The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXX, NO. 3 • JUNE 2015

Innovation, Growth, and Asset Prices

HOWARD KUNG and LUKAS SCHMID*

ABSTRACT

We examine the asset pricing implications of a production economy whose long-term growth prospects are endogenously determined by innovation and R&D. In equilibrium, R&D endogenously drives a small, persistent component in productivity that generates long-run uncertainty about economic growth. With recursive preferences, households fear that persistent downturns in economic growth are accompanied by low asset valuations and command high-risk premia in asset markets. Empirically, we find substantial evidence for innovation-driven low-frequency movements in aggregate growth rates and asset market valuations. In short, equilibrium growth is risky.

AN ECONOMY'S LONG-TERM GROWTH prospects reflect its innovative potential. At a fundamental level, innovation is a key source of sustained growth in aggregate productivity. Empirical measures of innovation, such as research and development (R&D) expenditures, tend to be volatile and quite persistent. Such movements affect the dynamics of growth. Indeed, in U.S. post-war data, productivity growth exhibits long and persistent swings.[1] Similarly, innovation-driven growth waves associated with the arrival of new technologies, such as telecommunication, computers, and the internet, to name a few, are well documented.[2] Stock prices reflect such changes in growth prospects. Moreover, if agents fear that a persistent slowdown in economic growth will lower asset prices, these movements will give rise to high-risk premia in asset markets.

In this paper, we develop a general equilibrium model of innovation and R&D to link asset prices and aggregate risk premia to endogenous movements

*Howard Kung is at the London Business School. Lukas Schmid is at Duke University and the University of California, Los Angeles. We thank the Acting Editor Dave Backus and two anonymous referees as well as Hengjie Ai, Ravi Bansal, Gian Luca Clementi, John Coleman, Diego Comin, Max Croce, Ian Dew-Becker, Bernard Dumas, Lorenzo Garlappi, João Gomes, Francois Gourio, Leonid Kogan, Lars Lochstoer, Pascal Maenhout, Stavros Panageas, Dimitris Papanikolaou, Adriano Rampini, Norman Schurhoff, Tom Tallarini, Amir Yaron, and Lu Zhang for helpful discussions. We have also benefited from the comments of seminar participants at Bocconi, Boston University, Carnegie Mellon University, Cornell University, Duke University, Federal Reserve Bank of San Francisco, Imperial College, INSEAD, London Business School, London School of Economics, MIT, Paris School of Economics, UCLA, and Universitat Pompeu Fabra, as well as from conference participants at Western Finance Association, Society for Economic Dynamics, CEPR Macro Finance Conference London, European Finance Association, Tel Aviv Finance Conference, Econometric Society, American Finance Association, and NBER Spring Asset Pricing Meeting.

[1] See, for example, Gordon (2010) and Jermann and Quadrini (2004).

[2] For example, see Helpman (1998) and Jovanovic and Rousseau (2005).

DOI: 10.1111/jofi.12241

in long-term growth prospects. Our setup has two distinguishing features. First, we embed a stochastic model of endogenous growth based on industrial innovation[3] into an otherwise standard production economy. In this model, productivity growth is endogenous and sustained by the creation of new patented technologies through R&D. Patents represent an endogenous stock of intangible capital. Second, we assume that households have recursive preferences, so that they care about uncertainty regarding long-term growth prospects.

When calibrated to match empirical evidence on innovation and long-run economic growth, our model can quantitatively replicate key features of asset returns in the data. In particular, our model rationalizes a sizeable equity premium and a low and stable risk-free interest rate. Moreover, our model predicts a sizeable spread between the returns on physical capital and intangible capital, which is related to the value premium in the data. In short, we find that equilibrium growth is risky.

We first show that in the model innovation and R&D endogenously drive a small but persistent component in the growth rate of productivity. In our general equilibrium setting, these low-frequency movements in productivity trigger long and persistent swings in aggregate growth rates, such as consumption and output, which we label growth cycles. Intuitively, shocks affect the incentives to innovate, which in turn impact long-term growth prospects. Notably, transitory shocks in this setting have long-lasting permanent effects through the innovation channel and generate endogenous persistence in growth rates.

Thus, a bad temporary shock not only lowers the level of consumption and cash flows today, but also depresses long-term growth rates. When agents have recursive preferences, they are sensitive to both short-run and long-run uncertainty about consumption growth. Growth cycles help rationalize sizeable risk premia in asset markets, as agents fear that such prolonged slumps in economic growth coincide with low asset valuations. Similarly, agents save for extended low growth episodes, driving down the real interest rate. Furthermore, in the model, physical capital is endogenously more exposed to predictable variation in growth than intangible capital, which generates a sizeable value spread.

An innovation-driven persistent component in productivity growth provides an equilibrium foundation of long-run risks in the spirit of Bansal and Yaron (2004). More precisely, in our model, long-run productivity risks, in the sense of Croce (2014), arise naturally in equilibrium. Furthermore, persistent movements in expected productivity affect all aggregate growth rates and therefore give rise to equilibrium long-run consumption risks and cash flow risks.

The model helps to identify economic sources of long-run risks in the data. In particular, the model predicts that R&D and innovation are equilibrium determinants of productivity growth. In line with the predictions of the model, we provide novel empirical evidence that measures of innovation have significant predictive power for aggregate growth rates including productivity, consumption, and output growth at horizons of one to five years.

[3] Here, we build on the seminal work of Romer (1990) and Grossman and Helpman (1991).

While predictability in growth rates is at the core of the long-run risk model, empirical evidence regarding this channel is still limited. The model provides novel theoretical and empirical support for the notion that movements in long-term growth prospects are a significant source of priced risk in asset markets. Moreover, our results suggest that extending macroeconomic models to account for the endogeneity of innovation and long-term growth can make progress toward an environment that jointly captures the dynamics of aggregate quantities and asset markets. We therefore view stochastic models of endogenous growth as a useful tool for macrofinance.[4]

Our paper is related to several strands of literature in asset pricing, economic growth, and macroeconomics. The economic mechanisms driving the asset pricing implications are similar to those in the consumption-based long-run risks model of Bansal and Yaron (2004). We contribute to this literature by showing that predictable movements in growth prospects are an equilibrium outcome of stochastic models of endogenous growth and by providing novel empirical evidence identifying economic sources of long-run risks.

A number of recent papers examine the link between technological growth and asset prices. Garleanu, Panageas, and Yu (2012) model technological progress as the arrival of large, infrequent technological innovations and show that the endogenous adoption of these innovations leads to predictable movements in consumption growth and expected excess returns. Garleanu, Kogan, and Panageas (2012) examine the implications of the arrival of new technologies for existing firms and their workers, and show that, in an overlapping-generations model, innovation creates a systematic risk factor labeled displacement risk. The asset pricing implications of displacement risk are further examined in a model of heterogeneous workers and firms in Kogan, Papanikolaou, and Stoffman (2012). Pástor and Veronesi (2009) explain bubble-like behavior of stock markets in the 1990s by the arrival of new technologies.

While our model has implications for consumption dynamics and asset returns that are related to these models, our approach is different but complementary. In the above models of technology adoption, the arrival of new technologies is assumed to be exogenous. In contrast, we examine the asset pricing and growth implications of the endogenous creation of new technologies through R&D, which leads to a distinct set of empirical predictions. Moreover, by embedding a model of endogenous technological progress into a real business cycle model, our paper provides a straightforward extension of the workhorse model of modern macroeconomics.

In this respect, the paper is closer to recent attempts to address asset pricing puzzles within versions of the canonical real business cycle model. Starting with the habit-based models of Jermann (1998) and Boldrin, Christiano, and Fisher (2001), recent examples, such as Tallarini (2000), Campanale, Castro, and Clementi (2008), Kuehn (2008), Kaltenbrunner and Lochstoer (2010),

[4] In a companion paper, Kung (2015) shows that a similar mechanism coupled with imperfect price adjustment quantitatively rationalizes many aspects of the term structure of interest rates in a production economy.

and Papanikolaou (2011) explore endogenous long-run consumption risks in real business cycle models with recursive preferences, while Gourio (2012, 2013) examines disaster risks. Particularly closely related are recent papers by Croce (2014), Backus, Routledge, and Zin (2007, 2010), Gomes, Kogan, and Yogo (2009), and Favilukis and Lin (2013a, 2013b), who examine the implications of long-run productivity risk with recursive preferences for equity market returns. However, while they specify long-run productivity risk exogenously, our model shows how such risk arises endogenously through innovation. Our cross-sectional return implications are related to Gala (2010), Kogan and Papanikolaou (2010), and Lin (2012), who examine the effects of technological progress on the cross-section of returns.

Methodologically, our paper extends recent work by Comin and Gertler (2006) and Comin, Gertler, and Santacreu (2009). Building on the seminal work of Romer (1990) and Grossman and Helpman (1991), these papers integrate innovation and the adoption of new technologies into a real business cycle model and show that the resulting stochastic endogenous growth model features rich movements at a lower-than-business-cycle frequency, which they label medium-term business cycles. We contribute to this literature by linking medium-term cycles to long-run risks and aggregate risk premia.

The paper is structured as follows. In Section I, we describe our benchmark model. In Section II, we qualitatively explore the equilibrium growth and productivity processes and relate them to a canonical real business cycle model. In Section III, we quantitatively examine the asset pricing implications of our benchmark model and detail a number of empirical tests. Section IV concludes.

I. Model

In our baseline framework, we embed a model of industrial innovation in the tradition of Romer (1990) into a fairly standard macroeconomic model with convex adjustment costs and Epstein-Zin preferences. In the model, rather than assuming exogenous technological progress, sustained growth arises through the accumulation of patented intermediate goods (henceforth, patents) that facilitate the production of a final consumption good. New patents are created through innovation, which requires investment in R&D, and can be stored. Therefore, patents in this model represent an endogenous stock of intangible capital.

We start by describing in detail the production sector and the innovation process in our economy. We then present the household sector and define the general equilibrium.

A. Production

The production process involves three sectors. The final consumption good is produced in a perfectly competitive sector, namely, the final goods sector, using physical capital, labor, and patents. Stationary shocks drive stochastic fluctuations in the production of the final consumption good. Patents are produced in the intangible goods sector, where firms have monopoly power due to product differentiation. New patents are created by innovation through R&D in the innovation sector, which is also perfectly competitive.

