

NBER WORKING PAPER SERIES

RISK PREMIUM SHOCKS CAN CREATE INEFFICIENT RECESSIONS

Sebastian Di Tella
Robert E. Hall

Working Paper 26721
<http://www.nber.org/papers/w26721>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020

Hall's research is supported by the Hoover Institution. We thank Chris Tonetti, Chad Jones, Emmanuel Farhi, Tom Winberry, Alp Simsek, Moritz Lenel, Rohan Kekre, and seminar participants in Princeton, LSE, Fed Board, AEA meetings, and Cleveland Fed. We thank Bernard Herskovic, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh for sharing their data with us. Neither author has any financial relationships calling for disclosure. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w26721.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Sebastian Di Tella and Robert E. Hall. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Risk Premium Shocks Can Create Inefficient Recessions
Sebastian Di Tella and Robert E. Hall
NBER Working Paper No. 26721
January 2020
JEL No. E21,E22,E32

ABSTRACT

We develop an equilibrium theory of business cycles driven by spikes in risk premiums that depress business demand for capital and labor. Aggregate shocks increase firms' uninsurable idiosyncratic risk and raise risk premiums. We show that risk shocks can create quantitatively realistic recessions, with contractions in employment, consumption, and investment. Business cycles are inefficient—output and employment fall too much during recessions, compared to the constrained-efficient allocation, and consumption should rise. Optimal policy involves stimulating employment and consumption during recessions.

Sebastian Di Tella
Stanford University
Graduate School of Business
655 Knight Way
Stanford, CA 94305
and NBER
sditella@stanford.edu

Robert E. Hall
Hoover Institution
Stanford University
Stanford, CA 94305-6010
and NBER
rehall@gmail.com

Risk Premium Shocks Can Create Inefficient Recessions*

Sebastian Di Tella
Stanford GSB

Robert Hall
Hoover Institution and Department of Economics, Stanford University

January 2020

Abstract

We develop an equilibrium theory of business cycles driven by spikes in risk premiums that depress business demand for capital and labor. Aggregate shocks increase firms' uninsurable idiosyncratic risk and raise risk premiums. We show that risk shocks can create quantitatively realistic recessions, with contractions in employment, consumption, and investment. Business cycles are inefficient—output and employment fall too much during recessions, compared to the constrained-efficient allocation, and consumption should rise. Optimal policy involves stimulating employment and consumption during recessions.

JEL: E32, E21, E22

Keywords: Business cycles, risk premium, recession, precautionary saving

1 Introduction

Market economies experience recurrent recessions with sharp contractions in economic activity. Why do employment, investment, and consumption suddenly fall, and what is the appropriate policy response? In this paper we explore a risk-premium view of business cycles: recessions are periods of heightened economic uncertainty when firms and investors shrink from risk. We propose a model of business cycles driven by spikes in risk premiums that depress business demand for labor and capital.

*sditella@stanford.edu and rehall@stanford.edu. Hall's research is supported by the Hoover Institution. We thank Chris Tonetti, Chad Jones, Emmanuel Farhi, Tom Winberry, Alp Simsek, Moritz Lenel, Rohan Kekre, and seminar participants in Princeton, LSE, Fed Board, AEA meetings, and Cleveland Fed. We thank Bernard Herskovic, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh for sharing their data with us.

The premise of our model is that businesses face significant uninsurable idiosyncratic risk which spikes during downturns, driving risk premiums up. Figure 1 shows idiosyncratic risk in stock returns, from Herskovic et al. [2016]. Spikes in idiosyncratic risk are clearly visible during recessions, especially during the Great Depression and the 2008 financial crisis. We show that risk shocks can create quantitatively realistic business cycles, with co-movement among employment, consumption, and investment. Furthermore, these economic fluctuations are inefficient—output and employment fall too much during recessions, compared to the constrained-efficient allocation, and consumption should rise. Optimal policy involves stimulating employment and consumption during recessions.

A long tradition attributes business cycles to time-varying risk premiums, dating back, at least, to chapter 12 of the General Theory (Keynes [1936]). The co-movement pattern of macroeconomic aggregates during business cycles, however, poses a challenge to this view. It is hard to explain why employment, consumption, and investment all simultaneously contract. As Barro and King [1984] noted, as long as the productive capacity of the economy is not affected by the shock, contractions in investment demand produce expansions in consumption in standard equilibrium models.

In response to this challenge, modern macroeconomics has focused on two broad views of business cycles. The RBC model highlights the role of TFP shocks that reduce productive capacity—recessions are periods of low productivity. The New Keynesian model highlights nominal rigidities and suboptimal monetary policy—recessions are periods when the central bank does not or cannot reproduce the flexible-price allocation, possibly because of a zero lower bound problem.

We aim to provide an alternative view of business cycles that highlights the role of risk premiums—recessions are periods when firms and investors demand large premiums to compensate them for risk. Our model features only risk shocks and no nominal rigidities. The only departure from the standard neoclassical growth model is that entrepreneurs cannot fully insure against the idiosyncratic risk in their firms. This friction enables the model to provide an account of business cycles driven by spikes in risk premiums.

The effects of risk shocks in our model can be understood in terms of risk premiums and precautionary saving. First, because the marginal products of capital and labor are uncertain, business demand for them carries a risk premium to compensate for uninsurable idiosyncratic risk. A spike in risk premiums therefore depresses demand for capital and labor symmetrically.

Second, higher idiosyncratic risk creates a precautionary saving motive that stimulates investment, because capital serves as a store of value that is attractive when risk is high. Investment trades off consumption today for consumption over a long period in the future. In contrast, labor does not serve as a store of value. It trades off disutility for output within a period, and is therefore not directly affected by the precautionary motive. Precautionary




Figure 1: HP-filtered annualized idiosyncratic risk in daily stock market returns, after extracting five principal components, from Herskovic et al. [2016]. Post-war mean is 0.28.

saving therefore affects capital and labor asymmetrically.

We can map the role of incomplete idiosyncratic risk sharing into capital and labor wedges relative to the first-best with perfect risk-sharing, in the tradition of Chari et al. [2007]. Risk shocks create a large increase in the labor wedge and a small reduction in the capital wedge. This configuration of wedges enables the model to deliver realistic business cycles. An increase in the labor wedge—like a labor tax—reduces employment and output, and thereby consumption and investment. It creates a recession. The distribution of the decline in output between lower consumption and lower investment depends on intertemporal substitution, the persistence of the shock, and capital adjustment costs.

An increase in the capital wedge—like a capital tax—reduces investment and employment, but increases consumption. To the extent that risk shocks show up as spikes in the capital wedge, they can't produce business cycles. The observation in our paper is that the precautionary motive compensates the effect of the risk premium on the capital wedge, but not on the labor wedge, as explained above. As a result, risk shocks generate business cycles with realistic co-movement among employment, consumption, and investment.

To evaluate the quantitative plausibility of the mechanism, we calibrate the model to US data. We find that it can successfully reproduce the business cycle behavior of the main macroeconomic aggregates. Idiosyncratic risk rises in recessions, both in terms of establishment-level TFP and sales, or stock returns as shown in Figure 1.¹ The model is stylized in the interest of theoretical clarity, but our claim is that the theoretical mechanism we describe is quantitatively plausible and a promising way to understand business cycles.

We also pursue a sufficient statistic approach that allows us to quantitatively evaluate

¹See Christiano. et al. [2014], Gilchrist et al. [2014], Herskovic et al. [2016], and Bloom et al. [2018].

the relationship between the labor and capital wedges in terms of measurable equilibrium objects, without having to specify many structural details of the environment. When we plug in data from the US, we find that while the labor wedge is positive and rises during recessions, the capital wedge is negative and falls even more during recessions. Besides giving us a quantitative sense of the mechanism, the sufficient statistic also allows us to see how it would change under different scenarios.

The view of business cycles we propose has important implications for the design of policy interventions. In Section 5 we study the efficiency of the competitive equilibrium. We solve for the constrained-efficient allocation that respects the incomplete risk-sharing environment. We find that the competitive equilibrium is inefficient. Under the optimal allocation, risk shocks create a small reduction in the labor wedge and a large increase in the capital wedge—the opposite of the competitive equilibrium. As a result, in the planner’s allocation risk shocks depress investment but produce an expansion in consumption, and output and employment fluctuations are smaller than in the competitive equilibrium. The constrained-efficient allocation can be implemented by stimulating employment and consumption during recessions. A simple way of doing this is subsidizing labor and taxing capital during recessions.

1.1 Relationship to the Literature

Our paper builds on a large literature that highlights the role of time-varying risk premiums in business cycles, such as Cochrane [2011]. A higher risk premium causes a negative demand shock to investment. However, these models are typically unable to deliver the co-movement pattern between consumption, investment, and employment that occurs over the business cycle. This is the essence of the problem identified in Barro and King [1984], and the reason that the wedge accounting analysis in Chari et al. [2007] minimizes the role of investment wedges. Christiano. et al. [2014], Di Tella [2017], Gilchrist et al. [2014], and Bloom et al. [2018] study the impact of aggregate shocks to idiosyncratic risk on investment dynamics. However, without nominal rigidities, risk shocks lead to consumption growth in these models. On their own, risk shocks cannot produce realistic recessions.

There is also a large literature on financial frictions where tighter financial constraints reduce asset prices and investment, including Kiyotaki and Moore [1997], Bernanke et al. [1999], Brunnermeier and Sannikov [2014], and He and Krishnamurthy [2012]. However, in the absence of movements in TFP or nominal rigidities, consumption must rise in recessions in these models. Di Tella [2019] shows how money can produce persistent slumps with excessively depressed investment when idiosyncratic risk premiums rise, but consumption increases on impact. On the other hand, Caballero and Simsek [2018] study the effect of risk shocks in a model with nominal rigidities which create a Keynesian aggregate demand channel that can deliver recessions. Kekre and Lenel [2019] highlight the role of heteroge-

nous risk tolerance in generating time-varying risk premiums, also in a model with nominal rigidities. In contrast, the mechanism in our paper does not require nominal rigidities.

It's worth distinguishing the mechanism we propose from the real options channel highlighted in the literature that follows from Bloom [2009]. The crucial ingredient in this class of models is non-convex adjustment costs, which create inaction regions, but risk sharing is complete. Risk shocks on their own do not produce business cycles, and the competitive equilibrium is typically efficient. In contrast, we assume convex adjustment costs. The crucial ingredient in our model is incomplete idiosyncratic risk sharing, which creates a risk premium on capital and labor and a precautionary motive for saving. Risk shocks produce realistic recessions and the competitive equilibrium is inefficient. We focus on the effect of time-varying idiosyncratic risk for simplicity. A time-varying price of risk would have similar effects. This could be a result of habits as in Campbell and Cochrane [1999] (exploited by Kehoe et al. [2019]), heterogenous risk aversion as in Longstaff and Wang [2012] and Gârleanu and Panageas [2015], or balance sheet effects as in He et al. [2015].

Hall [2017] treats employment as a form of investment, in the context of search models. Higher risk premiums act like a negative demand shock for labor, and has the potential to deliver recessions. His analysis is at the level of the labor market. Kilic and Wachter [2018] apply this idea in a general-equilibrium asset pricing model with rare disasters. Kehoe et al. [2019] incorporate persistent human capital accumulation to overcome the puzzle of low employment volatility uncovered in Shimer [2005]. These models largely abstract from investment, however. This limitation is important because, while a higher risk premium may depress job creation, it also reduces the demand for investment. In the short run, consumption may well rise in recessions in models that include investment. Kilic and Wachter [2018] briefly describe results for an extension of their model to include investment, but in that extension, a financial shock induces a substantial jump in consumption.

Arellano et al. [2019] study the role of uninsurable idiosyncratic risk for demand for labor in an open economy with financial frictions. In their model higher risk increases the probability of default, and, through higher credit spreads, reduces demand for labor. However, they abstract from investment. Jermann and Quadrini [2012] incorporate investment in a model with borrowing constraints where labor must be paid in advance, as working capital. When financial conditions tighten, firms that are up against the constraint must reduce employment. But this mechanism also affects investment, and in general equilibrium consumption spikes on impact. In a similar line, Occhino and Pescatori [2015] look at the effect of a debt-overhang problem, which also show up as both capital and labor wedges. It also produces countercyclical consumption.

In contrast to these papers, our paper focuses on the general equilibrium response to higher risk premiums and the resulting co-movement of consumption, investment, and employment. Higher idiosyncratic risk creates both a risk premium, which depresses demand

for capital and labor symmetrically, and a precautionary saving motive, which depresses interest rates. The difference in duration between capital and labor plays a central role in generating realistic co-movement. Because capital has a long duration, lower interest rates stimulate demand, dampening the effects of the risk premium. The result is a large countercyclical labor wedge, and a small procyclical capital wedge.

We use a simple neoclassical model with a frictionless spot market for labor. Uninsurable idiosyncratic risk is the only friction. While we abstract from search frictions, we believe that incorporating them into the framework is a promising approach for future work. The focus of our paper is on uninsurable idiosyncratic risk on the business side, an issue highlighted by Angeletos [2007] and Meh and Quadrini [2004]. We also abstract from uninsurable labor income risk, which is the topic of a large literature, for example, Aiyagari [1994].

