constraint.

[32] Additional TFP shocks will only have a direct impact on the level of the effective capital stock, but will have no further effects on the economy.

[33] Even in that case the financial constraints will never stop binding.

[34] I focus on the case with relative risk aversion $\gamma > 1$ and elasticity of intertemporal substitution $\psi > 1$, for which EZ preferences are necessary. I explore the role of both parameters below.

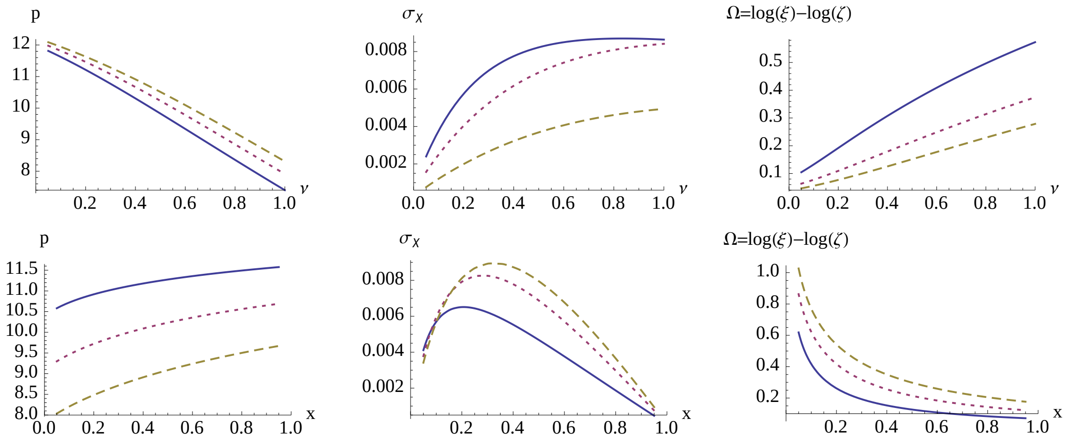

Figure 1: The price of capital p, volatility of x, σ_x, and the gap $\Omega = \log \xi - \log \zeta$, as functions of ν (above) for $x = 0.25$ (solid), $x = 0.5$ (dotted), and $x = 0.75$ (dashed), and as a function of x (below) for $\nu = 0.25$ (solid), $\nu = 0.5$ (dotted), and $\nu = 0.75$ (dashed). For parameter values see Numerical Solution below.

growth, amplified through a balance sheet channel. Figure 1 shows the price of capital is decreasing in idiosyncratic risk ν_t and increasing in the experts' share of aggregate wealth x_t. The growth rate has the same shape (I use $\iota(g) = Ag^2$, so $g_t = \frac{1}{2A}p_t$).[35]

Figure 1 also shows $\sigma_{x,t}$ is positive throughout, so experts' share of aggregate wealth x_t falls after an uncertainty shock raises ν_t, amplifying its effects through a balance sheet channel. The intuition is the following: an uncertainty shock endogenously increases the premium on idiosyncratic risk α_t, and this benefits experts compared to consumers because only experts perceive this premium. If the wealth effect dominates ($\gamma > 1$) agents want to stabilize their relative utility across states, so experts must have a smaller fraction of aggregate wealth (low x_t) after an uncertainty shock. This in turn makes α_t even larger, amplifying the effects of the uncertainty shock and inducing experts to take even more aggregate risk in a two-way feedback loop.

To see this in more detail, recall that aggregate risk sharing is given by

$$\sigma_{x,t} = (1 - x_t)x_t \frac{1-\gamma}{\gamma} vol(\Omega_t) \tag{16}$$

where $\Omega_t = \log \xi_t - \log \zeta_t$ is the investment opportunity gap between experts and consumers which depends on the aggregate shock dZ_t through both ν_t and x_t. When the wealth effect dominates ($\gamma > 1$) agents want to stabilize their relative utility across states, so they will share aggregate risk to give experts a smaller share of aggregate wealth when the gap Ω_t is large (because they are already relatively better off in this state, compared to consumers).

Figure 1 shows the investment opportunity gap $\Omega_t = \log \xi_t - \log \zeta_t$ is large when idiosyncratic

[35]Since this is an AK model, only TFP shocks can have effects on the current output level. I call periods of depressed growth recessions, although they have relatively high consumption. The model can be easily extended to allow consumers to use capital with a lower productivity $a_c < a$ as in Kiyotaki and Moore (1997) or Brunnermeier and Sannikov (2012). Under this variant, after an uncertainty shock experts will offload capital onto consumers, and output will go down since average productivity will be lower. Downturns will then have lower growth, output, and consumption.

volatility ν_t is high and experts' balance sheets x_t are weak. That is, experts are relatively better off than consumers during downturns, conditional on net worth. To grasp the intuition behind this, recall that in the fictitious price formulation, the only difference between experts' and consumers' problems is that experts perceive a positive price α_t for their own idiosyncratic risk (whereas consumers do not). In equilibrium experts go long on their own idiosyncratic risk ($\tilde{\sigma}_{n,t} = \frac{\phi}{x_t}\nu_t > 0$) so they benefit when α_t goes up. From equation (17) we see α_t is increasing in ν_t and decreasing in x_t.

$$\uparrow\uparrow \ \alpha_t = \gamma \frac{\phi\, \nu_t \ \uparrow}{x_t \ \downarrow} \tag{17}$$

Therefore, experts are better off (compared to consumers) during downturns, after an uncertainty shock raises ν_t and endogenously depresses x_t. Both effects work together to increase the investment opportunity gap Ω_t after an uncertainty shock[36] ($dZ_t < 0$)

$$vol(\Omega_t) = \sigma_{\xi,t} - \sigma_{\zeta,t} = \underbrace{\overbrace{\Omega_\nu}^{>0} \ \overbrace{\sigma_\nu \sqrt{\nu_t}}^{<0}}_{\text{exogen. } < 0} + \underbrace{\overbrace{\Omega_x}^{<0} \ \overbrace{\sigma_{x,t}}^{>0}}_{\text{endogen. } < 0} < 0$$

In line with equation (16), this induces experts to take a disproportionate share of aggregate risk, $\sigma_x > 0$, so x_t falls after an uncertainty shock. The more aggregate risk experts take (larger $\sigma_{x,t}$), the more x_t falls after an uncertainty shock and the bigger the gap in investment opportunities Ω_t in that state (through the endogenous component), which in turn induces them to take even more aggregate risk ex-ante, in a *two-way feedback loop*[37]

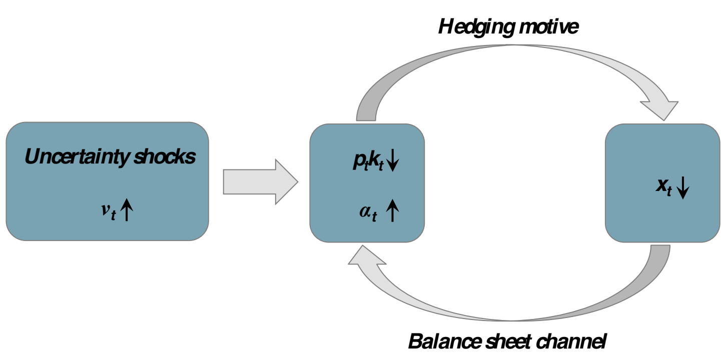

Experts are not necessarily better off during downturns. First, because they (endogenously) face large loses of net worth. But even conditional on net worth, experts might be worse-off after an

[36]The net worth multipliers ξ_t and ζ_t are forward looking and capture the expected present value of investment possibility sets, but we can get intuition about how the gap Ω_t behaves by studying α_t.

[37]Notice that the endogenous response of asset prices amplifies the effect of the exogenous shock on the balance sheets of experts, as in Kiyotaki and Moore (1997). In that paper, however, this happens ex-post because experts *cannot* hedge this risk. Here, instead, it happens ex-ante because they *can* hedge, and choose to increase their exposure to aggregate risk in anticipation of the response of asset prices to the exogenous shock.

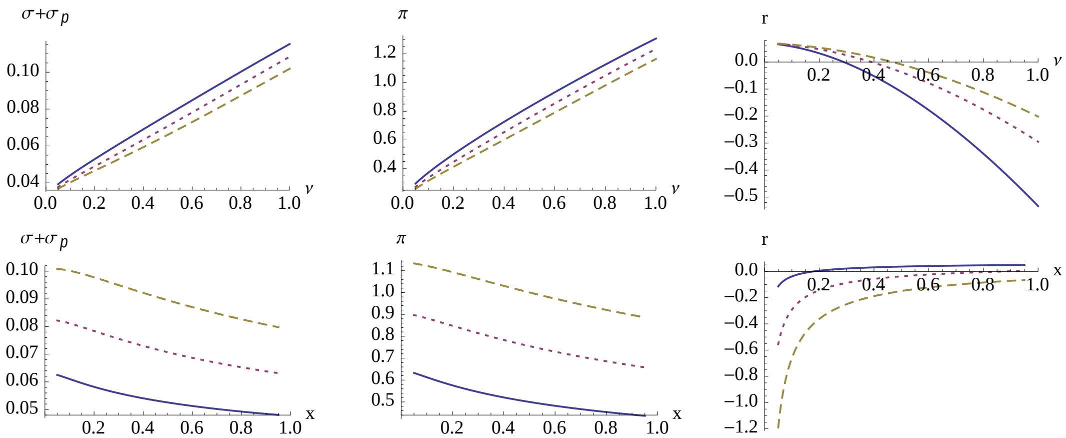

Figure 2: Aggregate risk $\sigma + \sigma_p$, the price of risk π, and risk free interest rate r as functions of ν (above) $x = 0.25$ (solid), $x = 0.5$ (dotted), and $x = 0.75$ (dashed), and as a function of x (below) for $\nu = 0.25$ (solid), $\nu = 0.5$ (dotted), and $\nu = 0.75$ (dashed). For parameter values see Numerical Solution below.

uncertainty shock because interest rates r_t and the risk premia π_t are also affected. It might very well be the case then that both experts and consumers are worse off after an uncertainty shock, even conditional on net worth (that is, both ξ_t and ζ_t lower). What matters for aggregate risk-sharing, however, is that the *gap* Ω_t between them rises after an uncertainty shock, because experts at least get higher premiums α_t on idiosyncratic risk. As a result, for a given price of aggregate risk π_t experts find taking aggregate risk more attractive than consumers, and in equilibrium the market concentrates a disproportionate share of aggregate risk on the balance sheets of experts.