Absent patents, decreasing returns to physical capital in the production function would imply that growth ceases in the long run without an exogenous trend component in the level of technology, which is a standard result from the Solow growth model. In contrast, in our setup, acquiring patents from the intangible sector facilitates production for a given stock of physical capital, allowing the final goods firm to grow and thereby creating demand for individual patents. That demand is met by the intangible sector, which earns profits from selling patents to the final goods firm by charging a markup over its marginal costs. Monopoly power is important as the associated profits provide rents for creating new patents. The innovation sector sells a newly developed patent to the intangible sector at the competitive price, which in equilibrium is equal to the present value of its profits.

Sustained growth is obtained in this economy because the demand for patents of the final goods firm creates new profit opportunities in the intangible sector and thus raises the incentives to create new patents through innovation. These new patents increase the efficiency of physical capital and thereby boost investment, creating even more demand for patents.

A.1. Final Good Sector

There is a representative firm that uses capital K_t, labor L_t, and a composite of patents G_t to produce final (consumption) goods according to the production technology

$$Y_t = (K_t^{\alpha}(\Omega_t L_t)^{1-\alpha})^{1-\xi} G_t^{\xi}, \tag{1}$$

where the composite G_t is defined as

$$G_t \equiv \left[\int_0^{N_t} X_{i,t}^{\nu} \, di \right]^{\frac{1}{\nu}} \tag{2}$$

and $X_{i,t}$ is the quantity of patent $i \in [0, N_t]$. Also, N_t is the measure of patents in use at date t, α is the physical capital share, ξ is the intangible capital share, and $\frac{1}{1-\nu}$ is the elasticity of substitution between patents with $\nu < 1$. We interpret N_t as the stock of intangible capital.

We introduce uncertainty into the model by means of an exogenous stochastic process Ω_t affecting the level of output. Importantly, Ω_t is assumed to follow a stationary Markov process by specifying that $\Omega_t = e^{a_t}$ and $a_t = \rho a_{t-1} + \epsilon_t$, with $\epsilon_t \sim N(0, \sigma^2)$ and $\rho < 1$. While Ω_t resembles labor-augmenting technology, it does not represent measured productivity in our setting, as we discuss in more detail below. Because of the stationarity of the forcing process, sustained growth arises endogenously from the development of new patents. We describe how new patents are developed by innovation below.

The firm's objective is to maximize shareholder value. This can be formally stated as

$$\max_{\{I_t, L_t, K_{t+1}, X_{i,t}\}_{t\geq 0, i\in[0,N_t]}} E_0\left[\sum_{t=0}^{\infty} M_t D_t\right], \tag{3}$$

where the firm's dividends are

$$D_t = Y_t - I_t - W_t L_t - \int_0^{N_t} P_{i,t} X_{i,t}\, di. \tag{4}$$

Here, M_t is the stochastic discount factor, I_t is investment in physical capital, W_t is the wage rate, and $P_{i,t}$ is the price per unit of patent i. The last term captures the costs of buying patents at time t. Prices $P_{i,t}$ are set by patent producers in the intangible sector, while the stochastic discount factor and the wage rate are determined in general equilibrium and are both taken as given by the final goods firm.

In line with the literature on production-based asset pricing, we assume that investment is subject to convex capital adjustment costs, so that the physical capital stock evolves as

$$K_{t+1} = (1-\delta)K_t + \Lambda\left(\frac{I_t}{K_t}\right)K_t, \tag{5}$$

where δ is the depreciation rate of physical capital and $\Lambda(\cdot)$ the capital adjustment cost function.[5]

A.2. Intangible Goods Sector

Patents are produced in the intangible goods sector. Patent producers have monopoly power. Given the demand schedules set by the final goods firm, monopolists producing the patents set the prices $P_{i,t}$ in order to maximize their profits $\Pi_{i,t}$. Patent producers transform one unit of the final good into one unit of their patented good. This fixes the marginal cost of producing one patent at unity.

Formally, monopolists solve the following static profit maximization problem each period

$$\Pi_{i,t} \equiv \max_{P_{i,t}} P_{i,t}\cdot X_{i,t}(P_{i,t}) - X_{i,t}(P_{i,t}). \tag{6}$$

The value $V_{i,t}$ of owning exclusive rights to produce patent i is equal to the present discounted value of the current and future monopoly profits, so that

$$V_{i,t} = \Pi_{i,t} + (1-\phi)E_t[M_{t+1}V_{i,t+1}], \tag{7}$$

[5] We specify $\Lambda(\cdot)$ as in Jermann (1998), $\Lambda\left(\frac{I_t}{K_t}\right) \equiv \frac{\alpha_1}{\zeta}\left(\frac{I_t}{K_t}\right)^{\zeta} + \alpha_2$. Here, $\frac{1}{1-\zeta}$ represents the elasticity of the investment rate with respect to Tobin's Q. The parameters α_1 and α_2 are set so that there are no adjustment costs in the deterministic steady state.

where ϕ is the probability that a patent becomes obsolete. This asset price is important in our model, as it provides the payoff to creating new patents through innovation as we describe next. This highlights the importance of monopoly power, as the associated profits provide the rents to innovation.

A.3. Innovation Sector

Innovators develop new patents used in the production of final output. They do so by conducting R&D, using the final good as input at unit cost. These newly developed patents can be sold to patent producers. Assuming free entry and perfect competition, the price of a new patent equals its value to the patent producer, namely, $V_{i,t}$.

We link the evolution of the intangible capital stock N_t to innovation as follows:

$$N_{t+1} = \vartheta_t S_t + (1-\phi)N_t, \tag{8}$$

where S_t denotes R&D expenditures (in terms of the final goods) and ϑ_t represents the productivity of the innovation sector, which is taken as exogenous by the R&D sector. In a similar spirit as Comin and Gertler (2006), we assume that this technology coefficient involves an externality effect

$$\vartheta_t = \frac{\chi \cdot N_t}{S_t^{1-\eta} N_t^{\eta}}, \tag{9}$$

where $\chi > 0$ is a scale parameter and $\eta \in [0, 1]$ is the elasticity of new patents with respect to R&D. This specification posits that concepts already discovered make it easier to come up with new ideas, $\partial\vartheta/\partial N > 0$, thus capturing positive spillovers of the aggregate stock of intangible capital as in Romer (1990), and that R&D investment has decreasing marginal returns, $\partial\vartheta/\partial S < 0$, capturing a congestion effect that raises the cost of developing new products as the aggregate level of R&D increases.[6]

B. Household

The household sector is standard. The representative household has Epstein-Zin preferences defined over consumption:

$$U_t = \left\{(1-\beta)C_t^{\theta} + \beta(E_t[U_{t+1}^{1-\gamma}])^{\frac{\theta}{1-\gamma}}\right\}^{\frac{1}{\theta}}, \tag{10}$$

[6] Similarly, this congestion externality can be thought of as giving rise to adjustment costs to investment in intangible capital, that is, R&D. Below, we will see that the optimality condition for R&D is $\frac{1}{\vartheta_t} = E_t[M_{t+1}V_{t+1}]$, equating the marginal cost of creating a new patent with its marginal benefit. Absent the congestion externality, this boils down to $1 = E_t[M_{t+1}V_{t+1}]$, a result analogous to Q-theory, in which case the absence of adjustment cost fixes marginal Q at unity.

where γ is the coefficient of relative risk aversion and $\psi \equiv \frac{1}{1-\theta}$ is the elasticity of intertemporal substitution. When $\psi \neq \frac{1}{\gamma}$, the agent cares about news regarding long-run growth prospects. We assume that $\psi > \frac{1}{\gamma}$, so that the agent has a preference for early resolution of uncertainty and dislikes uncertainty about long-run growth rates.

The household maximizes utility by participating in financial markets and by supplying labor. Specifically, the household can take positions $\mathcal{Z}_t$ in the stock market, which pays an aggregate dividend $\mathcal{D}_t$, and in the bond market B_t. Accordingly, the budget constraint of the household is

$$C_t + \mathcal{Q}_t \mathcal{Z}_{t+1} + B_{t+1} = W_t L_t + (\mathcal{Q}_t + \mathcal{D}_t)\mathcal{Z}_t + (1 + r_{f,t})B_t, \tag{11}$$

where $\mathcal{Q}_t$ is the stock price, $r_{f,t}$ is the risk-free rate, W_t is the wage, and L_t denotes hours worked.

We assume that stocks are claims to all the production sectors, namely, the final good sector, the intangible sector, and the innovation sector. Accordingly, we define the aggregate dividend as the net payout from the production sector,

$$\mathcal{D}_t = D_t + \int_0^{N_t} \Pi_{i,t}\, di - S_t. \tag{12}$$

C. Equilibrium and Asset Prices

We define an equilibrium for our economy in a standard way. In our setup, there is one exogenous state variable, Ω_t, and two endogenous state variables, the physical capital stock K_t and the intangible capital stock N_t. Given an initial condition $\{\Omega_0, K_0, N_0\}$ and the law of motion for the exogenous state variable Ω_t, an equilibrium is a set of sequences of quantities and prices such that (i) quantities solve producers' and the household's optimization problems and (ii) prices clear markets. Moreover, we focus on a symmetric equilibrium in which all patent producers make identical decisions. In the following, we describe the most important equilibrium conditions; we defer the complete list of all relationships characterizing the equilibrium to Appendix A.

The stochastic discount factor in the economy is given by

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{\theta-1} \left(\frac{U_{t+1}}{E_t\left(U_{t+1}^{1-\gamma}\right)^{\frac{1}{1-\gamma}}}\right)^{1-\gamma-\theta}, \tag{13}$$

where the second term, involving continuation utilities, captures preferences concerning uncertainty about long-run growth prospects. Optimality implies the following asset pricing conditions:

$$\mathcal{Q}_t = E_t[M_{t+1}(\mathcal{Q}_{t+1} + \mathcal{D}_{t+1})], \tag{14}$$

$$\frac{1}{1+r_t} = E_t[M_{t+1}]. \tag{15}$$

In equilibrium, the representative agent holds the entire supply of equities, while bonds are in zero net supply. The former is normalized to one (i.e., $\mathcal{Z}_t = 1 \quad \forall t$).

Since the agent has no disutility for labor, she will supply her entire endowment, which we normalize to unity, so that $L_t \equiv 1$.