2 Setting

There are two types of agents, workers and entrepreneurs. A representative worker supplies labor, and entrepreneurs use capital $k_{i,t}$ and labor $\ell_{i,t}$ to produce goods. Each entrepreneur is exposed to idiosyncratic risk. The output flow for entrepreneur i is

$$dY_{i,t} = f(k_{i,t}, \ell_{i,t})dt + f(k_{i,t}, \ell_{i,t})\sigma_t dB_{i,t}, \quad (1)$$

where $B_{i,t}$ is a Brownian motion specific to this entrepreneur, and $f(k, \ell) = k^\alpha \ell^{1-\alpha}$ is the usual Cobb-Douglas production function. Idiosyncratic risk washes out in the aggregate, so aggregate output flow is $f(k_t, \ell_t)dt$, as usual. As a result, the aggregate resource constraints are

$$c_t + \phi(x_t)k_t = f(k_t, \ell_t) \quad (2)$$

and

$$dk_t = (x_t k_t - \delta k_t)dt, \quad (3)$$

where $c_t = c_{w,t} + c_{e,t}$ is aggregate consumption, $c_{e,t} = \int c_{i,t} di$ is total consumption by entrepreneurs, $\phi(x_t)k_t$ is investment, δ is the rate of depreciation of capital, and $x_t - \delta$ is the rate of growth of the capital stock. The function $\phi(x)$ incorporates convex adjustment costs. We assume it takes the form $\phi(x) = (\exp(\epsilon(x - \delta)) - 1)/\epsilon + \delta$. In the steady-state $\phi(\delta) = \delta$ and $\phi'(\delta) = 1$, just like with quadratic adjustment costs. With $\epsilon = 0$ we have the case without adjustment costs.

The level of idiosyncratic risk σ_t follows a diffusion

$$d\sigma_t = \theta_\sigma(\bar{\sigma} - \sigma_t)dt + \sqrt{\sigma_t}v_\sigma dZ_t \quad (4)$$

driven by an aggregate Brownian motion Z that captures risk shocks. This is the only source of aggregate risk in this economy and the only exogenous driving force for business cycles.

Entrepreneurs cannot insure against their idiosyncratic risk B_i for incentive reasons, but markets are complete for aggregate risk. This is the only friction in the economy.

The representative worker's problem. The representative worker has impatience ρ_w and Frisch elasticity of labor supply ψ . The worker's problem is to choose values for consumption c_w and labor ℓ to solve

$$\max_{c_w, \ell} \mathbb{E} \left[\int_0^\infty e^{-\rho_w t} \left(\log(c_{w,t}) - \frac{\ell_t^{1+1/\psi}}{1+1/\psi} \right) dt \right] \quad (5)$$

subject to the intertemporal budget constraint

$$\mathbb{E} \left[\int_0^\infty \xi_t c_{w,t} dt \right] \leq n_{w,0} + \mathbb{E} \left[\int_0^\infty \xi_t \ell_t w_t dt \right], \quad (6)$$

where ξ_t is the pricing kernel, $d\xi_t/\xi_t = -r_t dt - \pi_t dZ_t$, r_t is the interest rate, and π_t is the price of aggregate risk, Z .

Entrepreneurs' problem. Entrepreneurs have impatience $\rho_e > \rho_w$. We take entrepreneurs to be more impatient than workers as a device to obtain a stationary wealth distribution. An entrepreneur's problem is to choose consumption, production, and risk sharing $(c_i, k_i, \ell_i, x_i, \sigma_{n,i})$ to solve

$$\max_{c_i, \ell_i, k_i, x_i, \sigma_{n,i}} \mathbb{E} \left[\int_0^\infty e^{-\rho_e t} \log(c_{i,t}) dt \right] \quad (7)$$

subject to the dynamic budget constraint

$$\begin{aligned} dn_{i,t} = & \left(n_{i,t} r_t + \pi_t n_{i,t} \sigma_{n,i,t} - c_{i,t} + f(k_{i,t}, \ell_{i,t}) - \phi(x_{i,t}) k_{i,t} - w_t \ell_{i,t} + q_t k_{i,t} [x_{i,t} - \delta + \mu_{q,t} \right. \\ & \left. - r_t - \pi_t \sigma_{q,t}] \right) dt + \sigma_t f(k_{i,t}, \ell_{i,t}) dB_{i,t} + \sigma_{n,i,t} n_{i,t} dZ_t \end{aligned} \quad (8)$$

and the solvency constraint $n_{i,t} \geq 0$, where q_t is the price of capital and $\mu_{q,t}$ and $\sigma_{q,t}$ its drift and volatility. We treat $\sigma_{n,i}$ as a choice variable because there are complete markets for aggregate risk sharing. Entrepreneurs can use financial instruments to choose any $\sigma_{n,i}$ independently of other choices. But they cannot share their idiosyncratic risk. If they could, they would perfectly insure and eliminate the $\sigma_t f(k_{i,t}, \ell_{i,t}) dB_{i,t}$ term. This is the only friction in this economy.

Competitive equilibrium. Total wealth is $n_{e,t} + n_{w,t} = q_t k_t$, where $n_{e,t} = \int n_{i,t} di$ is the total wealth of entrepreneurs. For a given initial distribution of wealth, a *competitive equilibrium* is a process for prices (r, π, w, q) , aggregate capital k , a plan for the representative worker (c_w, ℓ) , and a plan for each entrepreneur $(c_i, k_t, \ell_i, x_i, \sigma_{n,i})$ such that every agent optimizes taking prices as given; the aggregate resource constraints (2) and (3) hold; and markets clear: $\int \ell_{i,t} di = \ell_t$, $\int k_{i,t} di = k_t$ and $n_{e,t} + n_{w,t} = q_t k_t$.

2.1 Discussion of Assumptions

Technology. A crucial feature of the environment is that the marginal products of capital and labor are locally uncertain. An entrepreneur making decisions about k and ℓ has a probability distribution of how this will affect profits, but doesn't know for sure. The realized marginal product is uncertain when the factor quantity decision is made. The central premise of the paper is that economic activity involves uninsurable idiosyncratic risk of this type.

To clarify the technology, consider an entrepreneur who employs a constant amount of capital and labor for one year. The output produced over this year has a normal distribution:

$$Y_{i,1} = \int_0^1 f(k, \ell) dt + \int_0^1 f(k, \ell) \sigma dB_{i,t} = f(k, \ell)(1 + \sigma B_{i,1}), \quad (9)$$

where $B_{i,1}$ is distributed as standard normal. The continuous-time formulation says that this uncertain output is revealed gradually and the entrepreneur can continuously adjust labor and capital, but the marginal product is locally uncertain. This formulation allows us to map the entrepreneurs' problem into standard portfolio-choice theory.

Idiosyncratic output over a year can be negative, $Y_{i,1} < 0$. Economic activity may actually destroy resources—for example, a factory could burn down. However, because entrepreneurs can continually adjust k and ℓ , their wealth remains always positive, and has a lognormal distribution in equilibrium. Bankruptcy never occurs in equilibrium.

Idiosyncratic shocks are iid, which yields considerable tractability. Introducing persistent shocks would require keeping track of the cross-sectional joint distribution of productivity and wealth. But the crucial property remains that the marginal products are locally uncertain.

Aggregation. The setting with two types of agents is tractable because it won't be necessary to keep track of the distribution of wealth among entrepreneurs. Homothetic preferences, iid idiosyncratic shocks, and a linear budget constraint make the entrepreneur's problem homogeneous in the entrepreneur's wealth, so aggregation is easy. Crucially, entrepreneurs face uninsurable idiosyncratic risk, but don't have non-tradable labor income. The representative worker has non-tradable labor income, but doesn't face idiosyncratic risk. This

formulation has no role for uninsurable labor income risk, which is the focus of an extensive literature. The focus in this paper is on the role of uninsurable idiosyncratic risk on the entrepreneur’s side, and the resulting risk premium on labor and capital.

Entrepreneurs and risk sharing. Entrepreneurs in the model are insiders who retain a significant stake in the firm for incentive reasons.² They represent (1) founding entrepreneurs who often retain, for decades, substantial undiversified equity positions in large, successful startups, (2) members of families operating family-owned firms, (3) top executives whose incomes are linked to the success of their employers, and (4) private-equity fund managers who hold significant stakes in their portfolio companies. In the model workers do not take on idiosyncratic risk because they don’t have the same incentive problems as these insiders. The type of risk sharing that the first-best allocation requires does not involve workers, but rather risk sharing across entrepreneurs.

3 Characterizing the Competitive Equilibrium

The main departure of our model from the neoclassical growth model is time-varying uninsurable idiosyncratic risk. We can understand its effects in terms of a countercyclical risk premium, which depresses demand for capital and labor, and a precautionary motive for idiosyncratic risk which reduces interest rates.

Agents’ problem. The representative worker’s problem is completely standard. The entrepreneur’s problem can be mapped into a consumption-portfolio problem. Homothetic preferences and linear budget constraints imply that policy functions are linear in wealth, which avoids the need to keep track of the distribution of wealth across entrepreneurs. In what follows, hatted variables are normalized by the entrepreneur’s wealth, such as $\hat{c} = c/n$.

An entrepreneur’s value function takes the form $V_t(n) = \frac{1}{\rho} \log n + A_t$, where A_t captures the forward-looking investment opportunities facing entrepreneurs. The HJB equation is

$$\rho \cdot \left(\frac{1}{\rho} \log n + A_t \right) = \max_{\hat{c}, \hat{k}, \hat{\ell}, x, \sigma_n} \log \hat{c} + \log n + \frac{1}{\rho} \left[\mu_{n,t} - \frac{1}{2} \sigma_n^2 - \frac{1}{2} (\sigma_t f(\hat{k}, \hat{\ell}))^2 \right] + \mu_{A,t}, \quad (10)$$

where $\mu_{A,t}$ is the drift of A_t , and the drift of wealth is

$$\mu_{n,t} = r_t + \pi_t \sigma_n + f(\hat{k}, \hat{\ell}) - \phi(x) \hat{k} + (x_{i,t} - \delta + \mu_{q,t} - r_t - \pi_t \sigma_{q,t}) q_t \hat{k} - w_t \hat{\ell} - \hat{c}. \quad (11)$$

Idiosyncratic risk. The entrepreneur’s exposure to idiosyncratic risk plays a central role in our model, $\tilde{\sigma}_{n,e,t} = f(\hat{k}_t, \hat{\ell}_t) \sigma_t = y_t/n_{e,t} \times \sigma_t$. Using the FOC for consumption, $c_{i,t} =$

²See Angeletos [2007], Meh and Quadrini [2004], Moskowitz and Vissing-Jørgensen [2002], and Himmelberg et al. [2004].

$\rho n_{i,t}$, we get that idiosyncratic risk in entrepreneurs' consumption is $\tilde{\sigma}_{c,e,t} = \tilde{\sigma}_{n,e,t}$. Using the resource constraint, we obtain an expression for idiosyncratic risk in entrepreneurs' consumption,

$$\tilde{\sigma}_{c,e,t} = \frac{(k_t/\ell_t)^{\alpha-1}}{(k_t/\ell_t)^{\alpha-1} - \phi(x_t)} \frac{1+z_t}{z_t} \rho_e \sigma_t. \quad (12)$$

where $z_t = c_{e,t}/c_{w,t}$ is the consumption ratio of entrepreneurs relative to workers. We will later use z_t as a state variable in the recursive formulation of the problem.

Risk premium for idiosyncratic risk. Using the FOCs for $\hat{\ell}$ and \hat{k} we obtain Marshallian demand functions for labor and capital,

$$w_t = \overbrace{(1-\alpha)(k_t/\ell_t)^\alpha}^{\text{perfect risk sharing}} \left(1 - \overbrace{\tilde{\sigma}_{c,e,t}\sigma_t}^{\text{risk premium}} \right) \quad (13)$$

and

$$R_t = \alpha(k_t/\ell_t)^{\alpha-1} \left(1 - \tilde{\sigma}_{c,e,t}\sigma_t \right) \quad (14)$$

With perfect idiosyncratic risk sharing, we would get the usual expressions where the wage and the rental price of capital are equal to the marginal products of each factor. With incomplete risk sharing a risk premium emerges to compensate entrepreneurs for the uninsurable idiosyncratic risk they face when using capital and labor. The risk premium captures the covariance of the marginal product of capital or labor with the entrepreneur's marginal utility, $c_{i,t}^{-1}$, and reduces demand for capital and labor symmetrically.

Precautionary saving motive for idiosyncratic risk. Uninsurable idiosyncratic risk also shows up as a precautionary saving motive for entrepreneurs, which depresses equilibrium real interest rates. The entrepreneurs' Euler equation is

$$r_t = \rho_e + \mu_{c,e,t} - \sigma_{c,e,t}^2 - \tilde{\sigma}_{c,e,t}^2, \quad (15)$$

where $\tilde{\sigma}_{c,e,t}^2$ captures the precautionary motive for idiosyncratic risk, and $\sigma_{c,e,t}^2$ for aggregate risk. The workers' Euler equation is

$$r_t = \rho_w + \mu_{c,w,t} - \sigma_{c,w,t}^2. \quad (16)$$

Since they don't on take any idiosyncratic risk, they only have a precautionary motive for aggregate risk $\sigma_{c,w,t}^2$.

Complete aggregate risk sharing means that the consumption of entrepreneurs and workers have the same exposure to aggregate risk. From the FOCs for aggregate risk

sharing for entrepreneurs and workers, we obtain

$$\pi_t = \sigma_{c,e,t} = \sigma_{c,w,t} = \sigma_{c,t}. \quad (17)$$

Their precautionary motive for aggregate risk is therefore the same, $\sigma_{c,e,t}^2 = \sigma_{w,e,t}^2 = \sigma_{c,t}^2$.

Weighting the Euler equations by their consumption shares, $1/(1+z_t)$ and $z_t/(1+z_t)$, we obtain

$$r_t = \overbrace{\bar{\rho}_t + \mu_{c,t} - \sigma_{c,t}^2}^{\text{perfect risk sharing}} - \underbrace{\frac{z_t}{1+z_t}}_{c_{e,t}/c_t} \overbrace{\tilde{\sigma}_{c,e,t}^2}^{\text{lower interest rate prec. mot.}}, \quad (18)$$

where $\bar{\rho}_t = \frac{1}{1+z_t}\rho_w + \frac{z_t}{1+z_t}\rho_e$ is the consumption-weighted impatience rate. The first part of (18) is the expression for the real interest rate in a model with perfect risk sharing. With incomplete risk sharing, entrepreneurs' precautionary motive for idiosyncratic risk depresses the real interest rate, weighted by their consumption share.