Aggregate risk and flight to quality. After an uncertainty shock exogenously raises *idiosyncratic* risk ν_t in the economy, the economy experiences also endogenously high *aggregate* risk $\sigma + \sigma_{p,t}$ and a *flight to quality* event with low interest rates r_t and high risk premia π_t. Figure 2 shows aggregate risk rising when idiosyncratic risk ν_t is high and balance sheets x_t are weak. The model then provides an explanation for the observation that idiosyncratic and aggregate volatility seems to move together.[38]

The demand for aggregate risk from agents for hedging purposes also falls during downturns

$$\underbrace{\frac{\pi_t}{\gamma}}_{\text{myopic}\uparrow} + \underbrace{\frac{1-\gamma}{\gamma}\left(\sigma_{\xi,t}x_t + \sigma_{\zeta,t}(1-x_t)\right)}_{\text{hedging}\downarrow} = \underbrace{\sigma + \sigma_p}_{\text{supply}\uparrow}$$

This combination of reduced appetite for aggregate risk just when assets become more risky drives the price of aggregate risk π_t up and create a flight to the safety of risk free bonds that depresses the interest rate r_t. In fact, uncertainty shocks may drive the risk-free interest rate below zero, as Figure 2 shows. In a richer model with sticky prices this could lead to a "liquidity trap".

Stochastic risk premia have been extensively studied in the asset pricing literature. Campbell

[38] See Bloom et al. (2012).

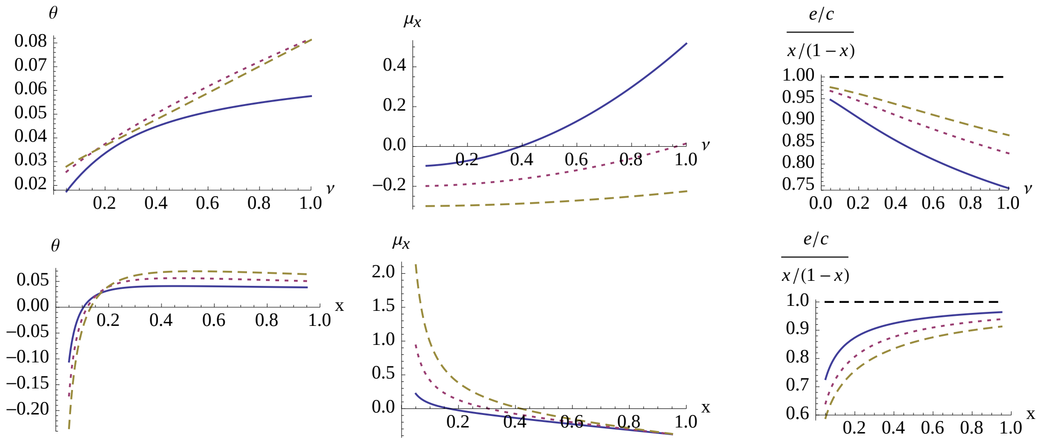

Figure 3: Experts' investment in the market index, θ, the drift of x, μ_x, and experts' consumption relative to their networth $\frac{e/c}{x/(1-x)}$, as functions of ν (above) for $x = 0.25$ (solid), $x = 0.5$ (dotted), and $x = 0.75$ (dashed), and as functions of x (below) for $\nu = 0.25$ (solid), $\nu = 0.5$ (dotted), and $\nu = 0.75$ (dashed). For parameter values see Numerical Solution below.

and Cochrane (1999), for example, introduce habit and obtain stochastic risk-premia. Here, instead, risk premia respond to aggregate shocks due to the presence of financial frictions. In the benchmark without moral hazard, risk premia are constant. He and Krishnamurthy (2011) obtain stochastic risk premia in a similar model where balance sheets play an important role. In their model, agents cannot write contracts on the aggregate state of the economy, so risk premia must rise after a negative shock to induce experts with weak balance sheets to take on aggregate risk. Here instead, agents can share aggregate risk freely, since optimal contracts separate leverage from aggregate risk sharing.

Implementation. Recall that optimal contracts can be implemented using standard financial instruments such as equity and a market index. The expert must keep a fraction ϕ of his own equity and this forces him to keep a fraction ϕ of his idiosyncratic risk, but he can adjust his exposure to aggregate risk using a market index. θ_t can be interpreted as his portfolio investment in a market index normalized with identity loading on the aggregate risk

$$\theta_t = \underbrace{(\sigma + \sigma_{p,t})\left(\frac{x_t - \phi}{x_t}\right)}_{\text{myopic risk-sharing}} + \underbrace{\frac{\sigma_{x,t}}{x_t}}_{\text{hedging}}$$

The first term corresponds to the myopic risk sharing motive. Since aggregate risk pays a premium, both experts and consumers want to buy some of it. They face the same price of risk π_t, so they have the same myopic incentives and they should share aggregate risk proportionally to their wealth. Experts must keep a fraction ϕ of their equity, which already exposes them to a fraction ϕ of aggregate risk. The first use of the market index is to adjust experts' exposure to aggregate risk to achieve proportional risk sharing. The second term captures the hedging motive

$\frac{\sigma_{x,t}}{x} > 0$. Experts want to increase their exposure to aggregate risk in order to take advantage of investment possibility sets. When the economy is hit only by a Brownian TFP shock, the hedging motive disappears and we are left with only the first term.

Figure 3 shows experts portfolio investment in a normalized market index θ as a function of ν and x. Its most striking feature is that θ is positive for most values of (ν, x). Not only are experts not getting insurance against aggregate risk, they are using the financial instruments at their disposal to buy even more aggregate risk than what the equity constraint forces them to. This is consistent with the empirical evidence on bank's balance sheets in Begenau et al. (2013), who show that banks do have large trading positions in derivatives (interest rate swaps) that could allow them to insure against the aggregate risk in their underlying traditional business (lending long and borrowing short), but instead use them to amplify their exposure to aggregate risk.

Short-run, medium-run, and long-run dynamics. In the short-run, an uncertainty shock that increases idiosyncratic risk weakens the balance sheets of experts, who are highly exposed to it ($\sigma_x > 0$). In the medium-run, however, higher idiosyncratic risk leads to stronger balance sheets. This is the result of two effects. First, after an uncertainty shock the premium on idiosyncratic risk α_t is high, so experts have large returns on their net worth compared to consumers. Second, with elasticity of intertemporal substitution greater than 1, $\psi^{-1} > 1$, since the value of net worth is higher for experts relative to consumers after an uncertainty shock ($\Omega_t = \log \xi_t - \log \zeta_t$ is large), experts postpone consumption more than consumers: $\hat{e}/\hat{c} = (\xi/\zeta)^{\frac{\psi-1}{\psi}}$. Figure 3 shows the fraction of aggregate consumption that corresponds to experts relative to the fraction of aggregate wealth that belongs to them, decreasing in ν and increasing in x. Both effects lead to faster accumulation of wealth for experts relative to consumers, and therefore to a gradual strengthening of their balance sheets as captured by x_t. Figure 3 also shows the drift of balance sheets μ_x is increasing in ν and decreasing in x. In fact, in the medium run it is possible for x_t to surpass the level it would have had if idiosyncratic risk had not increased, especially if the increase in idiosyncratic risk is very persistent.

In the long-run, however, both idiosyncratic risk ν and balance sheets x return to their long-run means. The fraction of aggregate wealth that belongs to experts x_t never reaches the boundaries 0 or 1 (i.e. $0 < x_t < 1$) because neither experts nor consumers ever want to have zero wealth. Figure 3 also shows the drift of balance sheets μ_x is positive and diverging to ∞ when $x \to 0$, and becomes negative as $x \to 1$ (so $\frac{\mu_x}{\sigma_x^2}$ diverges to $+\infty$ and $-\infty$ respectively).[39]

The role of elasticity of intertemporal substitution ψ^{-1} and relative risk aversion γ. EZ preferences separate agents' relative risk aversion γ from their elasticity of intertemporal substitution ψ^{-1}. Each plays a different role in the model.

Relative risk aversion controls aggregate risk sharing. If $\gamma < 1$ the substitution effect dominates, and experts have a larger share of aggregate wealth after an uncertainty shock, when the gap in

[39] Without turnover, $\tau = 0$, even though x never reaches the boundaries, the long-run distribution for x will be degenerate, accumulating around 1 if EIS $= \psi^{-1} > 1$.