The final good firm's optimality conditions are mostly standard. Denoting by $q_t = \frac{1}{\Lambda_t'}$ the shadow value of physical capital, the first-order condition for investment in physical capital is

$$1 = E_t\left[M_{t+1}\left\{\frac{1}{q_t}\left(\alpha(1-\xi)\frac{Y_{t+1}}{K_{t+1}} + q_{t+1}(1-\delta) - \frac{I_{t+1}}{K_{t+1}} + q_{t+1}\Lambda_{t+1}\right)\right\}\right]. \tag{16}$$

On the other hand, the final goods firm's demand for patent i is determined by

$$P_{i,t} = \left(K_t^{\alpha}(\Omega_t L_t)^{1-\alpha}\right)^{1-\xi} \xi \left[\int_0^{N_t} X_{i,t}^{\nu}\, di\right]^{\frac{\xi}{\nu}-1} X_{i,t}^{\nu-1}, \tag{17}$$

where it takes the price $P_{i,t}$ as given. In fact, $P_{i,t}$ is set by the monopolistically competitive producer of patent i. In a symmetric equilibrium, the monopolistically competitive characterization of the intangible goods sector à la Dixit and Stiglitz (1977) implies

$$X_{i,t} \equiv X_t \quad \text{and} \quad P_{i,t} \equiv P_t = \frac{1}{\nu}. \tag{18}$$

That is, each patent producer charges a markup $\frac{1}{\nu} > 1$ over unit marginal cost, so that its profits are

$$\Pi_{i,t} \equiv \Pi_t = \left(\frac{1}{\nu} - 1\right) X_t, \tag{19}$$

with $X_t = \left(\xi\nu\left(K_t^{\alpha}(\Omega_t L_t)^{1-\alpha}\right)^{1-\xi} N_t^{\frac{\xi}{\nu}-1}\right)^{\frac{1}{1-\xi}}$. Profits depend positively on the aggregate productivity shock Ω_t and thus are procyclical.

Discounted future profits on patents are the payoff to innovation. Thus, since the R&D sector is competitive, the optimality condition for R&D investment becomes

$$E_t[M_{t+1}V_{t+1}](N_{t+1} - (1-\phi)N_t) = S_t, \tag{20}$$

which says that the expected sales revenues equal costs, or equivalently, at the margin, $\frac{1}{\vartheta_t} = E_t[M_{t+1}V_{t+1}]$. By pinning down the amount of R&D investment, this condition is crucial in this model, as it ultimately determines the equilibrium growth rate of the economy. Importantly, the procyclicality of profits will be reflected in the dynamics of R&D.

C.1. Resource Constraint

Final output is used for consumption, investment in physical capital, factor inputs used in the production of patents, and R&D investment:

$$Y_t = C_t + I_t + N_t X_t + S_t, \tag{21}$$

$$= C_t + I_t + N_t^{1-\frac{1}{\nu}} G_t + S_t, \tag{22}$$

where the last equality exploits the optimality conditions and the term $N_t^{1-\frac{1}{\nu}} G_t$ captures the costs of patent production. Given that $\nu < 1$, reflecting monopolistic competition, it follows that a growing intangible capital stock increases the efficiency of patent production, as the costs fall with N_t growing.

C.2. Stock Market

Given our definition of stocks as claims to the net payout of all production sectors, in the symmetric equilibrium the aggregate dividend becomes

$$\mathcal{D}_t = D_t + \Pi_t N_t - S_t. \tag{23}$$

Defining the stock market value to be the discounted sum of future aggregate dividends and exploiting the optimality conditions, this value can be rewritten as

$$\begin{aligned} \mathcal{Q}_t = {} & q_t K_{t+1} + N_t(V_t - \Pi_t) \\ & + E_t\left[\sum_{i=0}^{\infty} M_{t+i+1}\big(V_{t+i+1}(N_{t+i+1} - (1-\phi)N_{t+i}) - S_{t+i+1}\big)\right], \end{aligned} \tag{24}$$

similar to Comin, Gertler, and Santacreu (2009). Thus, the stock market value consists of (i) the current market value of the installed capital stock (first term), (ii) the market value of currently used patents (second term), and (iii) the market value of patents to be developed in the future (third term). Therefore, the stock market values intangible capital and the option value of future intangibles in addition to the tangible capital stock.

II. Equilibrium Growth Risk

In our benchmark model, sustained growth is an equilibrium phenomenon resulting from agents' decisions. In contrast to variants of the workhorse stochastic growth model of dynamic macroeconomics, trend growth arises endogenously from the accumulation of patents rather than from an exogenous drift in productivity. In this section, we qualitatively examine the determinants of equilibrium growth and its dynamics. Most importantly, we document that in the model, movements in innovative activity generate predictable variation in growth rates. In the context of Bansal and Yaron (2004), the model identifies a novel source of long-run risks in the economy. In Section III, we quantitatively evaluate the ability of the model to rationalize aggregate asset risk premia and provide empirical evidence supporting this channel.

A. Endogenous Productivity

To start, it is convenient to represent the aggregate production function in our benchmark model in a form that permits straightforward comparison with specifications commonly used in macroeconomic models. Using the equilibrium conditions, final output can be rewritten as

$$Y_t = (\xi\nu)^{\frac{\xi}{1-\xi}} K_t^{\alpha}(\Omega_t L_t)^{1-\alpha} N_t^{\frac{\frac{\xi}{\nu}-\xi}{1-\xi}}. \tag{25}$$

To obtain sustained growth in this setting, we need to impose a parametric restriction. Technically, balanced growth requires the aggregate production function to be homogeneous of degree one in the accumulating factors K_t and N_t. In the following, we thus impose the restriction that $\alpha + \frac{\frac{\xi}{\nu}-\xi}{1-\xi} = 1$.[7] In this case, we obtain a standard neoclassical production function of the form $Y_t = Z_t K_t^{\alpha} L_t^{1-\alpha}$, where

$$Z_t \equiv \overline{A}(\Omega_t N_t)^{1-\alpha} \tag{26}$$

is the Solow residual, or productivity, with $\overline{A} \equiv (\xi\nu)^{\frac{\xi}{(1-\xi)}} > 0$.[8] The equilibrium productivity process thus contains a component driven by the exogenous forcing process, Ω_t, and an endogenous component reflecting the intangible capital stock, N_t. Importantly, while Ω_t is strictly stationary, productivity Z_t grows at an endogenous rate in equilibrium through the accumulation of patents. In this sense, technological progress is endogenous in our model.

B. Growth Cycles

Expression (26) highlights the importance of the accumulation of patents for productivity growth in the economy. The payoff to the creation of a new patent is its value, V_t. Naturally, the growth rate of intangible capital $\Delta N_{t+1} \equiv \frac{N_{t+1}}{N_t}$ reflects patent values and thus in equilibrium we have

$$\Delta N_{t+1} = (1-\phi) + E_t\left[\chi M_{t+1} V_{t+1}\right]^{\frac{\eta}{1-\eta}}. \tag{27}$$

This relationship, linking the growth rate of intangible capital to patent values, is central for the model and helps illustrate the main mechanisms determining

[7] Without that restriction, the economy will exhibit either decreasing or increasing returns to scale, so that growth rates either go to zero or will diverge in the long run. While positive growth would still obtain along the transition path in both cases, we impose the restriction to obtain positive growth in a stationary environment as does the bulk of the endogenous growth literature.

[8] Similar decompositions can be found in Ethier (1982), Comin and Gertler (2006), and others.

equilibrium growth. In particular, iterating this expression forward, we obtain

$$\Delta N_{t+1} = (1-\phi) + E_t\left[\chi^{\frac{1}{\eta}}\sum_{j=1}^{\infty} M_{t+j|t}(1-\phi)^{j-1}\Pi_{t+j}\right]^{\frac{\eta}{1-\eta}}, \qquad (28)$$

where $M_{t+j|t} \equiv \prod_s^j M_{t+s|t}$ is the j-step ahead stochastic discount factor and $M_{t|t} \equiv 1$. The equilibrium growth rate is therefore tied to discounted future profits on patents. We thus suggestively identify two major channels driving equilibrium growth, a *profit channel* and a *discount factor channel*.

The profit channel implies that the dynamics of the growth rate of the intangible capital stock reflect the dynamics of profits. Given $\Pi_t = \left(\frac{1}{\nu} - 1\right)\left(\xi\nu\left(K_t^{\alpha}(\Omega_t L_t)^{1-\alpha}\right)^{1-\xi} N_t^{\frac{\xi}{\nu}-1}\right)^{\frac{1}{1-\xi}}$, profit dynamics mirror the *level* of the forcing process Ω_t, and hence are procyclical and persistent. Accordingly, expression (28) suggests that the growth rate of intangible capital is procyclical and persistent. We label these cyclical movements of the equilibrium growth rate as *growth cycles*.

Figure 1 illustrates these dynamics. The figure displays the impulse responses of quantities in the intangible sector to a productivity shock. Importantly, after a positive shock, monopoly profits rise persistently. Intuitively, a positive shock in the final goods sector raises the demand X_t for patents, which translates directly into higher profits in the intangible sector. This raises the value of a patent, which triggers a prolonged increase in R&D and the growth rate of intangible capital. Finally, persistent R&D dynamics lead to a persistent expected productivity growth component.

The discount factor channel, on the other hand, points to the relevance of risk considerations for equilibrium growth dynamics, which can be seen in equation (17). With recursive preferences, the stochastic discount factor is sensitive to the entire intertemporal distribution of risk. Thus, a higher degree of predictability in consumption growth translates into a more volatile stochastic discount factor, and a more variable discount factor is reflected in a more volatile equilibrium growth rate. This feedback channel between the discount factor and endogenous growth dynamics implies a long-run growth amplification mechanism that generates quantitatively significant long-run risks.