The Euler equations and aggregate risk sharing conditions also give us a law of motion for the consumption ratio z_t ,

$$\mu_{z,t} = \rho_w - \rho_e + \tilde{\sigma}_{c,e,t}^2, \quad \sigma_{z,t} = 0 \quad (19)$$

The term $\tilde{\sigma}_{c,e,t}^2$ is entrepreneurs' precautionary saving motive. Here we can see that if entrepreneurs and workers had the same impatience rate, $\rho_e = \rho_w$, entrepreneurs would eventually account for all the consumption in the economy ($z_t \rightarrow \infty$) because of their precautionary saving motive for idiosyncratic risk. We assume $\rho_e > \rho_w$ to obtain a stationary distribution for z .

Labor and capital markets. The two main equilibrium conditions come from the labor and capital markets. The worker's FOC for labor supply is $\ell_t^{1/\psi} = c_{w,t}^{-1}w_t$. Plugging in (13) and using $c_w = c_t/(1+z_t)$, we obtain the equilibrium condition in the labor market:

$$\ell_t^{1/\psi} \times \overbrace{c_t/(1+z_t)}^{c_{w,t}} = \overbrace{(1-\alpha)(k_t/\ell_t)^\alpha(1-\tilde{\sigma}_{c,e,t}\sigma_t)}^{w_t}. \quad (20)$$

The optimality condition $\phi'(x_t) = q_t$ implies

$$x_t = \frac{\log q_t}{\epsilon} + \delta \quad \text{and} \quad \phi(x_t) = \frac{q_t - 1}{\epsilon} + \delta. \quad (21)$$

Next we use the entrepreneur's FOC for k to obtain an asset-pricing equation for q_t ,

$$\overbrace{\alpha(k_t/\ell_t)^{\alpha-1}(1-\tilde{\sigma}_{c,e,t}\sigma_t)}^{R_t} - \phi(x_t) + q_t(x_t - \delta + \mu_{q,t} - r_t - \pi_t\sigma_{q,t}) = 0. \quad (22)$$

The competitive equilibrium is pinned down by the equilibrium conditions for employment and investment, (20) and (21); prices (13), (17), (18), and (22); the resource constraint (2); entrepreneurs' idiosyncratic risk (12); and the laws of motion of k_t , σ_t , and z_t , given by (3), (4), and (19).

3.1 Recursive Equilibrium

We construct a recursive equilibrium with (k, σ, z) as state variables. Working with z_t as a state variable is useful because its law of motion is particularly simple. The two basic endogenous variables will be employment $\ell(k, \sigma, z)$ and the price of capital $q(k, \sigma, z)$ and will be characterized with a system of differential equations. All other endogenous variables can be derived from them and the state variables.

Equations (20) and (22) are our two main equilibrium conditions, and will define a system of differential equations for $q(k, \sigma, z)$ and $\ell(k, \sigma, z)$. From (12) and (21) we know how to express $\tilde{\sigma}_{c,e}$ and x in terms of q and ℓ , in addition to the states (k, σ, z) . Equations (17) and (18) give us expressions for π and r ; we only need to find expressions for the drift and volatility of aggregate consumption, μ_c and σ_c . Use the resource constraint to replace $c = k[(k/l)^{\alpha-1} - \phi(x)]$ in (20) and to obtain an expression for μ_c and σ_c in terms of the drift and volatility of q and ℓ :

$$\mu_c = (x - \delta) + \frac{(1 - \alpha)(k/\ell)^{\alpha-1} (\mu_\ell - (x - \delta) - \frac{\alpha}{2}\sigma_\ell^2) - (q/\epsilon)\mu_q}{(k/\ell)^{\alpha-1} - \phi(x)} \quad (23)$$

and

$$\sigma_c = \frac{(1 - \alpha)(k/\ell)^{\alpha-1}\sigma_\ell - (q/\epsilon)\sigma_q}{(k/\ell)^{\alpha-1} - \phi(x)}. \quad (24)$$

We already know the laws of motion of k , σ , and z , so we use Ito's lemma to compute the drift and volatility of q and ℓ :

$$\mu_q = \frac{q'_k(x - \delta)k + q'_\sigma\theta_\sigma(\bar{\sigma} - \sigma) + q'_z\mu_z z + \frac{1}{2}q''_{\sigma,\sigma}(\sqrt{\sigma}v_\sigma)^2}{q}, \quad \sigma_q = \frac{q'_\sigma\sqrt{\sigma}v_\sigma}{q}, \quad (25)$$

$$\mu_\ell = \frac{\ell'_k(x - \delta)k + \ell'_\sigma\theta_\sigma(\bar{\sigma} - \sigma) + \ell'_z\mu_z z + \frac{1}{2}\ell''_{\sigma,\sigma}(\sqrt{\sigma}v_\sigma)^2}{\ell}, \quad \sigma_\ell = \frac{\ell'_\sigma\sqrt{\sigma}v_\sigma}{\ell}. \quad (26)$$

Putting this together, equation (22) gives us a second order partial differential equation and equation (20) an algebraic equation, which pin down $q(k, \sigma, z)$ and $\ell(k, \sigma, z)$. We solve this system numerically using Smolyak interpolation, as described in Judd et al. [2014].

<i>Meaning</i>	<i>Parameter</i>	<i>Value</i>
Capital share	α	1/3
Frisch elasticity of labor supply	ψ	3
Capital adjustment cost	ϵ	5
Depreciation	δ	0.07
Impatience rate, workers	ρ_w	0.035
Impatience rate, entrepreneurs	ρ_e	0.0975
long run idiosyncratic risk	$\bar{\sigma}$	0.10
Mean-reversion of idiosyncratic risk	θ_σ	0.693
Aggregate volatility of idiosyncratic risk	v_σ	0.16

Table 1: Parameter Values

4 Main Results

4.1 Parameter Values

We solve the model with the parameter values shown in Table 1. We use an informal combination of calibration to existing research and rough matching of observed moments to arrive at these values. Our objective is to evaluate the quantitative plausibility of the theoretical mechanism we propose. In Section 4.6 we perform a sensitivity analysis to understand how results change under different parameter specifications.

The Cobb-Douglas capital elasticity, $\alpha = 1/3$, is standard. The Frisch elasticity of labor supply, $\psi = 3$, is standard in the macro literature when employment volatility is an issue. Hall [2009] suggests that an elasticity around that value is a reasonable working approximation for employment fluctuations in an economy with a search and matching setup and realistic equilibrium wage stickiness, in the sense of Hall [2005]. For the elasticity of the adjustment costs, which helps determine how lower output is split between lower investment and consumption, we use $\epsilon = 5$ to roughly match the observed split.

We set workers' impatience $\rho_w = 0.035$ and the depreciation rate $\delta = 0.07$ to match the steady-state ratios $k_{ss}/y_{ss} = 3$ and $c_{ss}/y_{ss} = 0.8$, which play an important role in the model, as we describe below in section 4.4. We set $\rho_e = \rho_w + 0.0625$ to obtain a steady-state idiosyncratic volatility of wealth $\tilde{\sigma}_{n,e,t} = \sqrt{\rho_e - \rho_w} = 25$ percent, roughly corresponding to the idiosyncratic risk in the stock market as in Herskovic et al. [2016].

For the stochastic process for σ_t we use $\bar{\sigma} = 0.1$, in line with evidence in Bloom et al. [2018] on idiosyncratic productivity risk at the establishment level; $\theta_\sigma = 0.692$ so risk shocks have a half-life of one year; and $v_\sigma = \frac{1}{2}\sqrt{\bar{\sigma}} = 0.158$ so that a two-standard deviation shock doubles idiosyncratic risk. We aim to capture transitory fluctuations in idiosyncratic risk at

<i>Variable</i>		<i>St. dev.</i>		<i>Rel. st. dev.</i>		<i>Corr. w/output</i>		<i>Autocorr.</i>	
		<i>Model</i>	<i>Data</i>	<i>Model</i>	<i>Data</i>	<i>Model</i>	<i>Data</i>	<i>Model</i>	<i>Data</i>
Output	y	1.7%	1.6%	1	1	1	1	0.91	0.85
Consumption	c	0.86%	0.80%	0.49	0.50	0.84	0.80	0.89	0.84
Investment	$\phi(x)k$	5.8%	6.34%	3.36	3.96	0.95	0.85	0.94	0.82
Employment	ℓ	2.1%	1.89%	1.24	1.18	0.87	0.87	0.87	0.91

Table 2: Business Cycle Moments from the Model Compared to HP-Filtered Data (1600), at Quarterly Frequency, Per-Capita, and in Logs, for the full postwar period.

the business cycle frequency. These numbers are broadly in line with evidence in Herskovic et al. [2016] for idiosyncratic risk in the stock market. The implied steady-state ratio of entrepreneurial consumption to worker consumption is $z_{ss} = 5$ percent.

4.2 Impulse-Response Functions and Business-Cycle Moments

Figure 2 shows the impulse response to a one-standard deviation risk shock. The economy starts in the steady state, the long run if aggregate shocks realize to zero, and we follow the response of the economy in the absence of further shocks.

The first panel shows the behavior of idiosyncratic risk σ_t , which spikes by 5.5 percentage points on impact and then returns to its long run value of 10 percent, and the behavior of the idiosyncratic risk of entrepreneurs’ wealth or consumption, $\tilde{\sigma}_{c,e,t}$, which spikes by 10 percentage points on impact and then returns to its long run value of 25 percent. The idiosyncratic risk premium, $\tilde{\sigma}_{c,e,t}\sigma_t$, therefore displays the same behavior. It spikes on impact by 3 percentage points, and then returns to its long run value of 2.5 percent.

Risk shocks create recessions in the model. The second panel shows the responses of consumption, investment, and employment. They all fall on impact and slowly recover afterwards. Investment falls by 5.5 percent, consumption by 1 percent and employment by 3 percent. This is broadly in line with stylized facts about US business cycles.

The third panel shows the behavior of real interest rates and wages. Wages fall on impact by around 2 percent, reflecting weaker labor demand—entrepreneurs demand a larger risk premium to compensate for uninsurable idiosyncratic risk. Interest rates are roughly acyclical, reflecting the interaction of an expected recovery of consumption, which raises interest rates, and a larger precautionary motive, which lowers them. The fourth panel shows the capital and labor wedges generated by uninsurable idiosyncratic risk, which we will discuss in the next section in detail.

Figure 3 shows output, consumption, investment, and hours of work, in the model and the data. For the data we use output, consumption of nondurables and services, investment (including durables), and hours in the non-farm business sector, all quarterly, in logs, and

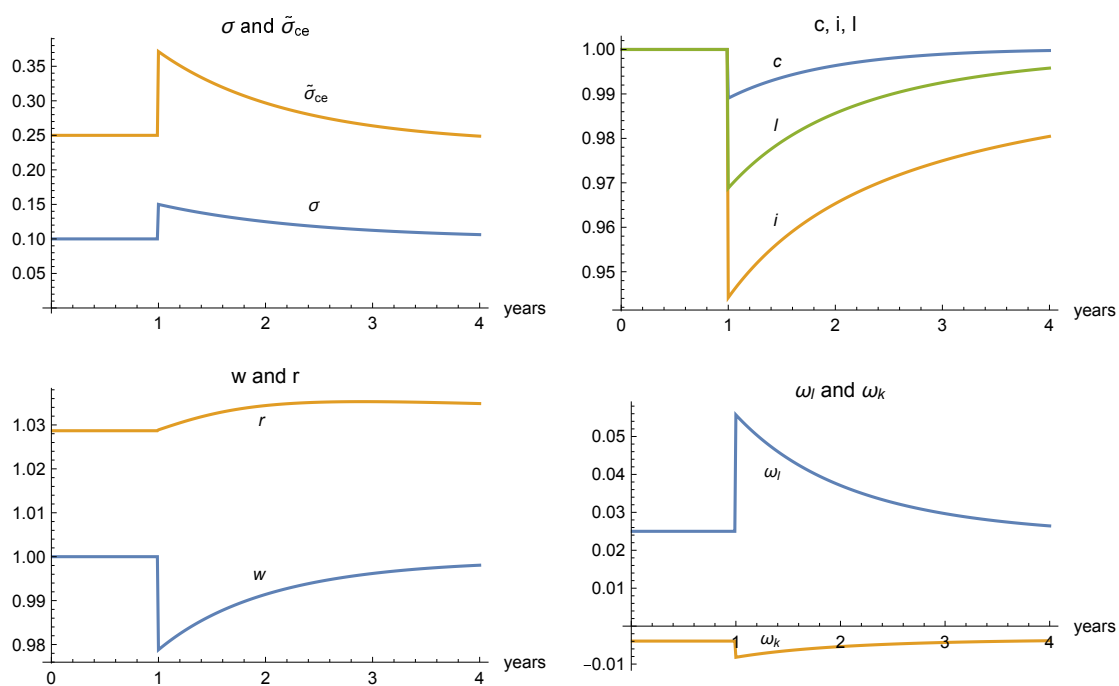


Figure 2: Impulse Response to a Risk Shock

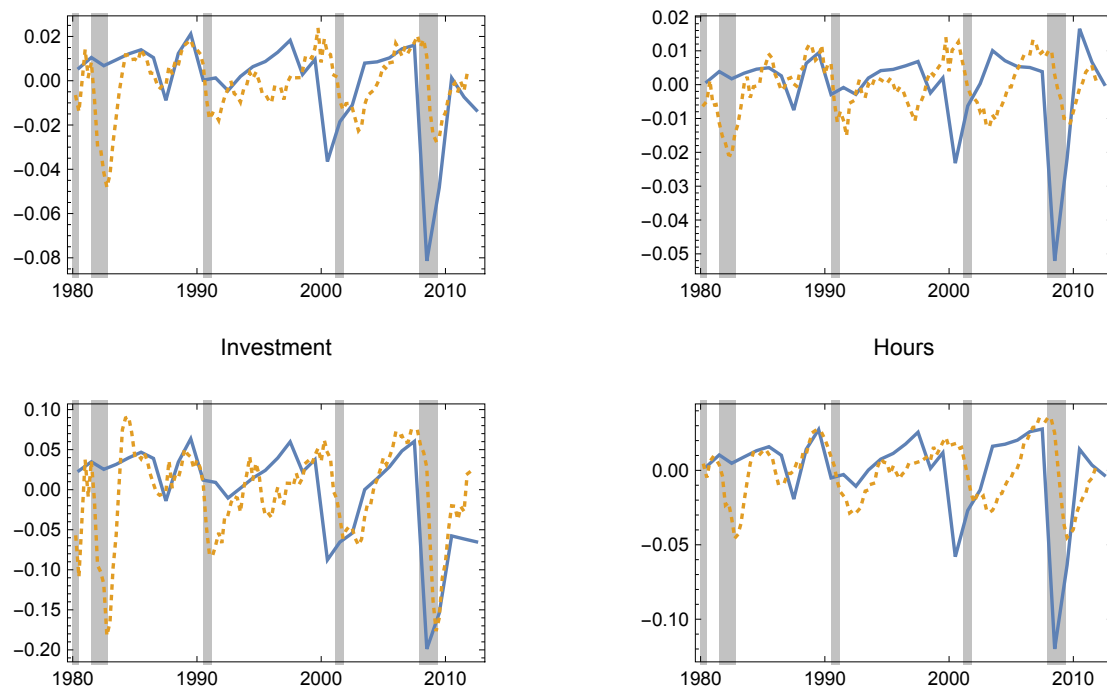


Figure 3: Output, Consumption, Investment, and Hours in the Model (solid) compared to HP-Filtered (1600) Data (dotted) Per-Capita and in Logs, from 1980q1-2012q4.