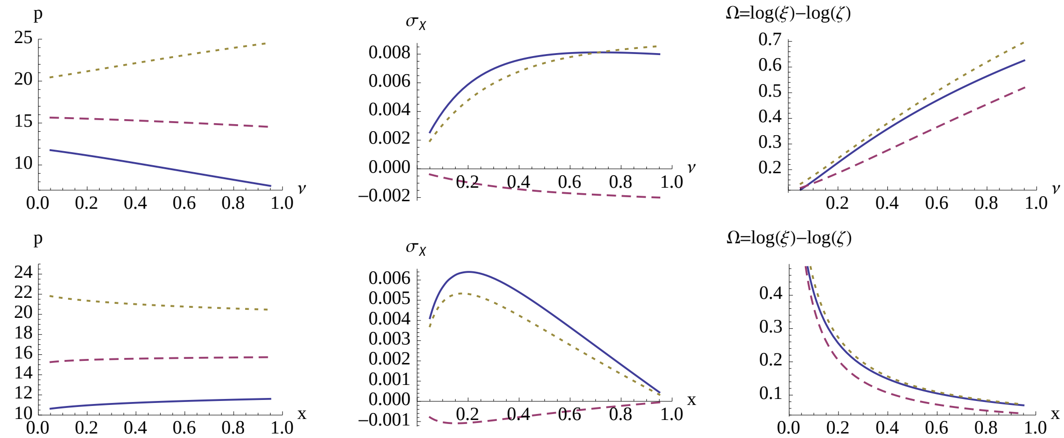

Figure 4: RRA and EIS. The price of capital p, the volatility of x, σ_x, and the gap $\Omega = \log \xi - \log \zeta$ as function of ν for $x = 0.2$ (above) and as functions of x for $v = 0.24$ (below). The solid line is the baseline parametrization with $\gamma = 5$ and $EIS = 1.5$; the dashed line has $\gamma = 0.8$ and $EIS = 1.5$; the dotted line has $\gamma = 3$ and $EIS = \frac{1}{3}$.

investment opportunities Ω_t is large and they prefer to get more "bang for the buck". If $\gamma > 1$, instead, the wealth effect dominates and experts have a smaller fraction of aggregate wealth after an uncertainty shock, since they are already relatively better off compared to consumers (gap Ω_t is large), and they prefer to stabilize utility. For the special log case with $\gamma = 1$ both effects cancel out and the balance sheet channel disappears: $\sigma_n = \sigma_w = \sigma + \sigma_p$ and $\sigma_x = 0$. A balance sheet channel where aggregate risk is concentrated on the balance sheets of experts ($\sigma_x > 0$) therefore requires $\gamma > 1$, which is in line with empirical evidence on risk aversion.

The elasticity of intertemporal substitution is important for growth and asset pricing. After an uncertainty shock idiosyncratic risk ν is high and experts' balance sheets x are endogenously weak. Both drive the idiosyncratic risk premium $\alpha = \gamma \frac{\phi \nu_t}{x_t}$ up. When the EIS $\psi^{-1} > 1$, the price of capital will fall, creating a downturn with low asset prices and growth. This is because agents are quickly convinced by changes in interest rates to adjust their consumption, so r does not move much in equilibrium. This results in the price of capital falling to raise its excess return. When EIS $\psi^{-1} < 1$ instead, interest rates adjust and the price of capital actually goes up after an uncertainty shock (even though ν is high and x low).[40] For the special log case with $\psi = 1$, the price of capital does not depend on ν and x at all: agents consume a constant flow of their net worth $\hat{c} = \hat{e} = \rho$, and the interest rate adjusts to keep p at a level consistent with market clearing for consumption goods[41]:

[40] In terms of wealth and substitution effects, when $\psi^{-1} > 1$ the substitution effect dominates intertemporally: since it's a bad time to invest, invest less now and more in the future. When $\psi^{-1} < 1$ the wealth effect dominates: since investment became less attractive we are poorer in certainty equivalence terms, so consume less (invest more) both today and in the future.

[41] This stark characterization is a result of the setup where only experts can use capital. If consumers are also allowed to use capital to produce with lower productivity, as in Brunnermeier and Sannikov (2012) or Kiyotaki and Moore (1997), the price of capital would fall after an uncertainty shock even in the $EIS = 1$ case, as experts offload capital onto consumers. This reduces aggregate output, so the price of capital must fall in order to reduce consumption and investment. He and Krishnamurthy (2011) instead assume experts are more patient than consumers, so investment falls when x_t falls, and so does the price of capital.

$p = \frac{a-\iota(g)}{\rho}$. A balance sheet recession (in the sense of depressed asset prices and growth amplified by weak balance sheets) therefore requires an EIS $\psi^{-1} > 1$, which is also in line with empirical evidence.[42]

Figure 4 shows the price of capital, σ_x, and the gap in investment opportunities $\Omega = \log \xi - \log \zeta$ for three different parameter values which illustrate these effects. The solid line is the baseline parametrization used above, with $\gamma = 5$ and $\psi^{-1} = 1.5$. As described above, the price of capital is decreasing in ν and increasing in x, and there is a balance sheet channel with $\sigma_x > 0$ everywhere. The dashed line is the second parametrization, with $\gamma = 0.8 \leq 1$ so the substitution effect dominates aggregate risk sharing. The balance sheet channel inverts, $\sigma_x < 0$. Aggregate risk is concentrated on the balance sheets of consumers, so experts' balance sheets strengthen after an uncertainty shock. Since $EIS = 1.5$ is still larger than one, the price of capital is still decreasing in ν and increasing in x, but the balance sheet channel dampens the effects of the uncertainty shock. Notice that the gap Ω is still increasing in ν and decreasing in x, since this reflects the behavior of $\alpha = \gamma \frac{\phi \nu_t}{x_t}$. The inversion of σ_x is caused by the substitution effect dominating the wealth effect in aggregate risk-sharing when $\gamma < 1$. In the third parametrization (dotted line), $\gamma = 3$ and $EIS = \psi^{-1} = 1/3$. Now $\sigma_x > 0$ as in the baseline case, since with $\gamma > 1$ the wealth effect dominates and the gap Ω is still increasing in ν and decreasing in x. However, with the EIS less than 1, the price of capital actually goes up with ν and falls with x.

Numerical solution. I use the following parameter values for illustration purposes. *Preferences*: The discount rate is $\rho = 0.05$, the risk aversion is $\gamma = 5$, while the EIS is set at 1.5 (i.e.: $\psi = 0.66$). Experts retire with Poisson arrival rate $\tau = 0.4$. *Technology*: capital productivity is normalized to $a = 1$. Capital exposure to aggregate risk is $\sigma = 0.03$. *Moral Hazard*: Hedge funds typically keep 20% of returns above a threshold, so following He and Krishnamurthy (2011), I set $\phi = 0.2$. The *investment function* takes the simple form $\iota(g) = 200g^2$, so $g(p) = \frac{p}{400}$ and $\iota(p) = \frac{p^2}{800}$. *Idiosyncratic volatility* is set with a long-run mean of $\bar{\nu} = 0.24$, mean reversion parameter $\lambda = 0.22$, and a loading on the aggregate shock $\sigma_\nu = -0.13$. I use data on idiosyncratic volatility from Campbell et al. (2001), and fit the monthly idiosyncratic standard deviation to a discretized version of (2):

$$\nu_{t+1} - \nu_t = \lambda(\bar{\nu} - \nu_t) + \sigma_\nu \sqrt{\nu_t} \epsilon_{t+1}$$

where the error terms ϵ_t are i.i.d. I then obtain the OLS estimators for the parameters λ, $\bar{\nu}$, and σ_ν.

3.5 General aggregate shocks

When contracts cannot be written on the aggregate state of the economy, the type of aggregate shock hitting the economy is not relevant for the purposes of the balance sheet channel. As long as the aggregate shock depresses asset values and experts are leveraged, their balance sheets will

[42] Campbell and Beeler (2009) use an EIS of 1.5. Gruber (2006) estimates an EIS of 2 based on variation across individuals in the capital income tax rate.

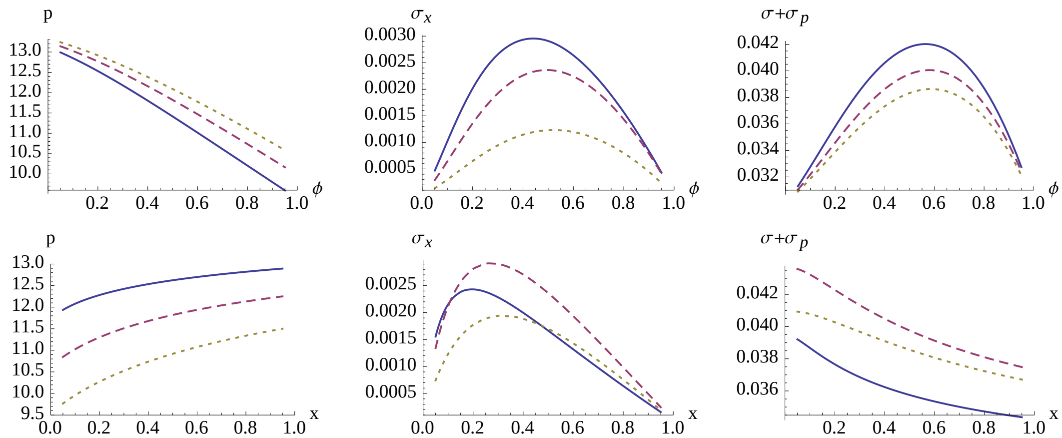

Figure 5: [Financial Shocks] The price of capital p, volatility of x, σ_x, and the total aggregate risk $\sigma + \sigma_p$, as functions of ϕ (above) for $x = 0.25$ (solid), $x = 0.5$ (dotted), and $x = 0.75$ (dashed), and as a function of x (below) for $\phi = 0.25$ (solid), $\phi = 0.5$ (dotted), and $\phi = 0.75$ (dashed).