C. Equilibrium Long-Run Productivity Risk

In our economy, persistent variation in the growth rate of the intangible stock is reflected in an endogenous persistent component in the dynamics of productivity Z_t by virtue of expression (26). Following Croce (2014), recent work in equilibrium asset pricing with production specifies productivity growth to contain a small persistent component, or, in other words, to exhibit long-run productivity risk (Gomes, Kogan, and Yogo (2009), Backus, Routledge, and

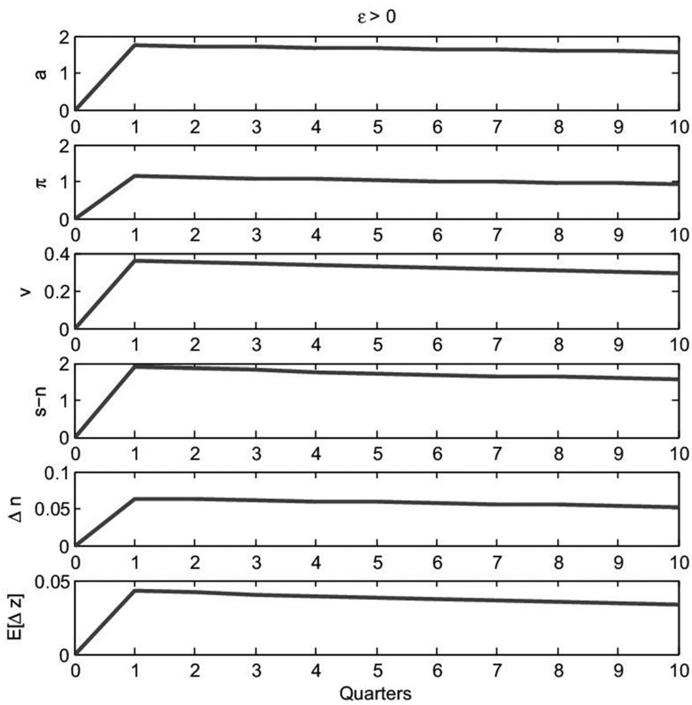

Figure 1. Endogenous growth mechanism. This figure plots impulse response functions in the benchmark model for the exogenous component of technology a, monopoly profits π, market value of patents v, R&D intensity $s-n$, R&D stock growth Δn, and expected productivity growth $E[\Delta z]$ to a positive productivity shock ($\epsilon > 0$).

Zin (2007, 2010), Favilukis and Lin (2013a, 2013b)). Our model thus gives an equilibrium interpretation to this channel.

Given the mean-reversion of Ω_t, the growth rate is negatively autocorrelated. However, if ρ is sufficiently close to one, the growth rate of Ω_t is close to i.i.d. and we may write

$$E_t\left[\Delta \log Z_{t+1}\right] \approx (1-\alpha)\chi\left(\frac{S_t}{N_t}\right)^{\eta}. \tag{29}$$

Any persistent movement in R&D thus provides an equilibrium source of long-run productivity risk. More precisely, such long-run productivity risk is driven by the dynamics of $\frac{S_t}{N_t}$, a ratio that we refer to as *R&D intensity*.

Quantitatively, the strength of this channel depends on the dynamics of R&D intensity. Figure 2 depicts a series of an empirical counterpart of R&D intensity, measured as the ratio of private business R&D investment (supplied by the National Science Foundation) to the stock of R&D (supplied by the Bureau of Labor Statistics) in the United States between 1953 and 2008. We discuss our data sources in more detail in the next section. Casual inspection suggests that

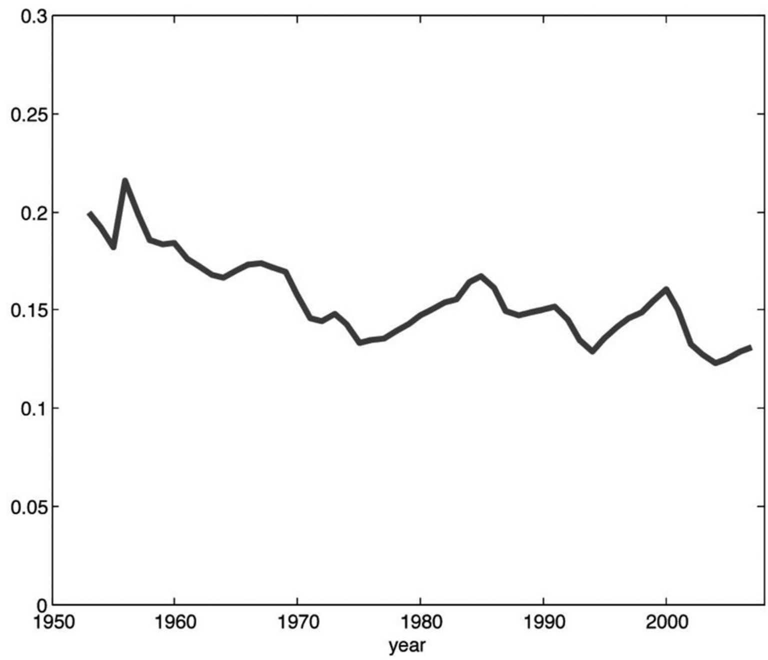

Figure 2. Empirical R&D intensity. This figure plots the annual R&D intensity, defined as the ratio of private business R&D investment to the stock of R&D, from 1953 to 2008. The data are from the National Science Foundation and the Bureau of Labor Statistics.

this driver of productivity growth expectations is volatile and highly persistent. In fact, in the data, the autocorrelation is 0.93. In our quantitative work, we calibrate our model carefully to be consistent with observed properties of innovation, and specifically R&D intensity. Moreover, we show through model simulations that exposure to such equilibrium long-run productivity risk helps in rationalizing sizeable risk premia in asset markets. This is because the ensuing predictable variation in productivity growth leads to both substantial long-run consumption risks as well as long-run cash flow risks. Thus, our model allows us to identify economic sources of long-run risks.

D. Growth Cycles versus Business Cycles

Most of the literature in equilibrium asset pricing with production operates in versions of the workhorse real business cycle model in which trend growth is specified exogenously. In our quantitative work, we contrast the implications of our benchmark endogenous growth model with those of a nested standard real business cycle model with exogenous growth. The real business cycle model we consider is a version of our benchmark model with constant R&D intensity. Specifying R&D intensity exogenously is equivalent to specifying an exogenous trend growth component in productivity, as in standard real business cycle models. We thus contrast the asset pricing implications of growth cycles and business cycles.

One way to interpret the differences between our model and the real business cycle framework is that the trend component of the productivity process, N_t, is endogenous and fluctuates in our setup, while it is exogenous and typically deterministic in real business cycle models. Through this channel, transitory shocks have permanent and persistent effects in our model. This is important from an asset pricing perspective, as Alvarez and Jermann (2005) show that, in an economy exclusively driven by transitory shocks, the term premium is the highest risk premium.

III. Quantitative Implications

In this section, we calibrate our model and explore its ability to replicate key moments of both macroeconomic quantities and asset returns. We view our endogenous growth model as a theory of long-run movements and therefore, rather than match standard business cycle moments, we parameterize it to be quantitatively consistent with long-run growth cycles by isolating the low-frequency components of growth rates using a bandpass filter. On the other hand, we find it instructive to compare our benchmark model with a version in which trend growth is given exogenously, in the spirit of a business cycle model. In the following, we refer to the benchmark endogenous growth model as the growth cycle model and the exogenous growth counterpart as the business cycle model.[9]

The models are calibrated at a quarterly frequency. The empirical moments correspond to the U.S. sample from 1953 to 2008. We focus on this particular period as R&D data become available only in 1953. The model is solved using second-order perturbation methods.

A. Calibration

Our benchmark model requires that we specify 13 parameters: three for preferences, seven relating to the final goods production technology, and three for the innovation technology. We focus on long-run growth cycles to help us determine key parameters. Note that we measure long-run growth cycles as movements at frequencies between 100 and 200 quarters that we isolate using a bandpass filter. In particular, we target the average growth rate of the economy and the growth cycle volatilities σ^{GC} of output, consumption, investment, and R&D intensity. Table I summarizes our parameter choices.

To construct our targets, we need to find empirical measures of innovation and R&D intensity, $\frac{S}{N}$. Our empirical series for S_t measures private business

[9] While there is no exactly corresponding model with exogenous growth, we find our choice natural as it facilitates comparison. The main conclusions are robust across a broad spectrum of exogenous growth models. Extensive robustness checks with other exogenous growth specifications are available from the authors on request.

Table I
Calibration

This table reports the benchmark quarterly calibration used for the benchmark growth cycle and business cycle models.

Parameter	Description	Growth Cycle	Business Cycle
β^4	Subjective discount factor	0.984	0.984
ψ	Elasticity of intertemporal substitution	1.85	1.85
γ	Risk aversion	10	10
ξ	Patent share	0.5	–
ν	Inverse markup	0.6	–
α	Capital share	0.35	0.35
ρ^4	Autocorrelation of Ω	0.95	0.95
χ	Scale parameter	0.332	–
ϕ	Patent obsolescence rate	0.0375	–
η	Elasticity of new patents with respect to R&D	0.83	–
δ	Depreciation rate of capital stock	0.02	0.02
σ	Volatility of exogenous shock ϵ	1.75%	0.97%
ζ	Investment adjustment cost parameter	3.3	3.3
$\mu * 4$	Trend growth rate	–	1.90%

R&D investment and comes from the National Science Foundation. The Bureau of Labor Statistics (BLS) constructs the R&D stock by accumulating R&D expenditures and allowing for depreciation, much in the same way as the physical capital stock is constructed. We thus use the R&D stock as our empirical counterpart for N_t to be consistent with the accumulation process in (8). For consistency, we use the same depreciation rate ϕ in our calibration as the BLS uses in its calculations. The R&D stock can be interpreted as measuring the economic benefits of R&D that spill over from the innovating firm to other firms. We provide further details on the data sources in Appendix A.

We start by discussing the less standard parameters. We set χ, which is a pure scaling parameter, to match the average growth rate of the U.S. economy in our sample. We pick η, the elasticity of new patents with respect to R&D, to match the growth cycle volatility of R&D intensity. This parameter can be thought of as an adjustment cost parameter for R&D. Furthermore, our choice of η is within the range of panel and cross-sectional estimates from Griliches (1990). Analogously, we set ζ to match the growth cycle volatility of investment. We choose σ, the volatility of the exogenous component of productivity, to match the growth cycle volatility of output. Finally, we choose ρ to be consistent with the autocorrelation of R&D intensity. This puts further discipline on the importance of movements in innovation as a source of long-run productivity risk, as becomes apparent in expression (29).

The choices of the remaining parameters follow the literature. Recursive preferences have been used extensively in recent work in asset pricing. We follow this literature and set preference parameters to standard values that are also supported empirically. The elasticity of intertemporal substitution ψ

is set to 1.85 and the coefficient of relative risk aversion γ is set to 10.[10] An elasticity of intertemporal substitution larger than one is consistent with the notion that an increase in uncertainty lowers the price-dividend ratio. Note that, in this parametrization, $\psi > \frac{1}{\gamma}$, which implies that the agent dislikes shocks to expected growth rates. The subjective discount factor β is set to an annualized value of 0.984 so as to be consistent with the level of the risk-free rate.

In the final goods sector, α is set to 0.35 to match the average capital share and the depreciation rate of capital δ is set to 0.02 to match the average capital investment rate, which are standard in the macroeconomic literature. The parameter ξ is set to 0.5 to accord with the choice in Comin and Gertler (2006).