per capita terms, HP-filtered (1600). We start the model at the steady state in 1980, and back out a yearly series for idiosyncratic risk σ_t to match the time series for idiosyncratic risk in stock returns in Figure 1 (adding a 0.25 mean), which corresponds to $\tilde{\sigma}_{c,e,t} = \tilde{\sigma}_{n,e,t}$ in the model. The model does a good job explaining the fluctuations in the data, with the salient exception of the 1981-82 recession. That recession was likely related to a monetary policy shock associated with the Volcker disinflation. The following three recessions are captured fairly well by the model, with the caveat that the model slightly leads the data. Idiosyncratic risk in stock returns spikes early in the recession, and it seems to take a short time for real variables to respond.

Table 2 summarizes key moments in the model and compares them to US data. It shows the standard deviation and correlation structure of cyclical output, consumption, investment, and hours. The model can successfully replicate business cycle behavior in the data. In Section 4.6 we explore the sensitivity of these results to alternative parametrizations. Our model is stylized in the interest of tractability and theoretical clarity. But our takeaway is that the mechanism we propose is quantitatively plausible and seems like a promising approach to understanding business cycles.

4.3 Wedges

During downturns idiosyncratic risk rises, making both capital and labor less attractive. The result is lower employment, consumption, and investment. The mechanism involves subtle general equilibrium effects. These can be understood from the wedges created by time-varying uninsurable idiosyncratic risk, in the spirit of Chari et al. [2007], where we take as the benchmark the efficient allocation with perfect risk sharing.

Risk shocks create time-varying labor and capital wedges. While higher risk reduces demand for labor and capital symmetrically, in general equilibrium lower interest rates affect labor and capital asymmetrically. The source of asymmetry is that capital provides a vehicle for precautionary saving, while labor does not. There is in addition a wedge in the law of motion of the consumption ratio z_t . While this plays an important role in the long run, it's a slow-moving state variable that does not respond on impact to risk shocks, and therefore does not play an important role in business cycles.

We start with the labor wedge $\omega_{\ell,t}$. From the market clearing condition for labor, equation (20), we write

$$\ell_t^{1/\psi} \frac{c_t}{1+z_t} = (1-\alpha)(k_t/\ell_t)^\alpha (1-\omega_{\ell,t}), \quad (27)$$

where the labor wedge $\omega_{\ell,t}$ is equal to the risk premium for idiosyncratic risk,

$$\omega_{\ell,t} = \tilde{\sigma}_{c,e,t} \sigma_t. \quad (28)$$

Higher risk σ_t raises the risk premium on labor, making employment less attractive.

With $\omega_{\ell,t} = 0$, equation (27) is the equilibrium condition for employment in an economy with perfect risk sharing. If $\omega_{\ell,t} > 0$ the economy behaves as if labor income is taxed. Increases in this type of wedge reduces equilibrium employment and output, as well as consumption and investment. The extent of consumption-smoothing and the magnitude of capital adjustment costs determine the distribution between lower investment and lower consumption.

It's important to highlight that the $1 + z_t$ term on the left side of equation (27) reflects only agent heterogeneity. We would have this even with perfect risk sharing, because only some agents actually provide labor. When z_t is large, workers' share of aggregate consumption is small so the income effect induces them to work more. Crucially, however, z_t does not react to the risk shocks, $\sigma_{z,t} = 0$. The direct impact of a risk shock on the equilibrium condition for labor is purely through the labor wedge $\omega_{\ell,t}$.

Second, there is a capital wedge, as if capital income was taxed or subsidized. Combining both agents' Euler equations and the asset pricing equation for capital, (22), we obtain

$$\alpha(k_t/\ell_t)^{\alpha-1}(1 - \omega_{k,t}) - \phi(x_t) + q_t(x_t - \delta + \mu_{q,t} - (\bar{\rho}_t + \mu_{c,t} - \sigma_{c,t}^2) - \sigma_{c,t}\sigma_{q,t}) = 0, \quad (29)$$

where $\omega_{k,t}$ is the capital wedge. With $\omega_{k,t} = 0$, equation (29) yields the equilibrium condition for investment in an economy with perfect risk sharing. If $\omega_{k,t} > 0$ it's as if capital income is taxed, reducing incentives for investment. In contrast to the labor wedge, this wedge does not create a recession. By itself, a larger wedge $\omega_{k,t}$ reduces investment but increases consumption.

The capital wedge is

$$\omega_{k,t} = \overbrace{(\tilde{\sigma}_{c,e,t}\sigma)}^{\omega_{\ell,t}} \times \left(1 - \frac{1}{\alpha} \frac{\rho_e k_t q_t}{c_t}\right). \quad (30)$$

It has the same form as the labor wedge, but the risk premium is dampened by the $1 - \frac{1}{\alpha} \frac{\rho_e k_t q_t}{c_t} < 1$ factor.

Two forces operate on the capital wedge. Higher risk σ_t raises the risk premium on capital, making investment less attractive, as shown in (14). However, it also creates a precautionary saving motive, which depresses the interest rate conditional on the behavior of consumption, as shown in (18), and makes investment more attractive. Which force dominates depends on the sign of $1 - \frac{1}{\alpha} \frac{\rho_e k_t q_t}{c_t}$. In general this term could be positive or negative.

To explain this, it's useful to go over a well-known benchmark. In a simple AK economy with $\alpha = 1$, and $\rho_e = \rho_w$, we know that $c_t = \rho_e k_t$, and so $1 - \frac{1}{\alpha} \frac{\rho_e k_t q_t}{c_t} = 0$. In this case the intertemporal wedge $\omega_{k,t} = 0$ for any k_t , σ_t , and z_t . The risk premium and the precautionary motive always exactly cancel out. This is in fact a well-known property of log preferences.

If, still in the AK environment, the intertemporal elasticity is greater than 1 (relative

risk aversion below 1), then the risk premium dominates. Higher risk depresses investment and raises consumption on impact. If instead the intertemporal elasticity is below 1 (relative risk aversion greater than 1), then the precautionary saving motive dominates. Higher risk raises investment and depresses consumption.

As we move away from the simple AK environment, we lose the simple characterization of the sign of the capital wedge. Angeletos [2007] studies this issue in an economy with uninsurable idiosyncratic risk on capital, but safe labor, and finds that the risk premium dominates if the intertemporal elasticity is greater than α . In our environment intertemporal elasticity is 1 (log preferences) but labor is also exposed to uninsurable idiosyncratic risk, and we don't have an analytical characterization. We therefore turn to numerical solutions, and below in section 4.4 we provide a sufficient statistic in terms of measurable equilibrium objects.

The fourth panel of Figure 2 shows the labor and capital wedge as an impulse response. The main takeaway is that risk shocks cause a large increase in the labor wedge, and a small contraction in the capital wedge. This constellation of wedges is what is needed to generate business cycles where consumption, investment, and employment co-move. The average labor wedge is 2.7 percent and has a standard deviation of 2.2 percent, while the average capital wedge is -0.35 percent and has a standard deviation of 0.25 percent. Their correlation is -0.9 .

Finally, there is also a wedge in the law of motion of the consumption ratio z_t . With perfect risk sharing $\mu_{z,t} = \rho_w - \rho_e$ and in the long run entrepreneurs' consumption share vanishes, $z_t/(1+z_t) \rightarrow 0$. With incomplete risk sharing entrepreneurs have a precautionary motive for saving, so $\mu_{z,t} = \rho_w - \rho_e + \tilde{\sigma}_{c,e,t}^2$. As a result, their consumption share $z_t/(1+z_t)$ has a non-degenerate ergodic distribution. However, because of perfect aggregate risk sharing z_t does not respond on impact to risk shocks, $\sigma_{z,t} = 0$. It's a slow-moving variable, analogous to the capital stock k_t . As a result, while z_t plays an important role in the long run, its role in business cycles is secondary.

Asymmetry between capital and labor. The first thing to note is that the risk-adjusted marginal product of capital and labor are affected symmetrically by the risk premium. The asymmetry comes from the long duration of capital, which makes investment a way to save as a precaution against future adverse shocks, and the short duration of the employment relationship (zero because of frictionless spot labor markets), which means that labor does not serve as a store of value. While the risk premium depresses investment, the precautionary motive drives the risk-free interest rate down. Because capital has a long duration, lower interest rates stimulate investment. The capital wedge is dampened relative to the labor wedge, and in the quantitative calibration is negative and stimulates investment in recessions.

To show this, we write the capital wedge as:

$$\omega_{k,t} = \underbrace{\tilde{\sigma}_{c,e,t}\sigma}_{\text{risk prem.}} - \frac{q_t}{\underbrace{\alpha(k_t/\ell_t)^{\alpha-1}}_{\text{price-dividend ratio}}} \overbrace{\frac{z_t}{1+z_t} \tilde{\sigma}_{c,e,t}^2}^{\text{lower interest rate}} \underbrace{\phantom{\tilde{\sigma}_{c,e,t}^2}}_{\text{prec. motive}}. \quad (31)$$

The first term captures the role of the risk premium on capital, which depresses demand for capital. This role is symmetric to labor. The second term captures the compensating role of lower interest rates because of entrepreneurs' precautionary saving motive for idiosyncratic risk, weighted by entrepreneurs' consumption share $z_t/(1+z_t) = c_{e,t}/c_t$. The price-dividend ratio of capital plays a central role. If capital had a zero duration, like labor, its price dividend ratio would be zero and we would recover $\omega_{k,t} = \tilde{\sigma}_{c,e,t}\sigma_t$, as for labor.

We conjecture that if employment relationships had a short but strictly positive duration, such as in a search model, then the precautionary saving motive would also dampen the labor wedge, but to a lesser extent than the capital wedge because of the shorter duration of employment (reflected in a lower price-dividend ratio). The proper duration in these models is not the expected length of an employment relationship, but the duration of the dividends from a marginal increase in recruiting activity—for example, creating a match 2 or 3 months earlier than otherwise.

There is also a technological asymmetry between labor and capital, in that capital supply is inelastic in the short-run, while labor has an elastic labor supply. This is why employment can fall on impact, while capital use is fixed. We conjecture that with variable capital utilization, capital use will also fall on impact. But this asymmetry is relevant to the different response of capital and labor to given wedges. It does not explain why the wedges themselves are different.

Summary. Risk shocks create time-varying labor and capital wedges. The labor wedge depresses employment, output, consumption, and investment. The capital wedge, instead, twists the economy by depressing investment and raising consumption. While higher risk reduces demand for labor and capital symmetrically, the lower interest rates produced by the precautionary saving motive affects them asymmetrically because of their different duration. Because capital has a long duration, lower interest rates dampen the capital wedge and may even stimulate investment. In contrast, employment has a short duration, so lower interest rates don't affect it. The result is a large increase in the labor wedge and a small reduction in the capital wedge after risk shocks, precisely what is needed to produce business cycles.

4.4 A Sufficient Statistic for the Sign of the Capital Wedge

The capital wedge arises out of two conflicting forces. The risk premium reduces demand for capital, while the precautionary motives stimulates it. While we don't have an analytical characterization of the sign of the capital wedge, we can express it in terms of meaningful measurable equilibrium objects that help us get a sense of its value. The sign of the capital wedge depends on whether the following expression is above or below 1:

$$\frac{1}{\alpha} \frac{\rho_e k_t q_t}{c_t} = \rho_e \times \frac{y_t}{c_t} \times \frac{1}{\alpha} \times \frac{k_t q_t}{y_t}. \quad (32)$$

We know that $\alpha = 1/3$, and that in steady-state the capital income ratio, $k_{ss}q_{ss}/y_{ss} = 3$ and the consumption-income ratio $c_{ss}/y_{ss} = 0.8$, approximately. We can understand this formula in light of equation (31). Notice that $\frac{1}{\alpha} \times \frac{k_t q_t}{y_t} = \frac{q_t}{\alpha y_t / k_t}$ is the price-dividend ratio of capital at the aggregate level (around 9 in steady-state), and $\rho_e y_t / c_t = \frac{\tilde{\sigma}_{c,e,t}^2}{\tilde{\sigma}_{c,e,t}\sigma} \times \frac{z_t}{1+z_t}$ is the ratio of precautionary motive to the risk premium, weighted by entrepreneurs' consumption share $z_t/(1+z_t) = c_{e,t}/c_t$.