be disproportionately affected and create a balance sheet recession. When we allow agents to write contracts on the aggregate state of the economy, on the other hand, the type of aggregate shock that hits the economy takes on a prominent role. Agents will share aggregate risk in order to take advantage of the endogenous investment possibility sets. When the wealth effect dominates, experts will have a smaller share of aggregate wealth after shocks that increase the investment opportunity gap between them and consumers $\Omega_t = \log \xi_t - \log \zeta_t$

$$\sigma_{x,t} = (1 - x_t)x_t \frac{1-\gamma}{\gamma} vol(\Omega_t)$$

The gap Ω_t is endogenous, but we can use this equation to understand how different types of aggregate shocks will be shared. As I've shown above, uncertainty shocks will create balance sheet recessions, while Brownian TFP shocks won't. Other aggregate shocks can be studied with the same tools developed above. Here I mention some possibilities. First, aggregate shocks to the long-run growth rate of the economy seem like a natural possibility. After financial crises, growth expectations are often said to have been unduly optimistic before the crash, and the economy to have been "living beyond its means", or in an irrational bubble. However, this could also be the result of negative aggregate shocks to the growth rate of the economy, possibly as a consequence of structural change, and maybe amplified by the balance sheet recession that follows.

Second, demand shocks play a prominent role in business cycle theory. In particular, during liquidity traps, the central bank is unable to stabilize the economy. Liquidity traps seem to occur after big financial crises, so it is natural to ask if the two phenomena are related. On the one hand, the balance sheet channel can help explain how small shocks get amplified and drive the natural rate of interest into negative territory (as it happens in this model). On the other hand, liquidity traps might help explain why the balance sheets of experts are so exposed to aggregate risk in the first place.

Shocks to financial frictions. Since big balance sheet recessions are usually times of financial distress, it makes sense to also explore whether they might be driven by shocks to the financial system itself. Several papers explore the effects of shocks to financial frictions, such as Eggertsson and Krugman (2010), Guerrieri and Lorenzoni (2011) and Buera and Moll (2012).

In the economy I study in this paper, shocks to idiosyncratic risk and shocks to financial frictions turn out to be equivalent. To see why, notice that due to a moral hazard problem, only a fraction $1-\phi$ of idiosyncratic risk is shared and aggregated away, as if idiosyncratic risk was actually $\phi\nu_t$. We may then take $\phi\nu_t$ as the exogenous state variable. Intuitively, it makes no difference to experts whether they must keep a fixed fraction of more idiosyncratic risk, or they must hold a larger fraction of the same underlying amount of idiosyncratic risk. A mathematically equivalent setup for the model takes ν fixed and lets the parameter ϕ_t follow a stochastic process.[43] The only variable affected is θ_t, experts' portfolio share in the market index. Otherwise, the equilibrium is characterized by the same set of equations.

However, it is natural to assume ϕ_t and ν_t follow different stochastic processes. For one thing, $\phi \in [0,1]$ is an important constraint. This will in general change the equilibrium. As an example, Figure 5 shows the equilibrium when $\nu = 0.2$ and the financial friction follows[44]

$$d\phi_t = \lambda(\bar{\phi} - \phi_t) + \sigma_\phi \phi_t (1-\phi_t) dZ_t$$

This process guarantees $\phi_t \in [0,1]$ always, and implies aggregate "financial" volatility vanishes when $\phi_t \to 1$; intuitively, at some point things can't really get any worse. The main features of the equilibrium are unchanged. A bad "financial shock" that increases ϕ_t is amplified by a balance sheet channel ($\sigma_x > 0$ always) and drives asset prices and growth down. However, as ϕ approaches its upper bound, the volatility of ϕ vanishes, and therefore σ_x vanishes too. Aggregate volatility is then non-monotonic in ϕ. This illustrates how different assumptions on the stochastic process driving the model can modify some features of the equilibrium.[45]

4 Financial Regulation

The model has several lessons for optimal policy. In standard models of balance sheet recessions driven by TFP shocks, where contracts cannot be written on the aggregate state of the economy, providing aggregate insurance to experts in order to eliminate the balance sheet channel is a Pareto improving policy. Brunnermeier and Sannikov (2012), for example, show how a social planner can achieve first best allocations in this way.

[43] He and Krishnamurthy (2011) consider a stochastic moral hazard setting in the model. Because they restrict agents' ability to share aggregate risk and assume log preferences, the stochastic moral hazard does not change agents' choices nor does it create a balance sheet channel (which already exists in their model due to the contracting constraints they assume).

[44] I use the same parameter values as in the baseline case, for ease of comparison. So for example, $\sigma_\phi = \sigma_\nu$.

[45] An alternative specification would have ϕ_t follow an arithmetic Brownian motion with reflective boundaries at $\phi_t = 0$ and $\phi_t = 1$. I also solved this case, which has similar features to the solution in Figure 5.

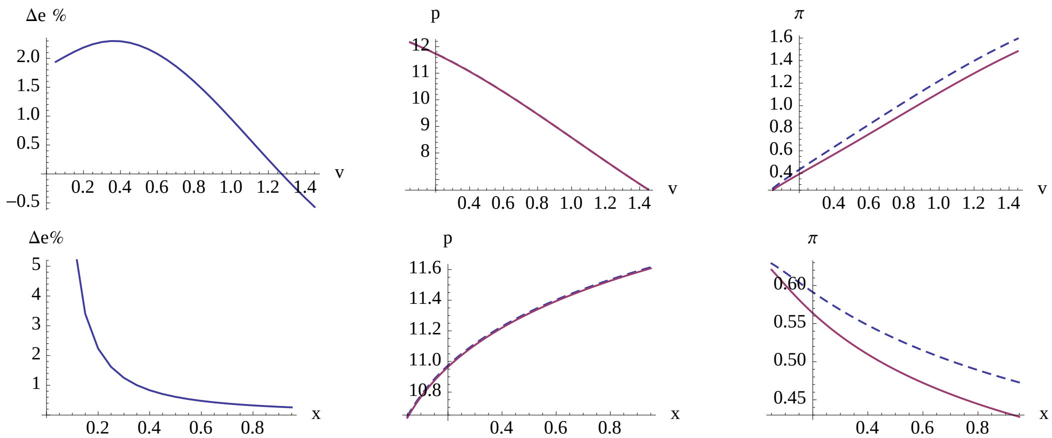

Figure 6: [Financial regulation] Left panels: Equivalent change in consumption for experts as a function of ν, for $x = \frac{1}{5}$ (above) and as a function of x for $\nu = 0.24$ (below). Middle and right panels: the price of capital p and the price of aggregate risk π in the unregulated equilibrium (solid) and under regulation (dashed) as a function of ν for $x = \frac{1}{5}$ (above) and as a function of x for $\nu = 0.24$ (below).

When we allow agents to write contracts on the aggregate state of the economy, two new issues arise. First, experts may react to the policy intervention by taking more risk. They were sharing aggregate risk optimally from an individual point of view, so if government policy somehow changes their underlying exposure to aggregate risk, they will simply try to undo its effects and achieve the same target exposure. Second, even if a regulatory agency could control agents' aggregate risk-sharing, understanding agents' incentives for taking aggregate risk is important for the design of optimal financial regulation. Experts may actually have good reasons to take aggregate risk. In constrast to the incomplete contracts case, eliminating the balance sheet may not be optimal at all.

To illustrate this, consider a policy aimed at eliminating the balance sheet channel by controlling experts' exposure to aggregate risk. The social planner forces experts to buy insurance against aggregate risk in the financial market, so they they end up with aggregate risk proportional to their wealth, i.e. force $\theta_t = \frac{x_t - \phi}{x_t}(\sigma + \sigma_{p,t})$ so that

$$\sigma_{n,t} = \sigma + \sigma_{p,t}$$

$$\implies \quad \sigma_{x,t} = 0$$

The government also carries out a one time wealth transfer between experts and consumers in order to keep consumers indifferent. Normalize the capital stock to $k_0 = 1$, and let $U^{reg}(\nu, x) = \frac{(\zeta^{reg}(1-x)p)^{1-\gamma}}{1-\gamma}$ be the present value of utility for consumers under this policy of financial regulation, and $U^{eq}(\nu, x)$ their value in the unregulated equilibrium. Likewise, $V^{reg}(\nu, x)$ and $V^{eq}(\nu, x)$ are the corresponding value functions for experts. Then the government changes the distribution of wealth from x_0 to x_1 such that:

$$U^{reg}(\nu, x_1) = U^{eq}(\nu, x_0)$$

We can then look at the utility of experts after the policy intervention. The left panels of Figure 6 show the change in the utility of experts, measured in equivalent percentage change in their consumption. This policy is Pareto improving if enacted when ν_t is low, but is counterproductive if enacted when ν_t is high. This allows us to draw two conclusions: 1) the competitive equilibrium is not efficient and can be improved upon by financial regulation, and 2) in contrast to the setting with incomplete contracts, a policy of eliminating the balance sheet channel is not only not optimal: it may even be worse than the unregulated equilibrium.