The inverse markup parameter in the intangible sector ν is set to 0.6 to be consistent with the balanced growth restriction. While markups are generally difficult to measure, and especially so on intangible capital, varying the parameter around a reasonable range does not change our key quantitative results.[11] Since we interpret the variety of patents as the stock of R&D, as discussed, we can interpret ϕ as the depreciation rate of the R&D stock. Hence, we set ϕ to 0.0375, which corresponds to an annualized depreciation rate of 15%, which is a standard value and assumed by the BLS in the R&D stock calculations.

We calibrate the real business cycle model to facilitate direct comparison with our benchmark model. To do so, we set a trend growth parameter μ equal to 1.90% to match average output growth, and we adjust the volatility of the forcing process to match the volatility of consumption growth in the benchmark model. The remaining parameter choices are identical to those of the benchmark model.

A.1. Implications for Growth and Cycles

To assess our benchmark calibration, we start by discussing its implications for steady-state growth. We then explore economic fluctuations at higher (business cycle) and lower (growth cycle) frequencies. The nature of fluctuations in the model will be a key determinant of asset prices.

A.1.1. Steady-State Growth

Since trend growth is an endogenous variable in our model, the deterministic steady-state growth rate is a function of the deep parameters of the model. While closed-form expressions for the steady-state growth rate are not available, Table II illustrates the relationship between the model parameters and trend growth through comparative statics analysis via our numerical solution.

[10] The choice of the elasticity of intertemporal substitution is consistent with the estimation evidence in Fernandez-Villaverde et al. (2012), while Bansal, Kiku, and Yaron (2013) use Euler conditions and a GMM estimator to provide empirical support for these parameter values.

[11] Competition has an effect on the equity premium in the model, as reducing the markup lowers the volatility of cash flows, which leads to a smaller risk premium. On the other hand, the quantitative effects in the present specification are small.

Table II

Determinants of Steady-State Growth

This table reports comparative statics analysis of the deterministic steady-state growth rate Δy_{ss} from the benchmark growth cycle model. Panel A reports the impact of varying preferences parameters (around the benchmark calibration) on the steady-state growth rate. Panel B reports the impact of varying technological parameters (around the benchmark calibration) on the steady-steady growth rate.

Panel A: Comparative Statics with Preference Parameters					
β	0.9945	0.995	0.9955	0.996	0.9965
Δy_{ss}	0.22%	0.30%	0.38%	0.45%	0.53%
ψ	1.70	1.75	1.80	1.85	1.90
Δy_{ss}	0.42%	0.43%	0.44%	0.45%	0.46%
γ	5	10	15	20	25
Δy_{ss}	0.45%	0.45%	0.45%	0.45%	0.45%
Panel B: Comparative Statics with Technological Parameters					
ξ	0.496	0.498	0.500	0.502	0.504
Δy_{ss}	0.62%	0.54%	0.45%	0.37%	0.29%
ν	1.63	1.64	1.65	1.66	1.67
Δy_{ss}	0.69%	0.57%	0.45%	0.34%	0.24%
χ	0.3314	0.3317	0.3320	0.3323	0.3326
Δy_{ss}	0.445%	0.449%	0.453%	0.457%	0.461%
η	0.8290	0.8295	0.8300	0.8305	0.8310
Δy_{ss}	0.464%	0.459%	0.453%	0.448%	0.443%
ϕ	0.9615	0.9620	0.9625	0.9630	0.9635
Δy_{ss}	0.36%	0.40%	0.45%	0.50%	0.55%
δ	0.0190	0.0195	0.0200	0.0205	0.0210
Δy_{ss}	0.53%	0.49%	0.45%	0.42%	0.38%
ζ	0.05	0.06	0.07	0.08	0.09
Δy_{ss}	0.45%	0.45%	0.45%	0.45%	0.45%

We begin with the preference parameters (top panel). A higher value for the time-preference parameter β implies that the agent values the future more relative to the present. Hence, the agent is willing to defer consumption and invest more, which leads to higher growth. An increase in the intertemporal elasticity of substitution ψ means that the agent is less concerned about smoothing the consumption path and therefore leads to higher growth. Since we are analyzing the deterministic steady-state, the coefficient of relative risk aversion does not affect trend growth.[12]

On the technology side, increasing the patent share ξ leads to a reallocation of resources from physical capital to intangible capital. As intangible capital has a higher depreciation rate than physical capital, this reduces production efficiency and therefore lowers growth. Increasing the parameter that determines the average markup ν has two opposing effects. First, increasing the markup, holding all else equal, increases monopoly profits in the intermediate sector.

[12] Note that, in the stochastic steady-state, higher risk aversion will increase the precautionary savings motive of the agent and raise the average growth rate.

Table III

Macro Moments

This table presents annualized macroeconomic moments from the data, the benchmark growth cycle model and the business cycle model. Panel A reports the average long-run growth rate and volatilities of low-frequency components of output growth Δy, consumption growth Δc, physical investment growth Δi, and R&D intensity S/N. The low-frequency components are obtained using the bandpass filter from Christiano and Fitzgerald (2003) and isolating frequencies between 100 and 200 quarters. Panel B reports short-run volatilities of output growth Δy, consumption growth Δc, physical investment growth Δi, R&D expenditures growth Δs, and productivity growth Δz. Panel C reports first autocorrelations for these macro growth rates and for Tobin's Q. The model statistics correspond to population moments.

	Data	Growth Cycle	Business Cycle
Panel A: Growth Cycle Statistics			
$E[\Delta y]$	1.90%	1.90%	1.90%
$\sigma_{\Delta y}^{GC}$	0.24%	0.22%	0.13%
$\sigma_{\Delta c}^{GC}$	0.28%	0.24%	0.15%
$\sigma_{\Delta i}^{GC}$	0.18%	0.17%	0.09%
$\sigma^{GC}(S/N)$	0.71%	0.72%	–
Panel B: Business Cycle Statistics			
$\sigma_{\Delta c}$	1.42%	1.42%	1.42%
$\sigma_{\Delta c}/\sigma_{\Delta y}$	0.61	0.61	1.13
$\sigma_{\Delta i}/\sigma_{\Delta c}$	4.38	2.23	0.79
$\sigma_{\Delta s}/\sigma_{\Delta y}$	2.10	1.64	–
$\sigma_{\Delta z}/\sigma_{\Delta y}$	1.22	1.52	1.54
Panel C: Autocorrelation			
$AC1(\Delta z)$	0.09	0.11	−0.020
$AC1(\Delta c)$	0.40	0.39	−0.002
$AC1(\Delta y)$	0.37	0.21	0.001
$AC1(\Delta i)$	0.25	0.14	0.012
$AC1(Q)$	0.95	0.96	0.89

Second, a higher markup depresses the demand for intermediate goods inputs, which reduces monopoly profits. In our benchmark calibration, the second effect dominates, and therefore a higher average markup lowers steady-state growth. A higher scale parameter χ directly raises the level of productivity in the R&D sector and therefore increases growth. A higher η increases the marginal returns to R&D, which raises growth. Increasing the obsolescence rate of the R&D stock ϕ reduces the returns to R&D and therefore growth declines. Analogously, a higher depreciation rate of physical capital stock δ lowers growth.

A.1.2. Growth and Business Cycles

Table III reports the main macro moments of the benchmark model and the corresponding business cycle model. As targeted, the benchmark model is quantitatively in line with the average growth rate of the economy and the growth cycle volatilities σ^{GC} of output, consumption, investment, and R&D intensity in the data. In contrast, while similarly calibrated, the business cycle counterpart generates quantitatively insignificant growth cycles.

Panel B reports standard business cycle statistics from simulations. The model is also reasonably consistent with basic business cycle properties of the U.S. economy. In particular, our benchmark model does just about as well as the business cycle model in explaining short-run fluctuations. Both specifications match the low volatility of consumption growth of the post-war era in the United States. On the other hand, they all predict investment to be too smooth. This is because the benchmark model is calibrated to generate realistic growth cycles, which, in sharp contrast to business cycle fluctuations in investment, are actually significantly smoother than the corresponding movements in consumption. This suggests that the pronounced movements of investment at business cycle frequencies are driven by a different set of shocks than the long-run movements our model readily captures.[13]

Moving beyond business cycle volatilities to autocorrelations of growth rates, as reported in Panel C, reveals a striking difference between the benchmark model and the business cycle model. While the former qualitatively and sometimes quantitatively captures the substantial autocorrelation of most macro variables in the data, the corresponding persistence implied by the business cycle model is virtually zero, and sometimes slightly negative. This is remarkable as the exogenous stochastic driver of productivity is the same across model specifications. The lack of persistence in growth rates has long been identified as a weakness of the real business cycle model (e.g., Cogley and Nason (1995)). In stark contrast, endogenous movements in R&D induce a strong propagation mechanism in our benchmark model that translates persistence in transitory shocks to (i) persistence in the levels of macro variables (around the trend) and (ii) persistence in the time-varying trend growth rate. Intuitively, a good shock encourages innovation and thereby boosts growth in the long run even further.

In sum, both model specifications exhibit a similar amount of variation at business cycle frequencies. In an asset pricing context, we refer to such movements as short-run risk. However, in sharp contrast to the business cycle model, our benchmark model exhibits significant persistence at lower frequencies, namely, growth cycles. In an asset pricing context, we relate such movements to long-run risks in the sense of Bansal and Yaron (2004), in a way made precise below. The presence of growth cycles in our economy leads to substantially different asset pricing implications, as we explore in the next section.

B. Asset Pricing Implications

Table IV reports the asset pricing implications of the benchmark model and alternative specifications. To understand these results, it is instructive to compare the asset pricing implications of the benchmark model with those of

[13] Similarly, we abstract from endogenous movements in the labor supply, as those drive a large proportion of the fluctuations at business cycle frequencies.

Table IV

Asset Pricing Implications

This table reports the asset pricing implications for the benchmark growth cycle model, a high volatility calibration of the growth cycle model, and the business cycle model. More specifically, the volatility calibration corresponds to calibrating the volatility parameter σ in the growth cycle model to match consumption volatility in the post–Great Depression sample (1930 to 2008). Also, the scale parameter χ and the subjective discount factor β are adjusted to match the average output growth rate and risk-free rate from the benchmark model, while all of the other parameters are kept the same as the benchmark calibration. Panel A reports the means of the risk-free rate r_f, the risk premium on the aggregate stock market $E[r_m^* - r_f]$, the risk premium on physical capital $E[r_d^* - r_f]$, and the spread between physical capital and intangible capital $E[r_d^* - r_{ic}^*]$. Panel B reports the standard deviation of these returns. The risk premiums are levered following Boldrin, Christiano, and Fisher (2001). The model statistics correspond to population moments.