The only remaining parameter is ρ_e . The model pins down its value in terms of workers' impatience rate $\rho_w = 0.035$ and the steady-state idiosyncratic risk of entrepreneurs' wealth (or consumption), $\tilde{\sigma}_{n,e,t} = 25$ percent, which is a core object of interest in this model:

$$\rho_e = \rho_w + \tilde{\sigma}_{n,e,t}^2 = 0.0975 \quad (33)$$

This expression is easy to recover from the law of motion of the consumption ratio, $\mu_z = \rho_w - \rho_e + \tilde{\sigma}_{c,e,t}^2 = 0$. Essentially, entrepreneurs' consumption will grow faster than workers' because of their idiosyncratic precautionary motive, until their exposure to risk $\tilde{\sigma}_{c,e,t} = \tilde{\sigma}_{n,e,t}$ falls to the point where it just compensates their higher impatience. In other words, expression (33) tells us how much more impatient entrepreneurs must be if we observe them in a steady-state, given their exposure to idiosyncratic risk. The value of ρ_e we recover is high but not unreasonable.

These numbers yield a steady-state value of $\frac{1}{\alpha} \frac{\rho_e k_{ss} q_{ss}}{c_{ss}} = 1.097$. The precautionary motive is slightly larger than the risk premium. In response to shocks, the capital-income ratio $k_t q_t / y_t$ and the consumption-income rate c_t / y_t will move a little, but not much. So we get a small reduction in the capital wedge after risk shocks, as shown in Figure 2.

While we have a good idea of how to discipline ρ_w , $k_{ss}q_{ss}/y_{ss}$, and c_{ss}/y_{ss} , the value of ρ_e is more uncertain because it's pinned down by the model in a roundabout way. However, this gives us a model-consistent benchmark we can use to understand the mechanism and think of comparative statics. For example, if we reduce ρ_e while still maintaining the k_{ss}/y_{ss} and c_{ss}/y_{ss} ratios, the precautionary motive will become weaker relative to the risk premium and the capital wedge will become positive. It will always be dampened (more

negative) compared to the labor wedge, $\omega_{\ell,t} > \omega_{k,t}$.

4.5 Excess Return of Capital, Markups, and Labor Share

Here we explore implications of our model for some salient equilibrium objects that have received significant attention in the macro literature.

Excess return of capital. Our model creates a time-varying excess return of capital. Part of it corresponds to a risk premium for aggregate risk, which can be identified with the equity premium. But the incomplete idiosyncratic risk sharing produces a further excess return, above the risk premium for aggregate risk. We can re-write the equilibrium condition for capital, (22), using the expected or average marginal product of capital,

$$f'_k(k_t, \ell_t) - \phi(x_t) + (x_t - \delta + \mu_{q,t} - r_t)q_t = q_t \times \left(\underbrace{\pi_t \sigma_{q,t}}_{\text{equity premium}} + \underbrace{\frac{f'_k(k_t, \ell_t)}{q_t} \times \tilde{\sigma}_{c,e,t} \sigma_t}_{\text{excess return}} \right). \quad (34)$$

The difference with the capital wedge in equation (30) is that here we are using the equilibrium interest rate r_t , while the wedge $\omega_{k,t}$ is defined using the interest rate in the model with perfect risk sharing. That is, the wedge helps us understand the effect of incomplete idiosyncratic risk sharing in terms of capital and labor taxes in a model with perfect risk sharing, while the excess return in (34) highlights the failure of the perfect-risk-sharing asset-pricing equation at equilibrium prices, ignoring that the equilibrium interest rate r_t is lower than what it would be with perfect risk sharing given the same allocation.

The advantage of equation (34) is that it's more directly related to the data. Farhi and Gourio [2018] point out that because the return to capital has remained roughly constant over the past decades, while interest rates have gone down, the excess return on capital has become larger. A rising risk premium is one possible explanation, together with rising market power and intangibles, but measures of the equity premium suggest it's stable or even decreasing. Although our model is not designed to address secular trends, the presence of an excess return above the equity premium is consistent with our mechanism.

Quantitatively, however, the total excess return attributable to idiosyncratic risk is small, 0.3 percent on average, with a standard deviation of 0.23 percent and a correlation with output of -0.84 . The equity premium produced by the model is also small, around 0.017 percent, with a standard deviation of 0.02 percent. Our model is essentially a consumption CAPM with log preferences. It is well known that such models cannot explain the equity premium quantitatively.

Likewise, our model produces an almost acyclical interest rate r_t . This is the result of two conflicting forces that play out over time. A larger precautionary saving motive lowers r_t after risk shocks, but a transitory contraction in consumption means agents expect higher

consumption in the future, which raises r_t .

The asset pricing literature has explored several avenues to explain the equity premium puzzle and the behavior of risk-free rates, such as habits as in Campbell and Cochrane [1999], heterogenous risk aversion as in Longstaff and Wang [2012], Gârleanu and Panageas [2015], or Kekre and Lenel [2019], or balance sheet effects as in He et al. [2015]. Our model does not have any of these ingredients, but we believe that exploring the interaction of asset pricing and business cycles is a fruitful area for future research.

Markups and labor share. Following Rotemberg and Woodford [1999], a common approach to business cycles is to focus on the cyclical properties of markups. In our model markups rise in recessions. These markups are not signs of market power, but rather are competitive rents arising from the limited risk bearing capacity of entrepreneurs. The marginal cost of goods is $w_t/f'_l(k_t, \ell_t) = (1 - \tilde{\sigma}_{c,e,t}\sigma_t)$, so we get markups over marginal cost

$$\mu_t = \frac{1}{1 - \tilde{\sigma}_{c,e,t}\sigma_t} - 1 \approx \tilde{\sigma}_{c,e,t}\sigma_t, \quad (35)$$

equal to the risk premium for idiosyncratic risk. Alternatively, we can compute the procyclical labor share of income,

$$\eta_\ell = \frac{w_t \ell_t}{f'_\ell(k_t, \ell_t)} = (1 - \alpha)(1 - \tilde{\sigma}_{c,e,t}\sigma_t). \quad (36)$$

It's worth stressing that the mechanism in our paper does not reduce to a time-varying markup because the precautionary saving motive for idiosyncratic risk also depresses the interest rate relative to the model with perfect risk sharing and a time-varying markup. Taking prices as given, the average user cost of capital is $q_t \times (r_t + \pi_t \sigma_{q,t} + \delta + \phi(x_t) - x_t - \mu_{q,t}) = f'_k(k_t, \ell_t) \times (1 - \tilde{\sigma}_{c,e,t}\sigma_t)$, which implies a markup of $f'_k(k_t, \ell_t)$ divided by the user cost of capital, equal to $1/(1 - \tilde{\sigma}_{c,e,t}\sigma_t) \approx \tilde{\sigma}_{c,e,t}\sigma_t$, the same as using the labor margin. The capital share of income is also procyclical.³ But we have to remember that the equilibrium interest rate is lower than what it would be with perfect risk sharing and a variable markup because of the precautionary saving motive for idiosyncratic risk. As a result, instead of a common capital and labor wedge (as we would get from adding markups to a perfect risk sharing model), our model delivers a large increase in the labor wedge and a small reduction in the capital wedge after risk shocks.

The average markup in the model is 2.7 percent, with a standard deviation of 2.2 percent and a correlation with output of -0.82 . The average labor share is roughly 64 percent, with

³When computing markups and capital and labor shares of income, we are counting the profits obtained by entrepreneurs, $f(k_t, \ell_t)\tilde{\sigma}_{c,e,t}\sigma_t$, as neither labor nor capital income. The pure profit share, $\tilde{\sigma}_{c,e,t}\sigma_t$, is countercyclical. In the data, depending on exactly how these profits are distributed, a fraction might show up as labor income (or example, bonuses) or capital income (for example, accounting profits for private firms).

a standard deviation of 1.4 percent, and a correlation with output of 0.82. In the data, an average markup of 15 percent is common in the literature.⁴ While the presence of markups is consistent with our model, quantitatively our model cannot explain such large markups. Common ingredients such as imperfect competition and distortionary taxes are required to account for markups in the data.

4.6 Sensitivity Analysis

To understand how the model works, we consider perturbations to the benchmark calibration. We focus on the standard deviations of output, consumption, investment, and employment. Because the objective is to understand how the model works, we don't recalibrate the model when changing each parameter. Table 3 summarizes the results.

Frisch elasticity ψ . A successful model of business cycles requires a relatively large elasticity of labor supply. The value $\psi = 3$ in the baseline calibration is reasonable in the macro literature. With a smaller Frisch elasticity around $\psi = 1$, more consistent with micro evidence, the effects of risk shocks are dampened, but would still look like business cycles. The second column of table 3 shows the standard deviations of all variables would be smaller.

In the limiting case with $\psi = 0$ employment and output would be fixed and could not respond to risk shocks. We would still get a countercyclical capital wedge, so risk shocks would create a small spike in investment and a small contraction in consumption, with no movement in employment or output on impact.

Persistence θ_σ and capital adjustment costs ϵ . The mean reversion parameter θ_σ and the curvature of the adjustment cost function ϵ are important in determining how lower output is split into lower consumption and lower investment in response to a larger labor wedge. The elasticity of intertemporal substitution would also play a role here, but it's pinned at one with log preferences.

We calibrate a transitory risk shock with a half-life of one year. Agents are averse to fluctuations in consumption in response to such transitory shocks, so investment must take the brunt of the adjustment. A convex capital adjustment cost function reduces fluctuations in investment and increases fluctuations in consumption.

We calibrate the model with $\epsilon = 5$ to obtain a ratio of standard deviation of consumption to investment roughly in line with the data. To put this number in context, $\epsilon = 5$ means that a contraction in investment of 10 percent is accompanied by a reduction in the marginal

⁴See Edmond et al. [2018], Hall [2019]. Recent work by De Loecker and Eeckhout [2017] finds an average markup of 60 percent. However, this is a sales-weighted markup. Edmond et al. [2018] report a cost-weighted markup using the same data of 25 percent.

	Benchmark⇒	$\psi = 3$	$\epsilon = 5$	$\theta_v = 0.693$	$v_\sigma = 0.16$	$\bar{\sigma} = 0.10$	$\rho_e = .0975$
	Variant⇒	$\psi = 1$	$\epsilon = 3$	$\theta_v = 0.462$	$v_\sigma = 0.10$	$\bar{\sigma} = 0.05$	$\rho_e = 0.0575$
y	1.7%	1%	1.8%	1.9%	1.3%	0.75%	1.2%
c	0.86%	0.5%	0.7%	0.94%	0.59%	0.55%	0.98%
i	5.8%	3.8%	6.6%	8.3%	4.6%	1.6%	2.4%
ℓ	2.1%	1.2%	2.3%	2.5%	1.5%	0.8%	1.2%

Table 3: Log Standard Deviations of Key Variables at Quarterly Frequency in the Benchmark Case and Variants

cost of capital of around 3 percent. Without adjustment costs, the effect on the marginal cost of capital would be zero.

The second and third columns of table 3 show standard deviations for smaller adjustment costs $\epsilon = 3$ and for more persistent risk shocks $\theta_v = 0.462$, which implies a half-life of two years. With smaller adjustment costs the standard deviation of consumption becomes smaller, and that of investment larger. With more persistent shocks all standard deviations become larger.

Volatility of risk shocks σ_v . Risk shocks are the only exogenous driving force in the model. With a smaller volatility of idiosyncratic risk, the standard deviation of all variables would become smaller. The fourth column of table 3 shows standard deviations for a smaller $\sigma_v = 0.1$.

Steady-state level of idiosyncratic risk $\bar{\sigma}$. The effects of risk shocks become larger with the long run level of idiosyncratic risk. The effect of risk enters the model through second moments, so if idiosyncratic risk is on average very small, an increase has only small effects. We calibrate the long run level of idiosyncratic risk $\bar{\sigma} = 10$ percent in line with Bloom et al. [2018]. A lower long run level of idiosyncratic risk $\bar{\sigma} = 5$ percent reduces the effects of risk shocks, and therefore reduces the standard deviation of all variables, as shown in the fifth column of table 3.

Steady-state idiosyncratic risk in entrepreneurs' consumption $\tilde{\sigma}_{c,e,ss}$. The effects of risk shocks also depend on the level of idiosyncratic risk in entrepreneurs' consumption or wealth. The sixth column of table 3 shows standard deviations for a smaller target for this idiosyncratic risk, $\tilde{\sigma}_{c,e,ss} = 15\%$. Hitting this target requires a lower impatience rate of entrepreneurs, $\rho_e = \rho_w + \tilde{\sigma}_{c,e,ss}^2 = 0.0575$.

5 Efficiency

To study the efficiency of the competitive equilibrium, we consider a planner who can use taxes on capital and labor to manipulate the capital and labor wedges in the competitive equilibrium. Because the role of risk shocks can be understood in terms of time-varying labor and capital wedges, this allows us to ask if the planner would like a different macroeconomic behavior in response to risk shocks. The planner has to live with the fundamental frictions in the model and cannot eliminate incomplete risk sharing or prevent agents from saving and sharing aggregate risk.

This planner’s problem can be fully microfounded in an environment with a moral hazard problem with hidden trade, as in Di Tella and Sannikov [2016] or Di Tella [2019]. An entrepreneur’s idiosyncratic shock, as well as the resulting capital and labor decisions and hidden savings cannot be observed. The optimal private contract takes the form of the reduced-form incomplete risk sharing problem we’ve used this far. We can then ask what is the best that a planner can do subject to the same contractual frictions. The resulting planner’s problem coincides with the planner’s problem in the next sub-section. In the interest of simplicity, we focus here on the reduced-form problem where the planner can use taxes to distort wedges, but the contractual microfoundation is useful to understand why this is the right planner problem to consider, and where the source of inefficiency lies.