To understand why the competitive equilibrium is not efficient, notice that there's a pecuniary externality because the price of capital appears in the "skin in the game" constraint. The higher the price of capital, the more idiosyncratic risk an expert with a given $k_{i,t}$ units of capital must be exposed to: $\tilde{\sigma}_{n,i,t} = \phi p_t k_{i,t} \nu_t$. If a social planner could set the price of capital to $p_t = 0$ always[46], the moral hazard problem would disappear, and the first best benchmark of Proposition 1 could be implemented with full idiosyncratic insurance and a balanced growth g^*.[47] Intuitively, there is no private benefit for the expert from stealing something that is worthless.

The competitive equilibrium is inefficient because experts don't internalize the fact that by competing for capital and bidding up its price, they create a moral hazard problem that hampers idiosyncratic risk sharing. When it comes to aggregate risk sharing, they don't internalize that by taking on aggregate risk, they relax or tighten the idiosyncratic risk-sharing problem across states of the world. A social planner that can control θ_t (but not the price of capital) must internalize this effect.[48]

5 Conclusions

In this paper I have shown how the type of aggregate shock hitting the economy can help explain the concentration of aggregate risk that drives balance sheet recessions and financial crises. While we have a good understanding of why the balance sheets of more productive agents matter in an economy with financial frictions, we don't have a good explanation for why these agents are so exposed to aggregate risk. Even if agents face a moral hazard problem that limits their ability to issue equity, this does not prevent them from sharing aggregate risk, which can we accomplished by trading a simple market index. In fact, I show that in standard models of balance sheet recessions driven by Brownian TFP shocks, the balance sheet channel completely vanishes when agents are

[46] for example by taxing away output net of investment costs: $\tau = a - \iota(g_t)$

[47] This assumes the planner can directly control g_t, so that the price of capital doesn't play any allocational role. If it cannot, then an interesting tradeoff appears, but this extreme case illustrates the pecuniary externality very starkly. I study optimal financial regulation in a similar setting in new work in progress.

[48] However, notice that in the first best with no moral hazard aggregate risk is shared proportionally, so it may seem surprising that the financial regulation policy that forces proportional aggregate risk sharing ($\theta_t = \frac{x_t - \phi}{x_t}(\sigma + \sigma_{p,t})$) can be worse than the competitive equilibrium, as shown above. The reason for this is that in this experiment the planner is not allowed to control the price of capital p_t to eliminate the moral hazard problem.

allowed to write contracts contingent on the aggregate state of the economy.

In contrast to TFP shocks, uncertainty shocks can generate an endogenous relative hedging motive that induces more productive agents to take on aggregate risk. Uncertainty shocks are then amplified through a balance sheet channel, depressing asset prices and growth, and triggering a "flight to quality" event with low interest rates and high risk premia. I also show that uncertainty shocks are isomorphic to financial shocks that tighten financial constraints. Finally, the model has lessons for the design of financial regulation. Most importantly, once we understand agents' aggregate risk sharing behavior, we realize they might have good reasons to be highly exposed to aggregate risk. I show how a naive policy of regulating agents' exposure to aggregate risk to eliminate the balance sheet channel can be counterproductive, even if the competitive equilibrium is not constrained efficient.

These results suggest three avenues for future research. The first is to think about optimal financial regulation more carefully. While completely eliminating the balance sheet channel is not optimal, neither is the competitive equilibrium. This suggests the question: how much concentration of aggregate risk is "right"? The second is to consider alternative aggregate shocks. While I have focused on uncertainty (and financial) shocks, the same tools developed in this paper can be used to study other kind of aggregate shocks. For example, shocks to the long-run growth possibilities of the economy can capture some features of financial crises. Indeed, pre-crisis growth projections are often judged unduly optimistic with hindsight. This could be the result of negative aggregate shocks to the growth rate of the economy. Liquidity traps seem to happen after big financial crisis. During liquidity traps, monetary policy is unable to stabilize the economy, so balance sheet problems become more severe. Integrating monetary phenomena into models of balance sheet recessions would allow us to study the interaction of balance sheet recessions and liquidity traps.

References

Adrian, T. and N. Boyarchenko. 2012. "Intermediary Leverage Cycles and Financial Stability." *Federal Reserve Bank of New York Staff Reports* 567.

Adrian, T., E. Etula, and T. Muir. 2011. "Broker-Dealer Leverage and the Cross Section of Stock Returns." Working paper.

Angeletos, M. 2006. "Uninsured Idiosyncratic Investment Risk and Aggregate Saving." Forthcoming in Review of Economic Dynamics.

Bachmann, R. and G. Moscarini. 2011. "Business Cycles and Endogenous Uncertainty." Working paper.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron. 2012. "Volatility, the Macroeconomy, and Asset Prices." NBER working paper 18104.

Bansal, R. and A. Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *The Journal of Finance* 59 (4): pp. 1481–1509.

Begenau, J., M. Piazzesi, and M. Schneider. 2013. "Banks' Risk Exposures." Stanford University.

Bernanke, B. and M. Gertler. 1989. "Agency Costs, Net Worth, and Business Fluctuations." *The American Economic Review* 79 (1): pp. 14–31.

Bernanke, B., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, edited by J.B. Taylor and M. Woodford, Vol. 1, chap. 21, 1341–1393. Elsevier.

Bloom, N. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): pp. 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry. 2012. "Real Uncertain Business Cycles." NBER working paper 18245.

Brunnermeier, M. and L. Pedersen. 2009. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22 (6): 2201 – 38.

Brunnermeier, M. and Y. Sannikov. 2012. "A Macroeconomic Model with a Financial Sector." Working paper, Princeton University.

Buera, F. and B. Moll. 2012. "Aggregate implications of a Credit Crunch." Working paper.

Campbell, J. and J. Beeler. 2009. "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment." *NBER* .

Campbell, J. and J. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–251.

Campbell, J., S. Giglio, C. Polk, and R. Turley. 2012. "An Intertemporal CAPM with Stochastic Volatility." Working paper.

Campbell, J., M. Lettau, Burton G., and Y. Xu. 2001. "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk." *The Journal of Finance* 56 (1): pp. 1–43.

Christiano, L., R. Motto, and M. Rostagno. 2012. "Risk Shocks." Working paper.

Cooley, T., V. Quadrini, and R. Marimon. 2004. "Aggregate Consequences of Limited Contract Enforceability." *Journal of Political Economy* 112 (4): 817–847.

DeMarzo, P. and M. Fishman. 2007. "Optimal Long Term Financial Contracting." *Review of Financial Studies* 20: 2079–127.

Di Tella, S. 2012. "Moral Hazard and Aggregate Risk-Sharing." Working paper.

Diamond, D. and R. Rajan. 2011. “Fear of Fire Sales and the Credit Freeze.” *QJE* forthcoming.

Duffie, D. and L. Epstein. 1992. “Stochastic Differential Utility.” *Econometrica* 60 (2): pp. 353–394.

Edmans, A. and X. Gabaix. 2011. “Tractability in Incentive Contracting.” *Review of Financial Studies* 24 (9): 2865 – 2892.

Eggertsson, G. and P. Krugman. 2010. “Debt, Deleveragin, and the Liquidity Trap: A Fisher-Minsky-Koo approach.” Working paper.

Fernández-Villaverde, J. and J. Rubio-Ramírez. 2010. “Macroeconomics and Volatility: Data, Models, and Estimation.” NBER working paper 16618.

Fernández-Villaverde, J., J. Rubio-Ramírez, A. Guerron-Quintana, and M. Uribe. 2011. “Risk Matters: The Real Effects of Volatility Shocks.” *American Economic Review* 101 (6): 2530–2561.

Fisher, I. 1933. “The Debt-Deflation Theory of of Great Depressions.” *Econometrica* 1 (7): 337–357.

Gabaix, X., A. Krishnamurthy, and O. Vigneron. 2007. “Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market.” *The Journal of Finance* 62 (2): pp. 557–595.

Geanakoplos, J. 2009. “The Leverage Cycle.” In *NBER Macroeconomics Annual*, edited by Keneth Rogoff Daron Acemoglu and Michael Woodford, Vol. 24, 1–65.

Gorton, G. and G. Ordoñez. 2013. “Collateral Crises.” *American Economic Review (forthcoming)* .

Gruber, J. 2006. “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution.” *NBER working paper* .

Guerrieri, V. and G. Lorenzoni. 2011. “Credit Crises, Precautionary Savings, and the Liquidity Trap.” Working paper.

He, Z. and A. Krishnamurthy. 2011. “A Model of Capital and Crises.” Forthcoming, Review of Economic Studies.

Holmstrom, B. and J. Tirole. 1996. “Modeling Aggregate Liquidity.” *The American Economic Review* 86 (2): pp. 187–191.

Kiyotaki, N., M. Gertler, and A. Queralto. 2011. “Financial Crises, Bank Risk Exposure and Financial Crises, Bank Risk Exposure and Government Financial Policy.” Working paper NYU and Princeton.

Kiyotaki, N. and J. Moore. 1997. “Credit Cycles.” *Journal of Political Economy* 105: 211–248.

Krishnamurthy, A. 2003. “Collateral Constraints and the Amplification Mechanism.” *Journal of Economic Theory* 111 (2): 277–292.

Merton, R. 1973. “An Intertemporal Capital Asset Pricing Model.” *Econometrica* 41 (5): 867–887.

Myerson, R. 2012. "A Model of Moral-Hazard Credit Cycles."

Rampini, A. and S. Viswanathan. 2010. "Collateral, risk management, and the distribution of debt capacity." *Journal of Finance* 65: 2293–2322.

Rampini, A. and S. Viswanathan. 2013. "Collateral and capital structure." *Journal of Fianancial Economcis* 109: 466–492.