	Growth Cycle	Growth Cycle—Volatility Calibration	Business Cycle
Panel A: Mean			
$E[r_f]$	1.21%	1.21%	2.61%
$E[r_m^* - r_f]$	2.92%	5.76%	0.12%
$E[r_d^* - r_f]$	4.10%	8.33%	0.12%
$E[r_d^* - r_{ic}^*]$	3.27%	6.89%	–
Panel B: Standard Deviation			
$\sigma_{\Delta c}$	1.42%	2.72%	1.42%
σ_{r_f}	0.30%	0.38%	0.05%
$\sigma_{r_m^* - r_f}$	4.86%	6.73%	2.27%
$\sigma_{r_d^* - r_f}$	7.08%	9.49%	2.27%
$\sigma_{r_d^* - r_{ic}^*}$	5.13%	7.81%	–

the business cycle specification, which are radically different. These differences are inherently linked to the presence of growth cycles in our benchmark economy. Growth cycles and persistent uncertainty about expected growth prospects are akin to long-run risks, and give rise to both *long-run consumption risks* and *long-run cash flow risks*.

Our benchmark model generates a low and smooth risk-free rate, while in the business cycle model the risk-free rate is counterfactually high. Intuitively, in the presence of long-run consumption risks, agents with a preference for early resolution of uncertainty save for persistent low-growth episodes, lowering the equilibrium interest rate. In contrast, in the business cycle model expected growth prospects are roughly constant, so that agents want to borrow against their future income, which can only be prevented by a prohibitively high equilibrium interest rate as documented in the table.

The presence of long-run cash flow risks, on the other hand, renders stocks risky. Consistent with the multisector structure of our benchmark model, the stock market is a claim to the net payout from production. Quantitatively, the benchmark model generates a sizeable excess return on the aggregate stock market of close to 3%, with a volatility close to 5%.

To facilitate comparison with the business cycle model, it is useful to decompose the aggregate stock return into its components. Equation (24) provides a decomposition of the value of this claim into the value of physical capital and patents (intangible capital). Accordingly, we can separately define the return on physical capital, the return on intangible capital, and the spread between the two.

The excess returns on physical capital in the benchmark and business cycle models are radically different. While in the presence of growth cycles the premium on physical capital is in excess of 4%, it is close to zero in the business cycle model and only a tiny fraction of the historical equity premium. In this case, in the presence of growth cycles, agents with a preference for early resolution of uncertainty fear that persistent slowdowns in cash flow growth coincide with a decrease in asset prices. In equilibrium, households thus require a sizeable risk premium on capital.

We suggestively relate the spread between physical and intangible capital to the value premium, which is defined as the return spread between high book-to-market stocks (value stocks) and low book-to-market stocks (growth stocks). Under this interpretation, the benchmark model generates a value spread close to the excess return on the aggregate stock market, with considerable volatility. The link is more suggestive, as growth firms in the data are likely intangibles intensive but also hold physical capital, while in our model they do not. Similarly, in our one-factor economy, a conditional CAPM holds. In other words, value firms in our model have higher expected returns because they have a higher conditional β. In this respect, this rationalization of a value spread follows the arguments in Gomes, Kogan, and Zhang (2003), Carlson, Fisher, and Giammarino (2004) and Zhang (2005).

While sizeable, the premia and volatilities of returns in the model do not rationalize their empirical counterparts entirely. In line with our interpretation of the benchmark model as a model of long-run growth cycles, we view the model-implied risk premia and volatilities as the components reflecting uncertainty about long-term growth prospects and productivity. As discussed earlier, our economy thus does not give a complete account of the relevant short-run risks, which are not likely to be entirely productivity driven. Indeed, Ai, Croce, and Li (2013) report that, empirically, the productivity-driven fraction of return volatility is around 6%, which is roughly consistent with our quantitative finding. On the other hand, the table also reports the asset pricing implications of a version of the endogenous growth model that is calibrated to match short-run consumption risks in a long sample starting from the Great Depression, as is customary in the literature. This calibration produces an overall equity premium of close to 6%, and a value premium of a similar magnitude.

Ultimately, in our model the dynamics of consumption and cash flows reflect endogenous movements in productivity, that is, long-run productivity risks. We now examine and quantify exposure to those risks, and relate them to their determinant, R&D intensity.

Table V

Properties of Consumption Growth and the Stochastic Discount Factor (SDF)

This table reports statistics for consumption growth dynamics (Panel A), expected consumption dynamics (Panel B), and SDF dynamics (Panel C). Panel A reports the annualized volatility and first autocorrelation of consumption growth and the annualized volatility of expected consumption growth for the data, the benchmark growth cycle model, and the business cycle model. In Panel B, we fit the the expected consumption growth process $E_t[\Delta c_{t+1}]$ from the growth cycle and business cycle models to an AR(1) process $x_t = \rho_x x_{t-1} + \sigma_x \epsilon_{x,t}$, where $\epsilon_{x,t} \sim N(0, 1)$, and compare it to the exogenous expected consumption growth component from Bansal and Yaron (2004). We report the persistence parameter ρ_x and the annualized volatility parameter $\tilde{\sigma_x}$ from the fitted AR(1) process, the relative volatility of expected consumption growth and realized consumption growth, and the correlation between expected consumption growth and realized consumption growth. Panel C reports the maximal Sharpe ratio and the one-period mean entropy as defined in Backus, Chernov, and Zin (2014) for the growth cycle and business cycle models. The model statistics correspond to population moments.

Panel A: Consumption Dynamics			
	Data	Growth Cycle	Business Cycle
$\sigma_{\Delta c}$	1.42%	1.42%	1.42%
$AC1(\Delta c)$	0.40	0.39	−0.002
$\sigma(E_t[\Delta c_{t+1}])$	–	0.51%	0.09%
Panel B: Expected Consumption Dynamics			
	Bansal-Yaron	Growth Cycle	Business Cycle
ρ_x	0.979	0.981	0.990
$\tilde{\sigma_x}$	0.12%	0.10%	0.03%
$\sigma(E[\Delta c)])/\sigma(\Delta c)$	0.345	0.359	0.065
corr$(E[\Delta c)], \Delta c)$	0.344	0.526	−0.041
Panel C: SDF Dynamics			
		Growth Cycle	Business Cycle
$\sigma(M)/E(M)$		0.326	0.015
$I(1)$		0.0520	0.0001

B.1. Long-Run Consumption Risks

Table V documents basic properties of consumption growth in the model. Panel A shows that the benchmark model matches the volatility and the annual autocorrelation of consumption growth in the data. While such persistence points to predictable variation in consumption growth in the benchmark model, the table also shows that the conditional mean of consumption growth $E_t[\Delta c_{t+1}]$ is quite volatile.

Panel B shows that this uncertainty about growth prospects in consumption is also very persistent. We quantify the persistence by fitting an AR(1) process to the expected consumption growth process from the benchmark model, in the spirit of Bansal and Yaron (2004). This procedure reveals that expected consumption growth is highly persistent and quite volatile in our model, close to the exogenous parameterization in Bansal and Yaron (2004). In particular, our process is slightly more persistent but slightly less volatile than their specification. In other words, there is a fair amount of persistent uncertainty about growth prospects in consumption, or long-run risks.

Our benchmark model thus exhibits quantitatively significant endogenous long-run risks in consumption. This is in sharp contrast to the companion business cycle specification, which, as noted before, counterfactually exhibits near-i.i.d. consumption growth and therefore minimal time-variation in expected consumption growth. On the other hand, both model specifications exhibit the same amount of short-run risks, as measured by the volatility of realized consumption growth. Note that, in contrast to the exogenous specification in Bansal and Yaron (2004), where innovations to consumption growth and expected consumption growth are uncorrelated, in our model long- and short-run risks are endogenously positively correlated. Bad news for the short run are thus bad news for the long run, reinforcing the endogenous risks in our model.

The presence of endogenous long-run risks in consumption has important implications for the stochastic discount factor, and especially measures of its volatility and dispersion. In Panel C, we report results for two such measures, namely, the maximal Sharpe ratio, $\sigma(M_t)/E(M_t)$, and the mean entropy, $E \log E_t M_{t+1} - E \log M_{t+1}$, following Backus, Chernov, and Zin (2014). Relative to both measures, the benchmark model generates much higher volatility and dispersion than the business cycle specification. In terms of the maximal Sharpe ratio, sometimes referred to as the price of risk, it is instructive to keep the implied value with power utility in mind. In that case, with the calibrated consumption volatility and risk aversion, the Sharpe ratio would be $0.145 \simeq 10 \times 1.45\%$. Deviations from that value reflect the dynamics of expected consumption growth captured by the continuation utility term in Epstein-Zin preferences. In the business cycle model, consumption growth exhibits slight negative autocorrelation, so that innovations to consumption growth and expected consumption growth tend to hedge each other, leading to a low price of risk. The opposite obtains in the benchmark model, resulting in a higher price of risk than the business cycle model. More importantly, the volatility of expected consumption growth is much lower in the real business cycle model.

The presence of long-run risks can also be characterized by inspecting the autocorrelation function for consumption growth and observing persistence. Figure 3 plots the first 10 autocorrelations for annual consumption growth from the benchmark model, the business cycle model, and the data. The benchmark model generates a sizeable first autocorrelation as in the data, but slightly more persistence at longer horizons. Importantly, the consumption dynamics from our model are broadly consistent with the data. In contrast, the business cycle model generates virtually no persistence in consumption growth.