The main takeaway from this section is that the response of the competitive equilibrium to a risk shock is inefficient. Output and employment fluctuations are too large, and consumption should rise in recessions. In the competitive equilibrium, risk shocks create a small reduction in the capital wedge that stimulates investment, and a large increase in the labor wedge that discourages employment. As a result, risk shocks create recessions with positive co-movement in output, employment, investment, and consumption. The planner, instead, wants a large increase in the capital wedge that discourages investment, and a small reduction in the labor wedge that stimulates employment. Output, employment, and investment fall, but consumption goes up. Because consumption and investment are negatively correlated, fluctuations in output and employment are smaller than in the competitive equilibrium.

5.1 The Planner’s Problem

We consider the case where the planner can use labor and capital taxes with lump-sum rebates. In this case, the planner can effectively control employment ℓ_t , investment x_t , and therefore aggregate consumption c_t . However, in addition to the resource constraints, equations (2) and (3), and the law of motion of the exogenous states σ_t , equation (4), the

planner must take as given entrepreneurs' idiosyncratic risk. Computing $\tilde{\sigma}_{c,e,t}$, we obtain:⁵

$$\tilde{\sigma}_{c,e,t} = \frac{f(k_t, \ell_t)}{c_{e,t}} \rho_e \sigma_t = \frac{k_t^\alpha \ell_t^{1-\alpha}}{c_t} \frac{1+z_t}{z_t} \rho_e \sigma_t. \quad (37)$$

In addition, because all agents have access to the same financial market, the planner must respect their Euler equations and their aggregate risk sharing. Combining them, we get a law of motion for the consumption ratio, $z_t = c_{e,t}/c_{w,t}$,

$$\mu_{z,t} = (\rho_w - \rho_e) + \tilde{\sigma}_{c,e,t}^2, \quad \sigma_{z,t} = 0. \quad (38)$$

This means that z_t will be a state variable for the planner (z_0 is chosen optimally). If the planner could control agents' access to the financial market, preventing entrepreneurs and workers from trading intertemporally or across aggregate states, then z_t would not be a state variable. It's important to note that while the planner must respect agents' Euler equations, this does not restrict the control of aggregate consumption, which is a choice variable, not a state variable. For any behavior of aggregate consumption, as long as we respect the law of motion of z_t , there is a process for the real interest rate r_t and the price of aggregate risk π_t that satisfies agents' Euler equations and aggregate risk sharing.

The planner's objective function is a weighted average of the utility of the representative worker and entrepreneurs, $\gamma U_w + (1-\gamma)U_e$. Strictly speaking, we should specify the Pareto weight on each entrepreneur, γ_i . However, entrepreneurs' allocations differ only in their scale, so for aggregate allocations we only need to know the Pareto weight on entrepreneurs as a whole. It is useful to write an individual entrepreneur's utility as⁶

$$\mathbb{E} \left[\int_0^\infty e^{-\rho_e t} \log c_{i,t} dt \right] = \mathbb{E} \left[\int_0^\infty e^{-\rho_e t} \left(\log c_{e,t} - \frac{1}{2} \frac{1}{\rho_e} \tilde{\sigma}_{c,e,t}^2 \right) dt \right]. \quad (39)$$

The planner's objective function is therefore

$$\mathbb{E} \left[\int_0^\infty \gamma e^{-\rho_w t} \left(\log c_{w,t} - \frac{\ell_t^{1+1/\psi}}{1+1/\psi} \right) + (1-\gamma) e^{-\rho_e t} \left(\log c_{e,t} - \frac{1}{2} \frac{1}{\rho_e} \tilde{\sigma}_{c,e,t}^2 \right) dt \right]. \quad (40)$$

⁵In the moral hazard microfoundation, the IC constraint requires $\sigma_{U_{it}} = \frac{f(k_{it}, \ell_{it})}{c_{it}} \sigma_t$, and consumption is $c_{it} = \hat{c}_t \exp(\rho_e U_{it})$, where \hat{c}_t depends only on the aggregate states. Computing $\tilde{\sigma}_{c,e,t}$, we obtain (37).

⁶An individual entrepreneur's consumption follows $dc_{i,t}/c_{i,t} = \mu_{c,e,t} dt + \tilde{\sigma}_{c,e,t} dB_{i,t} + \sigma_{c,e,t} dZ_t$, so $c_{i,t} = (c_{i,0}/c_{e,0}) \times c_{e,t} \times \exp \left(\int_0^t \tilde{\sigma}_{c,e,t} dB_{i,t} - \frac{1}{2} \tilde{\sigma}_{c,e,t}^2 dt \right)$. His utility is

$$\begin{aligned} \mathbb{E} \left[\int_0^\infty e^{-\rho_e t} \log c_{i,t} dt \right] &= \log(c_{i,0}/c_{e,0}) + \mathbb{E} \left[\int_0^\infty e^{-\rho_e t} \left(\log c_{e,t} - \int_0^t \frac{1}{2} \tilde{\sigma}_{c,e,s}^2 ds \right) dt \right] \\ &= \log(c_{i,0}/c_{e,0}) + \mathbb{E} \left[\int_0^\infty e^{-\rho_e t} \left(\log c_{e,t} - \frac{1}{\rho_e} \frac{1}{2} \tilde{\sigma}_{c,e,t}^2 \right) dt \right] \end{aligned}$$

The first term captures the initial inequality among entrepreneurs. If they all had the same Pareto weight we would pick $c_{i0} = c_{e0}$ for all entrepreneurs.

Notice that because entrepreneurs are more impatient than workers, $\rho_e > \rho_w$, in the first best with perfect risk sharing they would receive zero consumption in the long run, $\lim_{t \rightarrow \infty} z_t = 0$. Here because entrepreneurs can save on their own and have a precautionary motive, their consumption will not vanish in the long run. The consumption ratio z_t will have an ergodic distribution, and will converge to a steady state in the absence of shocks. However, it is still true that entrepreneurs' utility vanishes from the objective function in the long run. As we'll see below, there is a fundamental disagreement between the planner and private entrepreneurs about their consumption profile. They save according to their Euler equations, but the planner would like them to follow an inverse Euler equation, that eliminates the precautionary saving motive for idiosyncratic risk. Because we care about the properties of the planner problem in the ergodic distribution, we can ignore entrepreneurs' utility in the objective function, and maximize only workers' utility subject to a given z_t . The entrepreneur's Pareto weight $(1 - \gamma) > 0$ matters for how we choose z_0 and how we converge to the ergodic distribution.

The planner's HJB equation is therefore

$$\begin{aligned} \rho_w V(k, \sigma, z) = \max_{c, x, \ell} & \log c + \log \left(\frac{1}{1+z} \right) - \frac{\ell^{1+1/\psi}}{1+1/\psi} + V'_k k(x - \delta) \\ & + V'_z z(\rho_w - \rho_e + \tilde{\sigma}_{c,e}^2) + V'_\sigma \theta_\sigma (\bar{\sigma} - \sigma) + \frac{1}{2} V''_{\sigma\sigma} \sigma v_\sigma^2 \end{aligned} \quad (41)$$

subject to

$$c + \phi(x)k = k^\alpha \ell^{1-\alpha} \quad (42)$$

and (37). Notice that, in this formulation, the planner does not care directly about entrepreneurs' consumption or exposure to idiosyncratic risk. The planner only cares because if he exposes them to risk $\tilde{\sigma}_{c,e,t}$ they will save (higher z_t in the future), leaving less consumption for workers, $c_{w,t} = c_t/(1+z_t)$.

Instead of working with the HJB equation, it's easier and more revealing to work with the co-states, $m_k = V'_k$, $m_\sigma = V'_\sigma$, $m_z = V'_z$. Taking first-order conditions, we have

$$\frac{1}{c} (1 - 2m_z z \tilde{\sigma}_{ce}^2) = \lambda \quad [c] \quad (43)$$

$$(1 - \alpha)(k/\ell)^\alpha \left(\lambda + 2m_z z \frac{\tilde{\sigma}_{ce}^2}{k^\alpha \ell^{1-\alpha}} \right) = \ell^{1/\psi} \quad [\ell] \quad (44)$$

$$\lambda \phi'(x) = m_k, \quad [x], \quad (45)$$

where λ is the Lagrange multiplier on the resource constraint. The first condition says that giving consumption not only delivers utility, but also relaxes entrepreneurs' risk sharing, $\tilde{\sigma}_{c,e}$. This determines the marginal value of goods λ . The second condition equates the marginal disutility of labor, $\ell^{1/\psi}$, to its marginal product, taking into account that more

labor increases entrepreneurs' exposure to idiosyncratic risk. The third condition says that the marginal value of more capital should be equated to its marginal cost. Together with the resource constraint we can solve for the controls c , x , ℓ , and λ as a function of the states and co-states. It's hard to obtain closed-form expressions, but it's easy to eliminate c and λ and obtain a system of two equations for x and ℓ .

Now we differentiate the HJB equation with respect to each state to obtain a law of motion for m_k and m_z (m_σ is not directly used):

$$\rho_w m_k = m_k(x - \delta) + 2m_z z \tilde{\sigma}_{ce}^2 \frac{\alpha}{k} + \lambda (\alpha(k/\ell)^{\alpha-1} - \phi(x)) + \mu_{m_k} m_k \quad [m_k] \quad (46)$$

$$\rho_w m_z = -\frac{1}{1+z} - 2m_z z \frac{\tilde{\sigma}_{ce}^2}{1+z} \frac{1}{z} + m_z(\rho_w - \rho_e + \tilde{\sigma}_{c,e}^2) + \mu_{m_z} m_z \quad [m_z]. \quad (47)$$

The equation for m_k resembles an asset-pricing equation for capital. Capital delivers dividends net of new investment, transformed into utility using λ . It grows at rate $x - \delta$ but it's discounted more heavily because it exposes entrepreneurs to idiosyncratic risk. With this interpretation, the first-order condition for investment, x , can be understood as a Tobin's Q expression, properly taking into account the role of idiosyncratic risk. The equation for m_z captures the fact that higher z means that a smaller fraction of consumption goes to workers (who are all the planner cares about in the long run formulation). On the other hand, higher z reduces entrepreneurs' exposure to idiosyncratic risk, and therefore the future value of z .

We construct a pair of C^2 functions $m_k(k, \sigma, z)$ and $m_z(k, \sigma, z)$. Using Ito's lemma we transform equations (46) and (47), together with the first-order conditions (43) and the resource constraint, into a system of two second-order PDEs and two algebraic constraints, using the the resource constraint and first FOC to eliminate c and λ . We can solve this differential-equation system using numerical methods analogous to those used to solve for the competitive equilibrium.

Implied prices, wedges, and taxes. After finding the optimal allocation, we can back out prices and wedges. Set $q_t = \phi'(x_t)$, $w_t = \ell_t^{1/\psi} c_t / (1 + z_t)$, $r_t = \bar{\rho}_t + \mu_{c,t} - \sigma_{c,t}^2 - \frac{z_t}{1+z_t} \tilde{\sigma}_{c,e,t}^2$, and $\pi_t = \sigma_{c,t}$. The implied labor and capital wedges, $\omega_{\ell,t}$ and $\omega_{k,t}$ are defined by:

$$\ell_t^{1/\psi} \frac{c_t}{1+z_t} = (1-\alpha)(k_t/\ell_t)^\alpha (1-\omega_{\ell,t}) \quad (48)$$

$$\alpha(k_t/\ell_t)^{\alpha-1} (1-\omega_{k,t}) - \phi(x_t) + q_t (x_t - \delta + \mu_{q,t} - (\bar{\rho}_t + \mu_{c,t} - \sigma_{c,t}^2) - \sigma_{c,t} \sigma_{q,t}) = 0. \quad (49)$$

The optimal allocation can be implemented with a labor and capital tax, $\tau_{\ell,t}$ and $\tau_{k,t}$. Entrepreneurs' after-tax expected marginal product of labor is $f'_\ell(k_t, \ell_t)(1-\tau_{\ell,t})$; for capital

<i>Percent Standard Deviations</i>			
<i>Var.</i>	<i>Comp. Eq.</i>	<i>Social Planner</i>	<i>Comparison</i>
y	1.7	0.62	↓
c	0.88	0.96	=
i	5.8	5.5	=
ℓ	2.2	0.56	↓
$\omega_{\ell t}$	2.2	0.9	↓
ω_{kt}	0.3	24	↑

Table 4: Percent Standard Deviations of Main Variables under Competitive Equilibrium and Social Planner's Allocation.

it's $f'_k(k_t, \ell_t)(1 - \tau_{k,t})$. We get

$$\tau_{\ell,t} = \omega_{\ell,t} - \left(\frac{(k_t/\ell_t)^{\alpha-1}}{(k_t/\ell_t)^{\alpha-1} - \phi(x_t)} \frac{1+z_t}{z_t} \rho_e \sigma_t^2 \right) \quad (50)$$

and

$$\tau_{k,t} = \omega_{k,t} - \left(\frac{(k_t/\ell_t)^{\alpha-1}}{(k_t/\ell_t)^{\alpha-1} - \phi(x_t)} \frac{1+z_t}{z_t} \rho_e \sigma_t^2 \right) \left(1 - \frac{1}{\alpha} \frac{\rho_e k_t q_t}{c_t} \right). \quad (51)$$

The tax does not affect the idiosyncratic risk, it's levied before the idiosyncratic shock realizes, so that it does not interfere with the incomplete risk sharing problem.

5.2 Numerical Results

The main result is that the recession pattern of the responses of capital, employment, and consumption in the competitive equilibrium is inefficient. In the constrained efficient allocation, output and employment fluctuations are smaller, and consumption moves counter-cyclically. Figure 4 shows the impulse response to a risk shock in the planner's allocation. While employment and investment fall, consumption expands. As a result, employment falls less than in the competitive equilibrium.