Shleifer, A. and R. W. Vishny. 1992. "Liquidation Values and Debt Capacity: A Market Equilibrium Approach." *The Journal of Finance* 47 (4): 1343 – 1366.

Simsek, A. 2013. "Speculation and risk sharing with new financial assets." *Quaterly Journal of Economics (forthcoming)* .

Sraer, D., T. Chaney, and D. Thesmar. 2011. "The Collateral Channel: The Collateral Channel: How Real Estate Shocks affect Corporate Investment." Working paper.

Sun, Y. and Y. Zhang. 2009. "Individual risk and Lebesgue extension without aggregate uncertainty." *Journal of Economic Theory* (144): 432–443.

6 Appendix A: Contracting environment

In this Appendix I derive the "skin in the game" financial friction from a moral hazard problem. To make the contractual setting clear, I use a discrete time approximation to continuous-time. I introduce a moral hazard problem where the expert can divert capital and obtain a pecuniary private benefit. Notice that because experts are risk averse and the market does not price idiosyncratic risk, in the first best without moral hazard there is full insurance against idiosyncratic risk. In fact, if the expert could commit to long-term contracts that control his consumption, the first best would be implementable, even with unobservable capital diversion. The intuition is that the expert cannot do anything with his stolen funds, and his continuation utility does not depend on the observed return in the first best, so there is no incentive for the expert to steal. In order to obtain a binding moral hazard problem, I restrict agents to short-term contracts as in Brunnermeier and Sannikov (2012) or He and Krishnamurthy (2011). This has the advantage of yielding a tractable solution and making comparisons with the literature straightforward.[49] The same optimal contract arises if long-term contracts are allowed, but experts can offer new contracts to the principal at any time.

Time is discrete, with time interval Δ and infinite horizon: $t \in \mathcal{T} = \{0, \Delta, 2\Delta, ...\}$. I will later take the limit to continuous-time $\Delta \to 0$. A risk averse expert can use capital to produce, and wants to raise funds from and share risk with a principal who is risk-neutral with respect to idiosyncratic risk. At the beginning of period t, the agent consumes $e_t\Delta$, buys capital k_t and uses it to produce $ak_t\Delta$ and invest $\iota(g_t)k_t\Delta$ in order to make his capital grow at an expected rate $g_t\Delta$. The agent's

[49] An alternative would be to make consumption not observable. The form of the contract would change, but the main message would remain, i.e. incentive compatibility does not put constraints on aggregate risk sharing.

consumption e_t, capital k_t and growth g_t are observable and contractible. The expert then observes an aggregate shock[50] $z_t \in \{-1, 1\}$ and an idiosyncratic shock $w_t \in \{-1, 1\}$, independent and both i.i.d. with binomial distribution with equal probability of each node.[51] After observing these shocks, he decides how much capital to divert or steal $s_t \geq 0$.[52] As a result, capital at the end of the period is

$$\tilde{k}_t = k_t(1 + g_t\Delta + \sigma z_t\sqrt{\Delta} + \nu_t w_t\sqrt{\Delta} - s_t) \geq 0$$

where $\sigma \in \mathbb{R}^d$ and the idiosyncratic volatility $\nu_t \in \mathbb{R}_+$ follows a stochastic process I will introduce below. The consumption e_t, capital stock k_t, the growth g_t, and the resulting capital stock $\tilde{k}_t$ are observed by the principal, as well as the aggregate shock z_t. Since we will always want to implement no stealing, $s_t = 0$, the agent can effectively steal $s_t \in \{0, 2\nu_t w_t\sqrt{\Delta}\}$ only if $w_t = 1$, and cannot steal if $w_t = -1$. In addition, to ensure $\tilde{k}_t \geq 0$, I restrict g_t such that $1 + g\Delta - \sigma\sqrt{\Delta} - \nu_t\sqrt{\Delta} \geq 0$. Stealing works in the following way: for each unit of capital stolen, the agent gets only $\phi \in (0, 1)$ units, which he must immediately sell at price $p_{t+\Delta}$. The parameter ϕ captures the severity of the moral hazard problem. With $\phi = 0$, for example, there is no moral hazard. Because $\phi < 1$, stealing is inefficient: in the first best where s_t is observable there is no stealing ($s_t = 0$ always). In fact, no stealing will always be optimal.[53]

There is a complete financial market with state price density η_t. The financial market does not price idiosyncratic risk (η_t does not depend on the history of idiosyncratic shocks w^t), although cash flows contingent on w_t can be traded and priced. The price of capital is p_t, and idiosyncratic volatility ν_t also follows a stochastic process. I take the price of capital p_t, the state price density η_t, and the idiosyncratic volatility ν_t as exogenous functions of the history of aggregate shocks $z^t = (z_0, z_\Delta, ..., z_{t-\Delta})$, with $z^0 = \emptyset$, e.g. $p_t = p(z^t)$. To simplify notation, I will typically suppress their dependence on z^t. The particular process they follow is not important, but it can be chosen so that when we take the limit to continuous time, $\Delta \to \infty$ they converge to the processes in Section 2.

The history for the expert $h_t = (z^t, w^t, s^t)$ includes the history of aggregate shocks z^t, the history of idiosyncratic shocks $w^t = \{w_0, w_\Delta, ...w_{t-\Delta}\}$ and the history of stealing $s^t = \{s_0, s_\Delta, ..., s_{t-\Delta}\}$. At the beginning of each period t, after history h_t, the expert has a bank account with n_t funds. He sells a contract $\mathcal{C}_t(h_t) = (c_t, k_t, g_t, F_{t+\Delta})(h_t)$, where $F_{t+\Delta}(h_t; z_t, \tilde{k}_t)$ is a cash payment from the expert to the principal at the end of the period. After the contract is executed, the agent and the principal separate and never meet again. Assume his net worth is observable when he writes the contract, but this is wlog.

[50] I will focus on the case with only one aggregate shock, $z_t = \{-1, 1\}$. In general if we want d shocks we will have $d + 1$ nodes for z_t, i.e. $z_t \in Z$ with $\#Z = d + 1$.

[51] The fact that w_t has binomial distribution is not essential for the results. We could have $w_t \in [L, H]$ and a distribution with full support.

[52] This timing convention, used also in Edmans and Gabaix (2011), simplifies the analysis.

[53] The argument is standard, see DeMarzo and Fishman (2007)

We can replace $\tilde{k}_t$ by the return of capital

$$R_t(z^t; z_t, \tilde{k}_t) = \frac{p_{t+\Delta}(z^t, z_t)}{p_t(z^t)} \frac{\tilde{k}_t}{k_t(h_t)} = \frac{p_{t+\Delta}}{p_t}(1 + g_t\Delta + \sigma z_t\sqrt{\Delta} + \nu_t w_t\sqrt{\Delta} - s_t)$$

which carries the same information, and write $F_{t+\Delta}(h_t; z_t, R_t)$.

The market prices the contract $\mathcal{C}_t(h_t)$

$$J_t(h_t) = \mathbb{E}\left[\frac{\eta_{t+\Delta}}{\eta_t} F_{t+\Delta}(h_t; z_t, R_t)|h_t\right]$$

Under this contract, the net worth of the expert satisfies the budget constraints

$$n_t(h_t) + J_t(h_t) = p_t k_t(h_t) + e_t(h_t)\Delta \tag{18}$$

$$\begin{aligned} n_{t+\Delta}(h_{t+\Delta}) &= p_t k_t(h_t) R_t + k_t(h_t)(a - \iota(g_t(h_t)))\Delta - F_{t+\Delta}(h_t; z_t, R_t) && (19) \\ &\quad + \phi p_{t+\Delta} k_t(h_t) s_t(h_t, z_t, w_t) && (20) \end{aligned}$$

We can write wlog

$$F_{t+\Delta}(h_t; z_t, R_t) = \bar{F}_{t+\Delta}(h_t) + \sigma_{F,t}(h_t) z_t\sqrt{\Delta} + \tilde{\sigma}_{F,t}(h_t; z_t)(R_t - 1)$$

To make the problem well defined, we impose a solvency constraint: $n_t(h_t) \geq 0$ always.

The expert has recursive preferences

$$U_t = \left\{\rho e_t^{1-\psi}\Delta + (1-\rho\Delta)\mathbb{E}_t\left[U_{t+\Delta}^{1-\gamma}\right]^{\frac{1-\psi}{1-\gamma}}\right\}^{\frac{1}{1-\psi}} \tag{21}$$

His value function $U_t(h_t) = V(z^t, n_t(h_t))$ depends on aggregate conditions captured by the history z^t, and his net worth $n_t(h_t)$ (i.e. it does not depend on his previous history). A contract $\mathcal{C}$ is *feasible* at h_t with net worth $n_t(h_t)$ if it satisfies the solvency constraint $n_{t+\Delta}(h_t; z_t, R_t) \geq 0$ for all z_t, w_t, s_t. A feasible contract $\mathcal{C}_t$ is *incentive compatible* if it's optimal for the agent to not steal for all z_t and w_t.