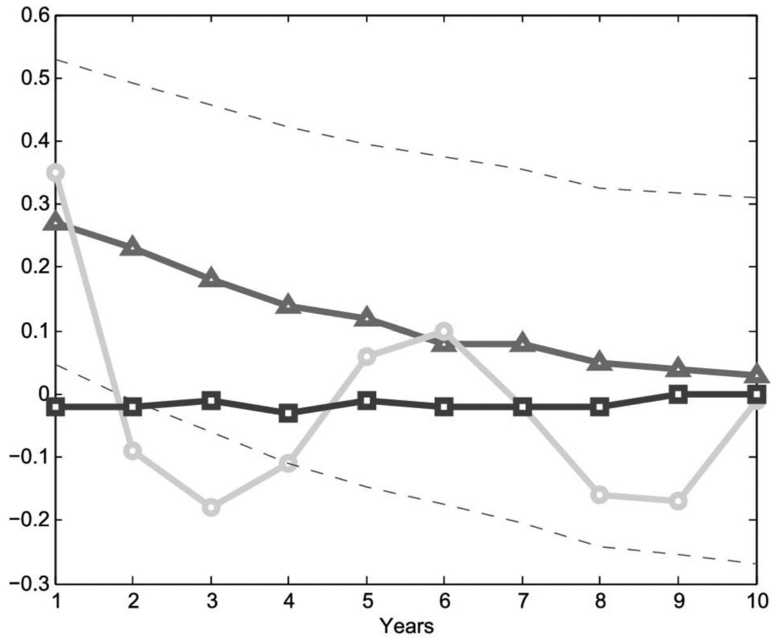

Figure 3. Consumption growth autocorrelations. This figure plots the first 10 autocorrelations of consumption growth. The line with circular markers plots the autocorrelations from the data for the 1953 to 2008 sample period. The line with the triangular markers plots the short-sample autocorrelations from the growth cycle model. The line with the square-like markers plots the short-sample autocorrelations from the business cycle model. From the models, we average across 100 simulations that are equivalent in length to the data sample. The dashed lines represent the lower and upper boundaries of the 95% confidence interval.

B.2. Long-Run Cash Flow Risks

The risk premium on equity is a reflection of both consumption and cash flow risks. Panel A of Table VI documents properties of dividend growth in the model. In line with most of the extant general equilibrium asset pricing literature, our model does not adequately capture the dynamics of stock market dividends obtained from Compustat data.[14] On the other hand, our model rationalizes some of the risks inherent in macroeconomic dividends as measured by the BLS. These dividends are measured as net corporate dividends from both publicly and privately held firms paid out to U.S. investors and arguably form a closer empirical counterpart to the notion of dividends entertained in the model. We provide further details on the data source in Appendix B.

The table reveals that our benchmark model endogenously generates more short-run risk as well as long-run risk in dividend growth than the business cycle specification. Several statistics indicate persistent uncertainty about

[14] In a recent paper, Favilukis and Lin (2013b) make progress on this front by explicitly accounting for wage rigidity, leverage, as well as nonconvexities in investment adjustment costs.

Table VI
Properties of Dividend Growth and the P/D Ratio

This table reports statistics for dividend growth dynamics (Panel A), expected dividend dynamics (Panel B), and price-dividend ratio dynamics (Panel C). Panel A reports the annualized volatility and first autocorrelation of dividend growth and the annualized volatility of expected dividend growth for the data, the benchmark growth cycle model, and the business cycle model. In Panel B, we fit the expected dividend growth process $E_t[\Delta c_{t+1}]$ from the growth cycle and business cycle models to an AR(1) process $x_t = \rho_x x_{t-1} + \sigma_x \epsilon_{x,t}$, where $\epsilon_{x,t} \sim N(0, 1)$, and compare it to the exogenous expected dividend growth component from Bansal and Yaron (2004). We report the persistence parameter ρ_x and the annualized volatility parameter $\tilde{\sigma_x}$ from the fitted AR(1) process, the relative volatility of expected dividend growth and realized dividend growth, and the correlation between expected dividend growth and realized dividend growth. Panel C reports the annualized volatility of the log price-dividend ratio, the annualized volatility of the low-frequency component of the log price-dividend ratio, and the first autocorrelation of the log price-dividend ratio for the data and the growth cycle and business cycle models. The low-frequency component is obtained using the bandpass filter from Christiano and Fitzgerald (2003) and isolating frequencies between 100 and 200 quarters. The model statistics correspond to population moments.

Panel A: Dividend Growth Dynamics			
	Data	Growth Cycle	Business Cycle
$\sigma(\Delta d)$	6.72%	3.21%	2.48%
$\sigma^{GC}(\Delta d)$	0.93%	0.77%	0.48%
$\sigma(E[\Delta d])$	–	0.78%	0.18%
$AC1(\Delta d)$	0.02	0.11	−0.01
Panel B: Expected Dividend Growth Dynamics			
	Bansal-Yaron	Growth Cycle	Business Cycle
ρ	0.979	0.971	0.990
$\tilde{\sigma}_{xd}$	0.36%	0.17%	0.05%
$\sigma(E[\Delta d])/\sigma(\Delta d)$	0.239	0.257	0.072
corr$(E[\Delta d)], \Delta d)$	0.236	0.101	−0.070
Panel C: P/D Ratio Dynamics			
	Data	Growth Cycle	Business Cycle
$\sigma(p-d)$	41.54%	23.26%	13.21%
$\sigma^{GC}(p-d)$	25.86%	7.65%	4.40%
$AC1(p-d)$	0.89	0.90	0.94

dividend growth in the benchmark model. In contrast to the business cycle model, and qualitatively in line with the data, dividend growth exhibits positive autocorrelation. Moreover, the conditional mean of dividend growth is quite volatile, pointing to persistent uncertainty about cash flow growth. Similarly, dividends exhibit fairly volatile growth cycles, as measured by σ^{GC}.

One difficulty that often arises in general equilibrium asset pricing models with production is that they predict dividends to be countercyclical. In these models, companies find it optimal to cut dividends in order to take advantage of productive investment opportunities, so that cash flows effectively end up

Table VII

Properties of Innovation and Productivity

This table reports statistics for variables pertaining to innovation and productivity from the data and the benchmark growth cycle model. Panel A reports volatilities and first autocorrelations of innovation-related measures (R&D expenditure growth Δs, R&D stock growth Δn, and R&D intensity S/N) and productivity growth Δz. Panel B reports the annual persistence and standard deviation of the expected growth rate component of productivity growth. The data estimates are taken from Croce (2014), where the expected growth rate component of productivity $\tilde{x}_{t-1}$ is a latent variable that is assumed to follow an AR(1). In contrast, in the growth cycle model the expected growth rate component is the growth rate of the variety of intermediate goods Δn_t, an endogenous structural variable of the model. The model statistics correspond to population moments.

	Data	Growth Cycle
Panel A: Innovation and Productivity Dynamics		
$\sigma_{\Delta s}$	4.89%	3.82%
$AC1(\Delta s)$	0.21	0.06
$AC1(\Delta n)$	0.90	0.94
$AC1(S/N)$	0.93	0.93
$AC1(\Delta z)$	0.09	0.11
$\sigma(E_t[\Delta z_{t+1}])$		0.38%
Panel B: Expected Productivity Dynamics		
$\rho_{\tilde{x}}$	0.93	0.95
$\sigma(\tilde{x})$	1.10%	1.20%

hedging consumption risks and reduce risk premia on equity. This effect is alleviated in our model due to strongly procyclical aggregate profits $N_t \Pi_t$ on patents. Moreover, in the model, aggregate profit growth exhibits a substantial amount of low-frequency variation itself, with the volatility of its conditional mean equal to 0.42%. Accounting for profits thus helps the benchmark model capture more realistic cash flow risks. However, while the model generates substantial low-frequency variation in profits, it underestimates the total volatility of profits. We view matching cash flow dynamics accurately as an interesting and important extension for future work.

More realistic cash flow dynamics also affect valuation ratios, as documented in Panel C of Table VI. In particular, endogenous long-run risks capture roughly half of the empirical volatility of price-dividend ratios.

B.3. Innovation and Long-Run Productivity Risks

Ultimately, in the benchmark model, long-run risks in both consumption and cash flows reflect innovation-driven movements in endogenous productivity, or, in other words, endogenous long-run productivity risks. As equation (29) above highlights, the significance of this channel depends crucially on the empirical properties of R&D intensity, $\frac{S}{N}$. Table VII documents properties of innovation and productivity in the benchmark model and in the data.

Panel A documents that the model is broadly consistent with volatilities and autocorrelations of R&D investment, the stock of R&D, and R&D intensity in the data. Crucially, in line with its empirical counterpart, R&D intensity is a persistent process and we match its annual autocorrelation of 0.93. Not surprisingly, such persistence is reflected in a positive autocorrelation of productivity growth in the model, close to the value in the data. Moreover, the model predicts considerable uncertainty about future productivity growth as measured by the volatility of its conditional mean.

These results suggest quantitatively significant long-run productivity risks driven by empirically plausible movements in R&D intensity. More formally, uncovering the persistent component in productivity growth as a latent variable in the data (as in Croce (2014)) yields an annual persistence coefficient of 0.93 for the expected growth rate of productivity, while our model closely matches this number with a persistence coefficient of 0.95. Moreover, the volatilities of expected productivity growth rates in the data and in the model roughly match. Note that, in contrast to our benchmark model, the business cycle specification implies that productivity growth is essentially i.i.d., which is inconsistent with empirical evidence.

Qualitatively, the model predicts that R&D intensity should closely track productivity growth. Figure 4 depicts these patterns in the model, using a simulated sample path, as well as in the data. The plots highlight the small but persistent component in productivity growth induced by equilibrium R&D activity. In our general equilibrium model, this persistent component shows up in consumption growth as well, as the figure also illustrates. Empirically, therefore, we expect R&D intensity to forecast productivity and consumption growth. We test this prediction in Section III.G.

B.4. Comovement

Our model also has realistic implications for comovement between prices and quantities at lower frequencies. In the following, we identify low-frequency movements in growth rates using a bandpass filter that isolates movements at frequencies between 100 and 200 quarters.

Figure 5 reveals the close match between the price-dividend ratio and productivity growth in the data and the benchmark model at low frequencies. This evidence strongly suggests the presence of slow productivity-driven movements in asset market valuations in the data. In the model, these movements are driven by variation in expected cash flows, induced by time-variation in R&D intensity. This is because the benchmark model generates little time-variation in risk premia. While there is evidence for time-variation in expected cash flows as discussed above, time-variation in price-dividend ratios is often related to time-varying risk premia. In an extension, we augment our model with stochastic volatility in the exogenous shock and find that it replicates the predictability evidence well.