We can probe these results by looking at the implied labor and capital wedges. The competitive equilibrium featured a large increase in the labor wedge and a small reduction in the capital wedge after risk shocks. In contrast, the planner wants a small reduction in the labor wedge and a large increase in the capital wedge. The capital wedge twists the economy in the direction of less investment but more consumption. In contrast to the competitive equilibrium, employment and total output fall because with more consumption, the income effect reduces labor supply. This effect is reflected in an increase in real wages.

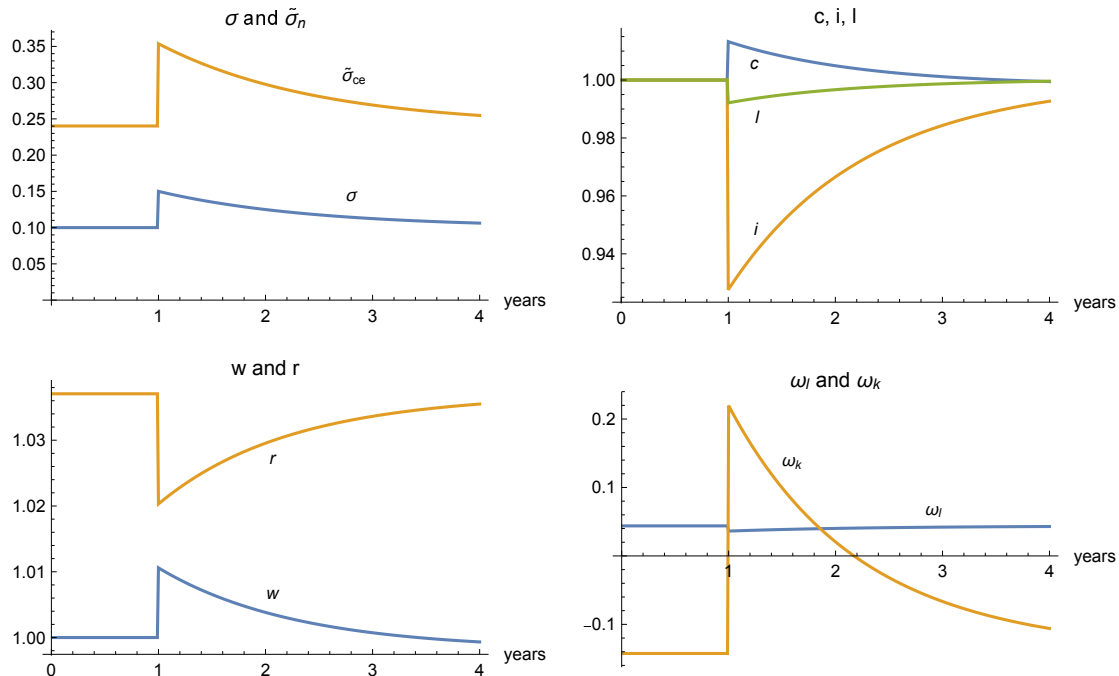


Figure 4: Impulse Response to a Risk Shock in the Planner's Solution

In contrast, in the competitive equilibrium the labor wedge reduced demand for labor, so wages fell on impact.

The real interest rate falls on impact. Lower expected consumption growth and a larger precautionary saving motive push interest rates down. In contrast, in the competitive equilibrium the larger precautionary motive lowered interest rates but it was compensated by higher expected consumption growth as the economy recovered. As a result, in the planner's allocation interest rates fall more than in the competitive equilibrium in response to risk shocks.

Table 4 compares the competitive equilibrium with the constrained efficient allocation. In the ergodic distribution the standard deviation of consumption and investment are roughly the same as in the competitive equilibrium, but their negative correlation (-0.88) means that the standard deviations of output and employment are lower. The average labor wedge (4.5 percent) and capital wedge (-8.4 percent) are larger than in the competitive equilibrium, but while the labor wedge becomes less volatile, the capital wedge becomes much more volatile.

The planner's allocation can be implemented by lowering labor taxes during recessions to stimulate employment and raising the capital tax to reduce investment and free more resources for consumption. Figure 5 shows the impulse response of the labor and capital tax to a risk shock. In the ergodic distribution $\tau_{\ell,t}$ has a mean of 1.7 percent (a tax on labor) with a standard deviation of 2.6 percent. It's positive in the steady-state, but becomes negative after risk shocks. In contrast, $\tau_{k,t}$ has a mean of -8 percent but a very




Figure 5: Impulse Response Function of the Labor and Capital Tax, τ_ℓ and τ_k

large standard deviation of 24 percent. It's negative in the steady-state but becomes very positive after risk shocks.

5.3 Understanding the Inefficiency

To understand why the competitive equilibrium is inefficient, it's useful to go over the role of hidden savings, and then use that to understand why the planner wants different wedges than in the competitive equilibrium.

The role of hidden savings. The inefficiency ultimately arises from the presence of hidden savings in the environment. Without hidden savings, the planner would like agents to follow the inverse Euler equation, which eliminates the precautionary saving motive. Understanding the optimal allocation without hidden savings helps us see what the planner is hoping to accomplish in the environment with hidden savings.

To see this issue in detail, consider the environment without hidden savings, so z_t is not a state variable. The planner can decide how much consumption goes to workers and entrepreneurs separately. The first-order conditions for consumption for finite t are

$$\gamma e^{-\rho_w t} c_{w,t}^{-1} = \lambda_t = (1 - \gamma) e^{-\rho_e t} c_{e,t}^{-1} \left(1 + \frac{1}{\rho_e} \tilde{\sigma}_{c,e,t}^2 \right). \quad (52)$$

Interpret λ_t be the pricing kernel (up to a constant), so workers satisfy their FOCs. Rearranging, we get that $\lambda_t e^{\rho_e t} c_{e,t} / (1 + \frac{1}{\rho_e} \tilde{\sigma}_{c,e,t}^2)$ must be a constant, and because an individual entrepreneur's $c_{i,t}$ is equal to $c_{e,t}$ times an idiosyncratic martingale, we get that $\lambda_t e^{\rho_e t} c_{i,t} / (1 + \frac{1}{\rho_e} \tilde{\sigma}_{c,e,t}^2)$ is a martingale. The term $c_{i,t} / (1 + \frac{1}{\rho_e} \tilde{\sigma}_{c,e,t}^2)$ is the marginal cost of delivering utility to the entrepreneur, so this expression says that the discounted marginal cost of utility must be a martingale. This is the essence of the inverse Euler equation. In the standard setting, the marginal cost of utility is the inverse of the marginal utility of consumption, $c_{i,t}^{-1}$. Here it's modified to take into account that higher consumption re-

duces the entrepreneur's exposure to idiosyncratic risk, as captured by (37). In a steady state without aggregate shocks, however, $\tilde{\sigma}_{c,e,t}$ is constant, so we recover the inverse Euler equation: $e^{(\rho_e - r_t)t} c_{i,t}$ is a martingale, or in flow form,

$$\mu_{c,e,t} = r_t - \rho_e. \quad (53)$$

The crucial difference with the Euler equation is that it's missing the precautionary savings term, $\tilde{\sigma}_{c,e,t}^2$.

Looking at (52), we see that, besides eliminating entrepreneurs' precautionary motive, the planner would like to front-load entrepreneurs' consumption relative to workers' (raise $z_t = c_{e,t}/c_{w,t}$) when idiosyncratic risk $\tilde{\sigma}_{c,e,t}$ is large, to relax idiosyncratic risk sharing. But when agents have access to hidden savings, he is forced to respect their Euler equation and aggregate risk sharing, captured by (38). If he wants to increase entrepreneurs' consumption, to relax the idiosyncratic risk sharing constraint (37), he must also increase workers' consumption.

Labor wedge. Coming back to the setting with hidden savings, we can study the inefficiency in terms of the implied labor and capital wedges, and how they compare to those in the competitive equilibrium. As in the competitive equilibrium, the difference in duration plays a central role in generating the asymmetric response of capital and labor. It is easier to focus first on the labor wedge, which does not involve dynamic considerations, and then extend the analysis to the capital wedge.

We can re-write the first-order condition for labor to obtain an expression for the labor wedge,

$$1 - \omega_{\ell,t} = \frac{\overbrace{-(2m_{z,t} z_t \tilde{\sigma}_{c,e,t}^2)}^{\text{improve risk sharing} > 0} + \overbrace{(2m_{z,t} z_t \tilde{\sigma}_{c,e,t}^2)}^{\text{worsen risk sharing} < 0} \times \frac{c_t}{y_t}}{\underbrace{1 + z_t}_{\text{income effect}}}. \quad (54)$$

The forces can be split into those that are primarily about the long run, and those that are important for business cycles.

Business cycles. The numerator of (54) captures the impact of employment on idiosyncratic risk. On the one hand, more employment increases entrepreneurs' exposure to idiosyncratic risk. Private agents also realize this, which is why they demand a risk premium to compensate them. But the planner also realizes that more aggregate consumption c_t relaxes the idiosyncratic risk sharing problem, as can be seen in equation (37). Essentially, if aggregate consumption is higher, entrepreneurs' consumption must also be higher, given the consumption ratio z_t . Because their consumption policy is $c_{e,t} = \rho_e n_{e,t}$, this means that

their wealth is higher and their exposure to idiosyncratic risk, $\tilde{\sigma}_{c,e,t} = \tilde{\sigma}_{n,e,t} = \frac{f(k_t, \ell_t)}{n_{e,t}} \sigma$ is lower, other things equal.

These countervailing forces create a labor wedge that goes down after risk shocks. Consider increasing employment by a small amount $d\ell > 0$, and using the extra output to increase aggregate consumption. Besides the usual considerations (marginal product vs. marginal rate of substitution between labor and consumption), this will improve idiosyncratic risk sharing. We can write entrepreneurs' exposure to idiosyncratic risk

$$\tilde{\sigma}_{c,e,t} = \frac{f(k_t, \ell_t) + f'_\ell(k_t, \ell_t)d\ell}{c_t + f'_\ell(k_t, \ell_t)d\ell} \frac{1 + z_t}{z_t} \rho_e \sigma_t. \quad (55)$$

Because $y_t = f(k_t, \ell_t) > c_t$, increasing output and consumption by the same amount, $f'_\ell(k_t, \ell_t)d\ell$, reduces entrepreneurs' exposure to idiosyncratic risk $\tilde{\sigma}_{c,e,t}$. Instead of increasing idiosyncratic risk, as private agents think, more employment actually improves idiosyncratic risk sharing. In consequence, the planner wants a lower labor wedge after risk shocks when σ_t is large. He wants to subsidize labor.

To formalize this, notice that the marginal value of an extra unit of aggregate consumption is not only that it delivers marginal utility, c_t^{-1} , but also the marginal improvement in idiosyncratic risk sharing, $-2c_t^{-1}m_{z,t}z_t\tilde{\sigma}_{c,e,t}^2 > 0$. This is reflected in the expression for λ , equation (43). It's useful to write

$$\lambda_t = \frac{1}{c_t} M_t, \quad (56)$$

where $M_t = 1 - 2m_{z,t}z_t\tilde{\sigma}_{c,e,t}^2 > 1$ captures the extra value of consumption from improving risk sharing.

When the planner considers the marginal value of labor, he takes into account that it adds to idiosyncratic risk, just like agents in the competitive equilibrium. That's what the term $2m_{z,t}z_t\tilde{\sigma}_{c,e,t}^2/(k_t^\alpha \ell_t^{1-\alpha})$ in equation (44) captures. But he also realizes that the extra output relaxes idiosyncratic risk sharing through higher consumption. That's why the marginal product of labor is weighted with λ_t . The numerator in the labor wedge equation (54) captures both these considerations. The two forces go in opposite directions and become larger after risk shocks, but it's clear that the improvement in idiosyncratic risk sharing dominates because $c_t/y_t < 1$. This gives us a labor wedge that decreases after risk shocks.

Long run. The denominator of (54) eliminates the $1 + z_t$ denominator in the definition of the labor wedge, equation (48), which captures an income effect on labor supply. The planner doesn't care about entrepreneurs' utility in the long run because they are more impatient. He would like to give them zero consumption, $z_t = 0$, but can't prevent them from saving following their Euler equation, which includes an inefficient precautionary saving motive, as explained above.

As a result, we have $z_t > 0$ and only a fraction $1/(1 + z_t)$ of the extra consumption goes to workers. But the planner gives workers only part of aggregate consumption not because he puts less than complete weight on them, in which case they would also be made to work more, but because of a constraint. He must give entrepreneurs some consumption, but he is not forced to make workers work more. The result is a positive labor wedge that undoes the income effect on labor supply. However, it's important to notice that because $\sigma_{z,t} = 0$, the denominator does not respond on impact to risk shocks, so this force does not play an important role in business cycles.

Capital wedge. The same forces apply to the capital wedge, but because capital is a long-duration asset, investment involves intertemporal tradeoffs. First, notice that $m_{k,t} = \lambda_t \phi'(x_t) = \lambda_t q_t$ is the marginal cost of new capital in utility terms. In standard models, it would be $c_t^{-1} q_t$, reflecting the forgone utility from consumption. Here, instead, the lower aggregate consumption also worsens idiosyncratic risk sharing, which is taken into account by $\lambda_t = c_t^{-1} M_t$. So the cost of investment for the planner is larger than what private agents realize, and moves with risk shocks. On the other hand, the marginal product of capital will not only be risky, as private agents realize, but also improve risk sharing through increased aggregate consumption, analogously to labor. So it will also be weighted by λ_t in the future.