$$0 \in \arg \max_{s(h_t; z_t, w_t)} \left\{V(z^{t+\Delta}, n_{t+\Delta}(h_t; z_t, R_t))\right\}$$

where R_t depends on $s_t(h_t; z_t, w_t)$, and $s = 0$ denotes $s_t(h_t; z_t, w_t) = 0\, \forall z_t, w_t$. Since the expert can only steal when $w_t = 1$, we have for all z_t

$$p_t k_t(h_t)\frac{p_{t+\Delta}}{p_t}(1 + g_t(h_t)\Delta + \sigma z_t\sqrt{\Delta} - \nu_t\sqrt{\Delta}) - \bar{F}_{t+\Delta}(h_t) - \sigma_{F,t}(h_t) z_t\sqrt{\Delta}$$

$$-\tilde{\sigma}_{F,t}(h_t; z_t)(\frac{p_{t+\Delta}}{p_t}(1 + g_t(h_t)\Delta + \sigma z_t\sqrt{\Delta} - \nu_t\sqrt{\Delta}) - 1) + \phi p_{t+\Delta} k_t(h_t) 2\nu_t\sqrt{\Delta} \leq$$

$$p_t k_t(h_t)\frac{p_{t+\Delta}}{p_t}(1+g_t(h_t)\Delta+\sigma z_t\sqrt{\Delta}+\nu_t\sqrt{\Delta})-\bar{F}_{t+\Delta}(h_t)-\sigma_{F,t}(h_t)z_t\sqrt{\Delta}$$
$$-\tilde{\sigma}_{F,t}(h_t;z_t)(\frac{p_{t+\Delta}}{p_t}(1+g_t(h_t)\Delta+\sigma z_t\sqrt{\Delta}+\nu_t\sqrt{\Delta})-1)$$

which rearranging yields

$$\tilde{\sigma}_{F,t}(h_t;z_t)\leq p_t k_t(h_t)(1-\phi)$$

This is the "skin in the game" constraint: the expert can't offload all of his return on the market. He must keep a fraction ϕ, so that stealing costs him more than what it earns him. Notice, however, that it imposes no constraints on $\sigma_{F,t}(h_t)$. Moral hazard does not restrict the expert's ability to share aggregate risk.

Defining $\theta_t(h_t)=-\frac{\sigma_{F,t}(h_t)}{n_t(h_t)}$ and $\tilde{\phi}_t(h_t;z_t)=1-\frac{\tilde{\sigma}_{F,t}(h_t;z_t)}{p_t k_t(h_t)}$, the IC constraints says that $\tilde{\phi}_t(h_t;z_t)\geq\phi$ for all h_t and z_t. Now write the budget constraints in terms of z and w (under $s=0$) as follows

$$n_{t+\Delta}(h_{t+\Delta})=pk_t(h_t)+k_t(h_t)(a-\iota(g_t(h_t)))\Delta+X(h_t,z_t)+Y_t(h_t,z_t)w_t\sqrt{\Delta} \tag{22}$$

and

$$n(h_t)+J_t(h_t)=pk_t(h_t)+e_t(h_t)\Delta \tag{23}$$

with

$$X_t(h_t,z)=n_t(h_t)\theta_t(h_t)z_t\sqrt{\Delta}+p_t k_t(h_t)\tilde{\phi}_t(h_t,z_t)(\frac{p_{t+\Delta}}{p_t}(1+g_t\Delta+\sigma z_t\sqrt{\Delta})-1)-\bar{F}_{t+\Delta}(h_t)$$

$$Y_t(h_t,z_t)=p_t k_t(h_t)\tilde{\phi}_t(h_t,z_t)\frac{p_{t+\Delta}}{p_t}\nu_t$$

and

$$J_t(h_t)=\mathbb{E}\left[\frac{\eta_{t+\Delta}}{\eta_t}p_t k_t(h_t)(R-1)\right]-\mathbb{E}\left[\frac{\eta_{t+\Delta}}{\eta_t}X_t(h_t,z_t)\right]-\underbrace{\mathbb{E}\left[\frac{\eta_{t+\Delta}}{\eta_t}Y_t(h_t,z_t)w_t\sqrt{\Delta}\right]}_{=0}$$

where the last term is 0 because $\frac{\eta_{t+\Delta}}{\eta_t}$ is a function of h_t and z_t only (does not involve w_t).

An incentive compatible contract $e_t(h_t)$, $k_t(h_t)$, $g_t(h_t)$, $\bar{F}_{t+\Delta}(h_t)$, $\theta_t(h_t)$ and $\tilde{\phi}_t(h_t,z_t)$ is *optimal* at h_t with net worth $n_t(h_t)$ if it maximizes

$$V(z^t,n_t(h_t))=\max\left\{\rho e_t(h_t)^{1-\psi}\Delta+(1-\rho\Delta)\mathbb{E}_t\left[V(z^{t+\Delta},n_{t+\Delta}(h_{t+\Delta}))\right]^{\frac{1-\psi}{1-\gamma}}\right\}^{\frac{1}{1-\psi}} \tag{24}$$

subject to the budget constraints (22) and (23), the solvency constraint $n_{t+\Delta}(h_{t+\Delta})\geq 0$ and the IC constraint $\tilde{\phi}(h_t,z_t)\geq\phi$. The expert can therefore use $\tilde{\phi}_t(h_t,z_t)$ to control $Y_t(h_t,z_t)$ without affecting either J_t or $X_t(h_t,z_t)$, which he completely controls with $\theta_t(h_t)$ and $\bar{F}_{t+\Delta}(h_t)$. Since he is risk averse he will want to make $Y_t(h_t,z_t)$ as small as possible, and since $p_{t+\Delta}k_t(h_t)\nu_t\geq 0$, this is accomplished by setting

$$\tilde{\phi}_t(h_t, z_t) = \phi \, \forall h_t, z_t$$

This makes sense: because the market does not price idiosyncratic risk and the agent is risk averse, he wants to get rid of as much of it as possible. Without moral hazard this would mean offloading all of his return R_t on the market. With moral hazard, the "skin in the game" IC constraint is binding in all states.

We assumed that his net worth is observable when writing the contract, but this is wlog. Notice that since the IC constraint does not depend on the expert's net worth, the expert will always choose to fully reveal his net worth if he had the chance to hide it, because it relaxes the solvency constraint.

The budget constraints actually eliminate $\bar{F}_{t+\Delta}(h_t)$, so the Bellman equation (24) characterizes the solution to the sequential portfolio problem of choosing $e_t(h_t)$, $k_t(h_t)$, $g_t(h_t)$, $\theta_t(h_t)$, and $\tilde{\phi}(h_t, z_t)$ to maximize U_0, defined recursively according to (21), subject to the budget constraints (22) and (23), the solvency constraint $n_t(h_t) \geq 0$ (this is included in the definition of feasible), and the IC constraint $\tilde{\phi}_t(h_t; z_t) \geq \phi$ for all h_t and z_t.

Finally, using $\tilde{\phi}(h_t, z_t) = \phi$ always, we can take the limit to continuous-time $\Delta \to 0$, eliminating terms smaller than Δ, to obtain the experts' problem in Section 2[54]

7 Appendix B: omitted proofs

7.1 Proof of Proposition 1

Proof. Without any financial frictions, idiosyncratic risk can be perfectly shared and has zero price in equilibrium. Capital then must be priced by arbitrage

$$g_{i,t} + \mu_{p,t} + \sigma\sigma'_{p,t} + \frac{a - \iota\left(g_{i,t}\right)}{p_t} - r_t = \pi_t\left(\sigma + \sigma_{p,t}\right) \tag{25}$$

and experts face the same portfolio problem as consumers, with the exception of the choice of the growth rate g, pinned down by the static FOC

$$\iota'\left(g_t\right) = p_t$$

We have, in effect, a standard representative agent model with a stationary growth path with risk-free interest rate:

$$r_t = \rho + \psi g_t - \frac{1}{2}\left(1 + \psi\right)\gamma\sigma^2$$

and price of aggregate risk

$$\pi_t = \gamma\sigma$$

[54] Of course, for this we need to chose the stochastic process for η_t, p_t and ν_t so that they converge to the continuous-time processes of Section 2.

In a stationary equilibrium the price of capital is constant so we have $\mu_{p,t} = \sigma_{p,t} = 0$, and replacing all of this in (25) gives (7). For the agents' problem to be well defined we need $\rho - (1-\psi)g^* + (1-\psi)\frac{\gamma}{2}\sigma^2 > 0$. $\square$

7.2 Proof of Proposition 2

Proof. Standard from homothetic preferences and taking first order conditions in the HJB equations. $\square$

7.3 Proof of Proposition 3

Proof. From (15) we see that if $\sigma_\nu = 0$ then $\sigma_{x,t} = 0$. Furthermore, the idiosyncratic volatility of capital, ν_t is then deterministic because it is the solution to an ODE (2). We can replace ν_t with t in the Markov equilibrium (and obtain a time-dependent equilibrium). The only possibly stochastic state variable is x_t, but we have seen that it can only have a stochastic drift. However, since all equilibrium objects are functions of x and time t, then by (9) we see that x_t is the deterministic solution to a time-dependent ODE. $\square$

7.4 Proof of Proposition 4

Proof. We solve the experts' problem with the added constraint that $\theta_t = \tilde{\theta}_t B$. The FOC for $\tilde{\theta}_t$ yields:

$$B\sigma'_{n,t} = B\frac{\pi_t}{\gamma} - \frac{\gamma-1}{\gamma}B\sigma'_{\xi,t}$$

with an analogous condition for consumers. Using as before $\sigma_{x,t} = x_t(1-x_t)(\sigma_{n,t} - \sigma_{w,t})$ we obtain

$$B\sigma'_{x,t} = 0$$

In addition, recall the formula

$$\sigma_{n,t} = \phi\frac{1}{x_t}(\sigma + \sigma_{p,t}) + \tilde{\theta}_t B = \phi\frac{1}{x_t}(\sigma + \frac{p_x}{p}\sigma_{x,t}) + \tilde{\theta}_t B$$

Because σ is in the row space of B, we can write $\sigma = \kappa B$ for some $\kappa \in \mathbb{R}^{d'}$, and using $\sigma_{x,t} = x_t(\sigma_{n,t} - \sigma - \sigma_{p,t})$ we obtain

$$\sigma_{x,t} = x_t\frac{\frac{\phi}{x_t}\kappa + \tilde{\theta}_t - \kappa}{(1 - \frac{p_x}{p}(\phi - x))}B = m_t B$$

So we get

$$B\sigma'_{x,t} = BB'm'_t = 0$$

and because B is of full rank this implies $\sigma_{x,t} = 0$. By the same argument as in Proposition 3, x_t is deterministic.

Corollary: Experts' equity has a loading on aggregate risk $\sigma + \sigma_{p,t} = \sigma$, since $\sigma_{p,t} = 0$ in this equilibrium. If agents can trade a market index, σ is in the row space of B. $\square$

8 Appendix C: solving for the equilibrium

The strategy to solve for the equilibrium when uncertainty shocks hit the economy is to first use optimality and market clearing conditions to obtain expressions for equilibrium objects in terms of the stochastic processes for p, ξ, ζ, and then use Ito's lemma to map the problem into a system of partial differential equations. In order to obtain a non-degenerate stationary long-run distribution for x, I also introduce turnover among experts: they retire with independent Poisson arrival rate $\tau > 0$. When they retire they don't consume their wealth right away, they simply become consumers. Without turnover, experts want to postpone consumption and approach $x_t \to 1$ as $t \to \infty$. Turnover modifies experts' HJB slightly:

$$\frac{\rho}{1-\psi} = \max_{\hat{c},g,\hat{k},\theta} \frac{\hat{e}^{1-\psi}}{1-\psi}\rho\xi_t^{\psi-1} + \frac{\tau}{1-\gamma}\left(\left(\frac{\zeta}{\xi}\right)^{1-\gamma} - 1\right) + \mu_n - \hat{e} + \mu_\xi - \frac{\gamma}{2}\left(\sigma_n^2 + \sigma_\xi^2 - 2(1-\gamma)\sigma_n\sigma_\xi + \tilde{\sigma}_n^2\right) \tag{26}$$

With Poisson intensity τ the expert retires and becomes a consumer, losing the continuation utility of an expert, but gaining that of a consumer. For this reason, consumers' wealth multiplier ζ appears in experts' HJB equation. Consumers have the same HJB equation as before. The FOC for consumption for experts and consumers are:

$$\hat{e} = \rho^{\frac{1}{\psi}}\xi^{\frac{\psi-1}{\psi}}$$

$$\hat{c} = \rho^{\frac{1}{\psi}}\zeta^{\frac{\psi-1}{\psi}}$$

So market clearing in the consumption goods market requires:

$$\rho^{\frac{1}{\psi}}\left(\xi^{\frac{\psi-1}{\psi}}x + \zeta^{\frac{\psi-1}{\psi}}(1-x)\right) = \frac{a-\iota}{p} \tag{27}$$

Equation (15) provides a formula for σ_x using $\Omega_\nu = \frac{\xi_\nu}{\xi} - \frac{\zeta_\nu}{\zeta}$ and $\Omega_x = \frac{\xi_x}{\xi} - \frac{\zeta_x}{\zeta}$:

$$\sigma_x = \frac{(1-x)x\frac{1-\gamma}{\gamma}\left(\frac{\xi_\nu}{\xi} - \frac{\zeta_\nu}{\zeta}\right)}{1-(1-x)x\frac{1-\gamma}{\gamma}\left(\frac{\xi_x}{\xi} - \frac{\zeta_x}{\zeta}\right)}\sigma_\nu\sqrt{\nu_t}$$

We can use Ito's lemma to obtain expressions for

$$\sigma_p = \frac{p_\nu}{p}\sigma_\nu\sqrt{\nu_t} + \frac{p_x}{p}\sigma_x, \quad \sigma_\xi = \frac{\xi_\nu}{\xi}\sigma_\nu\sqrt{\nu_t} + \frac{\xi_x}{\xi}\sigma_x \quad \sigma_\zeta = \frac{\zeta_\nu}{\zeta}\sigma_\nu\sqrt{\nu_t} + \frac{\zeta_x}{\zeta}\sigma_x$$

and the definition of σ_x from (9) to obtain and expression for

$$\sigma_n = \sigma + \sigma_p + \frac{\sigma_x}{x}$$

Then we use experts FOC for aggregate risk sharing (11) to obtain an expression for the price of aggregate risk

$$\pi = \gamma\sigma_n + (\gamma - 1)\,\sigma_\xi$$

Consumers' exposure to aggregate risk is taken from (12):

$$\sigma_w = \frac{\pi}{\gamma} - \frac{\gamma - 1}{\gamma}\sigma_\zeta$$

Experts' exposure to idiosyncratic risk is given by $\tilde{\sigma}_n = \frac{\phi}{x}\nu$. We can now use consumers' budget constraint to obtain the drift of their wealth (before consumption)

$$\mu_w = r + \pi\sigma_w$$

and plugging into their HJB equation we obtain an expression for the risk-free interest rate

$$r = \frac{\rho}{1-\psi} - \frac{\psi}{1-\psi}\rho^{\frac{1}{\psi}}\zeta^{\frac{\psi-1}{\psi}} - \pi\sigma_w - \mu_\zeta + \frac{\gamma}{2}\left(\sigma_w^2 + \sigma_\zeta^2 - 2(1-\gamma)\sigma_w\sigma_\zeta\right)$$

where the only term which hasn't been solved for yet is μ_ζ. We use the FOC for capital (10) and the expression for the risk-free interest rate and plug into the formula for μ_n from equation (5) to get

$$\mu_n = r + \underbrace{\gamma\frac{1}{x^2}(\phi\nu)^2}_{\alpha_t p\hat{k}(\phi\nu)} + \pi\sigma_n$$

In equilibrium experts receive the risk-free interest on their net worth, plus a premium for the idiosyncratic risk they carry through capital, $\gamma\frac{1}{x^2}(\phi\nu)^2$, and a risk premium for the aggregate risk they carry, $\pi\sigma_n$. This allows us to compute the drift of the endogenous state variable x in terms of known objects, from (9) (appropriately modified for turnover) and (10)

$$\mu_x = \mu_n - \hat{e} - \tau + \frac{a-\iota}{p} - r - \pi(\sigma + \sigma_p) - \frac{\gamma}{x}(\phi\nu)^2 + (\sigma + \sigma_p)^2 - \sigma_n(\sigma + \sigma_p)$$

Turnover works to reduce the fraction of aggregate wealth that belongs to experts through the term $-\tau$. Using Ito's lemma we get expressions for the drift of p, ξ, and ζ:

$$\mu_p = \frac{p_\nu}{p}\mu_\nu\nu + \frac{p_x}{p}\mu_x + \frac{1}{2}\left(\frac{p_{\nu\nu}}{p}\sigma_\nu^2\nu + 2\frac{p_{\nu x}}{p}\sigma_\nu\sqrt{\nu_t}\sigma_x + \frac{p_{xx}}{p}\sigma_x^2\right)$$

$$\mu_\xi = \frac{\xi_\nu}{\xi}\mu_\nu\nu + \frac{\xi_x}{\xi}\mu_x + \frac{1}{2}\left(\frac{\xi_{\nu\nu}}{\xi}\sigma_\nu^2\nu + 2\frac{\xi_{\nu x}}{\xi}\sigma_\nu\sqrt{\nu_t}\sigma_x + \frac{\xi_{xx}}{\xi}\sigma_x^2\right)$$

$$\mu_\zeta = \frac{\zeta_\nu}{\zeta}\mu_\nu \nu + \frac{\zeta_x}{\zeta}\mu_x + \frac{1}{2}\left(\frac{\zeta_{\nu\nu}}{\zeta}\sigma_\nu^2 \nu + 2\frac{\zeta_{\nu x}}{\zeta}\sigma_\nu \sqrt{\nu_t}\sigma_x x + \frac{\zeta_{xx}}{\zeta}\sigma_x^2\right)$$

Finally, experts' HJB (26) and their FOC for capital (10) provide two second order partial differential equations in p, ξ, and ζ. Together with the market clearing condition for consumption (27) they characterize the Markov equilibrium. The idiosyncratic volatility ν_t moves exogenously in $(0, \infty)$ and $x \in (0, 1)$. The system never reaches any of its boundaries.

Numerical Algorithm. The system of partial differential equations can be solved in several ways. I use a finite difference scheme. I start with a finite horizon problem which modifies the HJB equations and FOC for capital by adding a time derivative when computing the drifts. Now we must look for p, ξ, and ζ as functions of (ν, x, t). Notice, however, that if we find a stationary point such that the time derivatives vanish, we have found a solution for the infinite horizon equilibrium.

Starting from some terminal values for p, ξ, and ζ at $t = T$, we can solve the system backwards using any standard integrator such as Runge-Kutta 4. Because the market clearing condition for consumption is an algebraic constraint, it is easier to differentiate with respect to time and obtain a differential equation. Terminal values at $t = T$ are not particularly important as long as they satisfy the market clearing condition.

As we move backwards in time, the solution should approach the solution for the time-homogenous system that characterizes the infinite horizon equilibrium. This suggests that the algorithm will converge to the desired solution for a wide variety of terminal conditions. In any case, as long as it converges to the stationary solution, we have found the infinite horizon equilibrium, and this can be verified.