Now we can use the law of motion for m_k , equation (46), and rewrite it in terms analogous to the competitive equilibrium condition, equation (22),

$$\alpha(k_t/\ell_t)^{\alpha-1} (1 - (m_{z,t} z_t 2\tilde{\sigma}_{c,e,t}^2)(1 - c_t/(k_t^\alpha \ell_t^{1-\alpha}))) - M_t \phi(x_t), \quad (57)$$

$$+ M_t q_t (x_t - \delta + \mu_{q,t} - (\bar{\rho}_t + \mu_{c,t} - \sigma_{c,t}^2) - \sigma_{c,t} \sigma_{q,t}), \quad (58)$$

$$+ M_t q_t (\mu_{M,t} + \sigma_{M,t}(\sigma_{q,t} - \sigma_{c,t}) + (\bar{\rho}_t - \rho_w)) = 0 \quad (59)$$

The differences from equation (22), help explain the inefficiency. First, the marginal product of capital is multiplied by $(1 - (2m_{z,t} z_t \tilde{\sigma}_{c,e,t}^2)(1 - c_t/(k_t^\alpha \ell_t^{1-\alpha})))$ instead of $(1 - \tilde{\sigma}_{c,e,t} \sigma_t)$. The planner internalizes that the marginal product of capital is risky but also helps improve idiosyncratic risk sharing through higher aggregate consumption. This is symmetric with labor.

Second, the cost of creating new capital $\phi(x_t)$, and the value of capital in the future q_t , is weighted by M_t because those resources could also be used to relax idiosyncratic risk sharing. If M_t is high today—for example, if σ_t is large—but is expected to be lower in the future ($\mu_{M,t} < 0$)—today is a bad time to invest. In this case, investment uses goods when their value from relaxing idiosyncratic risk sharing is high, and will deliver goods when their value relaxing idiosyncratic risk sharing is expected to be low. The term $\sigma_{M,t}(\sigma_{q,t} - \sigma_{c,t})$ captures the covariance between M_t and the marginal-utility-weighted value of capital $c_t^{-1} q_t$ to properly incorporate aggregate risk. In a steady-state without aggregate shocks M_t is

constant, $\mu_{M,t} = \sigma_{M,t} = 0$, and these dynamic issues vanish.

The last term, $\bar{\rho}_t - \rho_w > 0$, reflects that the planner doesn't care directly about entrepreneurs (in the long run), and so would like to discount the future using ρ_w rather than $\bar{\rho}_t = \frac{z_t}{1+z_t}\rho_e + \frac{1}{1+z_t}\rho_w$. The planner would like to give entrepreneurs zero consumption, $z_t = 0$, in which case $\bar{\rho}_t = \rho_w$, but cannot because they save to prevent this. This issue is analogous to the role of the $1 + z_t$ denominator in the expression for the labor wedge, (54). Ultimately, the disagreement reflects the inefficiency of entrepreneurs' precautionary saving motive, as we explained above. However, while this inefficiency may be important in the steady-state, it does not play a role in the impact of risk shocks, because z_t and therefore $\bar{\rho}_t$ do not respond on impact.

We can use equation (57) to obtain an expression for the capital wedge:

$$1 - \omega_{k,t} = \left(\frac{\overbrace{1 - (2m_{z,t}z_t\tilde{\sigma}_{c,e,t}^2)}^{\text{improve risk sharing}>0} + \overbrace{(2m_{z,t}z_t\tilde{\sigma}_{c,e,t}^2) \times \frac{c_t}{y_t}}^{\text{worsen risk sharing}<0}}{1 - \underbrace{(2m_{z,t}z_t\tilde{\sigma}_{c,e,t}^2)}_{\text{worse risk sharing}>0}} \right) \quad (60)$$

$$+ \underbrace{\frac{q_t}{\alpha(k_t/\ell_t)^{\alpha-1}}}_{\text{price-dividend ratio}} \left(\underbrace{\mu_{M,t} + \sigma_{M,t}\sigma_{q,t} - \sigma_{M,t}\sigma_{c,t}}_{\text{dynamic ineff.}} + \underbrace{(\bar{\rho}_t - \rho_w)}_{z_t>0} \right).$$

If we compare the expressions for the labor and capital wedge, equations (54) and (60), we see that if the price-dividend ratio for capital, $\frac{q_t}{\alpha(k_t/\ell_t)^{\alpha-1}}$ was zero, the dynamic part of the capital wedge would disappear. We wouldn't quite obtain the labor wedge $1 - \omega_{\ell,t}$, however, because instead of $1 + z_t$ in the denominator we have M_t . The difference is that investment requires goods which could be used to relax idiosyncratic risk sharing, captured by M_t , while labor requires only forgoing workers' utility from leisure, which cannot relax idiosyncratic risk sharing. In the ergodic distribution, most of the variation in the capital wedge comes from the second term in equation (60). The first term has a standard deviation of 2.5 percent, compared to 25 percent for the total capital wedge.

The role of $\rho_e > \rho_w$. We assumed entrepreneurs are more impatient than workers to obtain a stationary distribution of consumption and wealth, in both the competitive equilibrium and the planner's allocation. It may seem that this assumption is driving the result that in the long run, the planner does not care about entrepreneur's utility directly. That is, their utility vanishes from the planner's objective functions (40).

Here we'll show that $\rho_e > \rho_w$ is not essential to this issue. The fundamental issue is that the planner and private entrepreneurs disagree on their appropriate consumption profile. The planner would like to eliminate the precautionary saving motive for idiosyncratic risk and front-load their consumption, so in the long run the only reason he cares about their

consumption and risk sharing is to satisfy the hidden savings constraint.

An analogous situation would arise with $\rho_e = \rho_w$. In this case, because of the precautionary saving motive for idiosyncratic risk, entrepreneurs would accumulate all the wealth and consumption in the long run, $\lim_{t \rightarrow \infty} z_t = \infty$, both in the competitive equilibrium and the planner's allocation. The planner would be forced to give almost all of the consumption to entrepreneurs because of the hidden savings constraint. So while the utility of entrepreneurs would not vanish from the objective function (40), the allocation would be so far away from the distribution of consumption that optimizes the objective function in the absence of the hidden savings constraint, that on the margin the planner would not value entrepreneurs' utility directly (it vanishes relative to the utility of workers). The only motive to give entrepreneurs consumption or reduce their exposure to risk would come from the hidden savings constraint, as in the $\rho_e > \rho_w$ case.

Summary. In the planner's allocation the labor wedge goes down after risk shocks because, while the marginal product of labor has idiosyncratic risk, the planner internalizes that the extra output can also be used to improve idiosyncratic risk sharing through higher aggregate consumption. Both forces go in opposite directions and both become larger after risk shocks, but the improvement in idiosyncratic risk sharing dominates.

The same logic applies to investment, but the resources used to produce capital could also be used to improve idiosyncratic risk sharing through higher aggregate consumption. Periods when idiosyncratic risk is high, and especially if it's expected to be lower in the future, are particularly bad periods for investment. The planner does not want to use resources when their value relaxing idiosyncratic risk sharing is high, and get more output or capital in the future when the value of relaxing idiosyncratic risk sharing is low. So the capital wedge goes up after risk shocks.

The inefficiency arises from the presence of hidden savings, which are important in the microfoundation of the incomplete idiosyncratic risk-problem. The planner would like entrepreneurs to follow an inverse Euler equation which front-loads their consumption, especially when idiosyncratic risk is high. However, he must respect their Euler equations. This creates inefficiency both in the long run and in response to risk shocks.

6 Conclusions

In this paper we propose a theory of business cycles driven by spikes in risk premiums that act like negative demand shocks for capital and labor. In this view, recessions are periods of heightened uncertainty, when businesses and investors shrink from risky economic activity. There is a long tradition that attributes business cycles to time-varying risk premiums, going back to chapter 12 of the *General Theory* (Keynes [1936]). However, the co-movement

pattern of recessions poses a long-standing challenge to this view. It is hard to explain why employment, consumption, and investment contract simultaneously. This is essentially the reason why the macro literature has focused on productivity shocks and monetary shocks with nominal rigidities as drivers of business cycles. In this paper we aim to provide a novel theoretical framework to explain business cycles.

Our model is parsimonious, quantitatively plausible, and policy-relevant. The only deviation from the neoclassical growth model is uninsurable idiosyncratic risk on the business side. We show that risk shocks can generate quantitatively realistic business cycles, with co-movement between employment, consumption, and investment. The mechanism hinges on the interaction of risk premiums and the precautionary saving motive, and the different duration of capital and labor plays a central role. The resulting business cycles are not efficient. Employment and output fluctuations are too large, and consumption should move countercyclically. Optimal policy involves stimulating employment and consumption during recessions. This can be achieved by subsidizing labor and taxing capital during recessions.

References

- S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, pages 659–684, 1994.
- George-Marios Angeletos. Uninsured idiosyncratic investment risk and aggregate saving. *Review of Economic Dynamics*, 10(1):1–30, 2007.
- Cristina Arellano, Yan Bai, and Patrick J Kehoe. Financial frictions and fluctuations in volatility. *Journal of Political Economy*, 127(5):2049–2103, 2019.
- Robert J. Barro and Robert G. King. Time-separable preferences and intertemporal-substitution models of business cycles. *The Quarterly Journal of Economics*, 99(4): 817–839, 11 1984.
- Ben Bernanke, Mark Gertler, and Simon. Gilchrist. The financial accelerator in a quantitative business cycle framework. In John B. Taylor and Michael Woodford, editors, *Handbook of Macroeconomics*, volume 1, chapter 21, pages 1341–1393. Elsevier, 1999.
- Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.
- Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry. Really uncertain business cycles. *Econometrica*, 86(3):1031–1065, 2018.
- Markus Brunnermeier and Yuliy Sannikov. A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421, 2014.
- Ricardo Caballero and Alp Simsek. A risk-centric model of demand recessions and macro-prudential policy. Technical Report 23614, NBER, 2018.
- John Y Campbell and John H Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of political Economy*, 107(2):205–251, 1999.
- Varadarajan V Chari, Patrick J Kehoe, and Ellen R McGrattan. Business cycle accounting. *Econometrica*, 75(3):781–836, 2007.
- Lawrence Christiano., Roberto Motto, and Massimo Rostagno. Risk shocks. *American Economic Review*, 104(1):27–65, 2014.
- John H Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4): 1047–1108, 2011.
- Jan De Loecker and Jan Eeckhout. The rise of market power and the macroeconomic implications. Technical report, National Bureau of Economic Research, 2017.

- Sebastian Di Tella. Uncertainty shocks and balance sheet recessions. *Journal of Political Economy*, 125(6):2038–2081, December 2017.
- Sebastian Di Tella. Risk premia and the real effects of money. November 2019.
- Sebastian Di Tella and Yuliy Sannikov. Optimal asset management contracts with hidden savings. Stanford GSB, 2016.
- Chris Edmond, Virgiliu Midrigan, and Daniel Yi Xu. How costly are markups? Technical Report 24800, National Bureau of Economic Research, 2018.
- Emmanuel Farhi and François Gourio. Accounting for macro-finance trends: Market power, intangibles, and risk premia. Technical Report 25282, National Bureau of Economic Research, 2018.
- Nicolae Gârleanu and Stavros Panageas. Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685, 2015.
- Simon Gilchrist, Jae W Sim, and Egon Zakrajšek. Uncertainty, financial frictions, and investment dynamics. Technical Report 20038, National Bureau of Economic Research, 2014.
- Robert E. Hall. Employment fluctuations with equilibrium wage stickiness. *American Economic Review*, 95(1):50–65, March 2005.
- Robert E. Hall. Reconciling cyclical movements in the marginal value of time and the marginal product of labor. *Journal of Political Economy*, 117(2):281–323, April 2009.
- Robert E Hall. High discounts and high unemployment. *American Economic Review*, 107(2):305–30, 2017.
- Robert E. Hall. Using empirical marginal cost to measure market power in the us economy. Technical Report 25251, National Bureau of Economic Research, 2019.
- Zhigu He and Arvind Krishnamurthy. A model of capital and crises. *Review of Economic Studies*, 79(2):735–777, 2012.
- Zhiguo He, Bryan T Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. 2015.
- Bernard Herskovic, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics*, 119(2):249–283, 2016.

- Charles P Himmelberg, R Glenn Hubbard, and Inessa Love. Investor protection, ownership, and the cost of capital. *World Bank Policy Research Working Paper*, (2834), 2004.
- Urban Jermann and Vincenzo Quadrini. Macroeconomic effects of financial shocks. *American Economic Review*, 102(1):238–271, 2012.
- Kenneth Judd, Lilia Maliar, Serguei Maliar, and Rafael Valero. Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. *Journal of Economic Dynamics and Control*, 44:92–123, 2014.
- Patrick J Kehoe, Pierlauro Lopez, Virgiliu Midrigan, and Elena Pastorino. Asset prices and unemployment fluctuations. Working Paper 26580, National Bureau of Economic Research, December 2019.
- Rohan Kekre and Moritz Lenel. Redistribution, risk premia, and the macroeconomy. University of Chicago, May 2019.
- J.M. Keynes. *The General Theory of Employment, Interest and Money*. Harcourt, Brace, 1936.
- Mete Kilic and Jessica A Wachter. Risk, unemployment, and the stock market: A rare-event-based explanation of labor market volatility. *The Review of Financial Studies*, 31(12):4762–4814, 2018.
- Nobuhiro Kiyotaki and John Moore. Credit cycles. *Journal of Political Economy*, 105: 211–248, 1997.
- Francis A Longstaff and Jiang Wang. Asset pricing and the credit market. *The Review of Financial Studies*, 25(11):3169–3215, 2012.
- Césaire Assah Meh and Vincenzo Quadrini. Uninsurable investment risks. 2004.
- Tobias J Moskowitz and Annette Vissing-Jørgensen. The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review*, 92(4):745–778, 2002.
- Filippo Occhino and Andrea Pescatori. Debt overhang in a business cycle model. *European Economic Review*, 73:58–84, 2015.
- Julio J Rotemberg and Michael Woodford. The cyclical behavior of prices and costs. *Handbook of Macroeconomics*, 1:1051–1135, 1999.
- Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):24–49, 2005.