At lower frequencies we also find strong cross-correlations between stock returns and consumption growth. This is displayed in Figure 6, indicating the

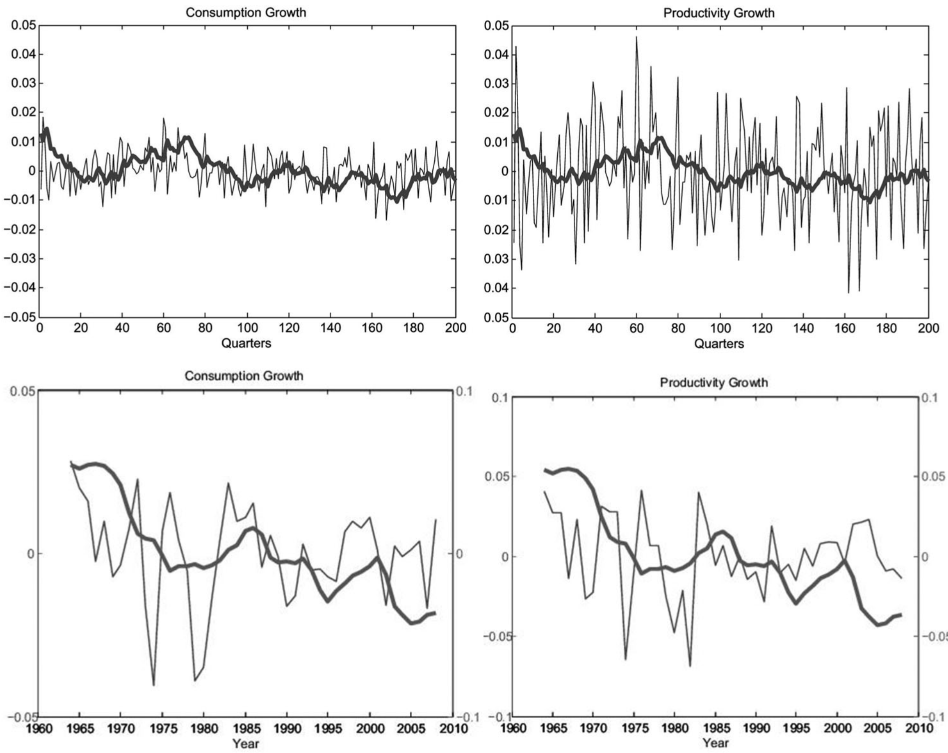

Figure 4. Growth rates and R&D intensity. The top left panel plots demeaned log consumption growth Δc_t (thin line) and R&D intensity $\frac{S_{t-1}}{N_{t-1}}$ (thick line) from the growth cycle model, while the bottom left panel shows the same plot for the data. The top right panel plots demeaned log output growth Δy_t (thin line) and R&D intensity $\frac{S_{t-1}}{N_{t-1}}$ (thick line) from the growth cycle model, while the bottom right panel shows the same plot for the data.

lag-lead structure between returns and consumption growth. In the data and at low frequencies, returns lead consumption growth by several quarters and the lead correlations die away more slowly relative to the lag correlations. In other words, lower-frequency movements in returns contain important information regarding long-run movements in future growth. In contrast to the business cycle specification, the benchmark model replicates this feature of the data quite well. This important divergence between the two models is due to the endogenous predictable component in productivity growth, which is absent in the business cycle model. In sum, the benchmark model is able reconcile to the long-term relationship between returns and growth quite well.

C. Empirical Evidence

Time-varying growth prospects in consumption are at the core of the long-run risks literature following Bansal and Yaron (2004). However, the empirical

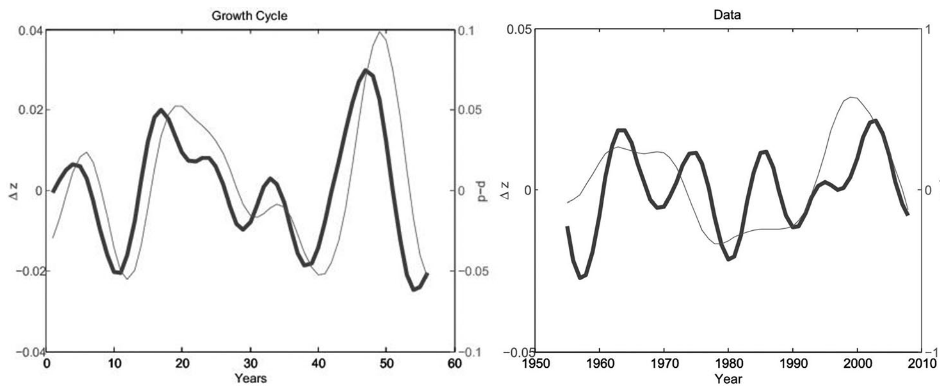

Figure 5. Low-frequency component of productivity growth and price-dividend ratio. This figure plots the low-frequency components for productivity growth (thick line) and for the price-dividend ratio (thin line). The left panel corresponds to a sample simulation from the growth cycle model and the right panel corresponds to the data. The low-frequency component is obtained by applying the bandpass filter from Christiano and Fitzgerald (2003) and selecting a bandwidth of 100 to 200 quarters. The correlation between the two series is 0.46 in the data and 0.67 in the model.

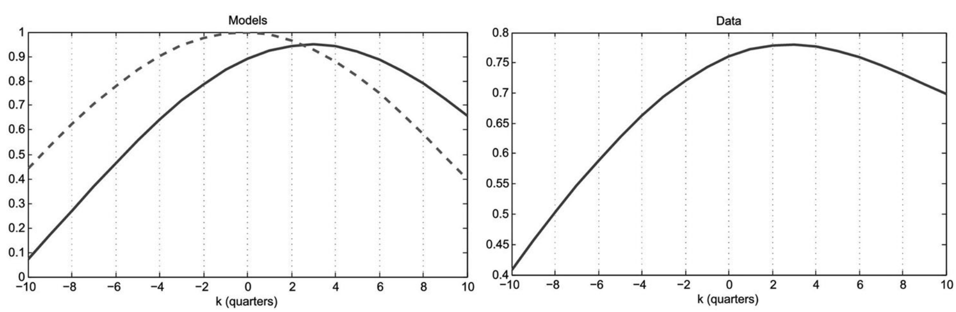

Figure 6. Low-frequency cross-correlation of returns and consumption growth. The left panel plots cross-correlations of the medium-frequency component of the equity return and the low-frequency component of consumption growth for the growth cycle (thick line) and business cycle (dashed line) models: corr($r_{d,t}$, Δc_{t+k}). The right panel plots the same cross-correlations from the data. The low-frequency component is obtained using the bandpass filter from Christiano and Fitzgerald (2003) and selecting a bandwidth of 100 to 200 quarters.

evidence regarding this channel is still controversial. In particular, few instruments have been shown to successfully predict consumption growth over longer horizons. Our benchmark model implies that R&D intensity should predict consumption growth. We now present empirical evidence supporting this channel, based on quarterly data from 1953 to 2008. Innovation-related measures are thus economically meaningful predictors of aggregate growth rates.

Panel A of Table VIII documents the results from projecting future consumption growth for horizons of one to five years on log R&D intensity, both in the data and in model simulations. Empirically, the slope coefficients are positive,

Table VIII
Consumption Growth Forecasts

This table presents short-sample annual consumption growth forecasting regressions from the data and from the benchmark growth cycle model for horizons (k) of one to five years. In Panel A, log consumption growth is projected on log R&D intensity, $\Delta c_{t,t+1} + \cdots + \Delta c_{t+k-1,t+k} = \alpha + \beta(s-n)_t + \nu_{t,t+k}$. In Panel B, log consumption growth is projected on log R&D stock growth, $\Delta c_{t,t+1} + \cdots + \Delta c_{t+k-1,t+k} = \alpha + \beta \Delta n_t + \nu_{t,t+k}$. The regressions are estimated via OLS with Newey-West standard errors with $k-1$ lags and overlapping annual observations. The estimates from the model regression are averaged across 100 simulations that are equivalent in length to the data sample.

	Horizon (Years)				
	1	2	3	4	5
Panel A: Forecasts with R&D Intensity					
β (Data)	0.017	0.034	0.048	0.062	0.077
S.E. (Data)	0.006	0.012	0.017	0.023	0.030
R^2 (Data)	0.070	0.105	0.131	0.163	0.202
β (Growth cycle)	0.068	0.118	0.168	0.200	0.224
S.E. (Growth cycle)	0.028	0.052	0.073	0.095	0.116
R^2 (Growth cycle)	0.141	0.161	0.179	0.175	0.168
Panel B: Forecasts with R&D Growth					
β (Data)	0.217	0.395	0.540	0.703	0.842
S.E. (Data)	0.084	0.178	0.276	0.347	0.401
R^2 (Data)	0.094	0.115	0.132	0.168	0.198
β (Growth cycle)	0.573	1.012	1.437	1.750	1.993
S.E. (Growth cycle)	0.189	0.356	0.526	0.704	0.878
R^2 (Growth cycle)	0.158	0.189	0.207	0.203	0.193

increasing with horizon, and statistically significant. The R^2s are between 0.07 and 0.2, and are monotonically increasing with horizon. Less surprisingly, we find a similar pattern in our model simulations. For completeness, Panel B reports results from projecting future consumption growth on a related measure of innovation, namely, the growth rate of the log R&D stock, $\log N_{t+1} - \log N_t$. A similar pattern obtains. In the data, R&D stock growth forecasts consumption growth over horizons of one to five years with statistically significant and positive slope coefficients and sizeable R^2s, in line with our model. These regressions give empirical support for the notion of innovation-driven low-frequency variation in consumption growth, consistent with the implications of our benchmark model.

In the model, just as changing growth expectations in consumption reflect movements in innovative activity, so too do changing productivity and output growth expectations. Therefore, we expect measures of innovation to forecast productivity and output growth. In Table IX, we provide empirical evidence supporting these predictions. The table documents that in the data both R&D intensity and R&D stock growth forecast productivity and output growth over

Table IX

Output and Productivity Growth Forecasts

This table presents annual output and productivity growth forecasting regressions from the data and from the benchmark growth cycle model for horizons (k) of one to five years. In the top two panels, log output growth is projected on log R&D intensity, $\Delta y_{t,t+1} + \cdots + \Delta y_{t+k-1,t+k} = \alpha + \beta(s-n)_t + \nu_{t,t+k}$ (Panel A) and on log R&D stock growth, $\Delta y_{t,t+1} + \cdots + \Delta y_{t+k-1,t+k} = \alpha + \beta \Delta n_t + \nu_{t,t+k}$ (Panel B). In the bottom two panels, log productivity growth is projected on log R&D intensity, $\Delta z_{t,t+1} + \cdots + \Delta z_{t+k-1,t+k} = \alpha + \beta(s-n)_t + \nu_{t,t+k}$ (Panel C) and on log R&D stock growth, $\Delta z_{t,t+1} + \cdots + \Delta z_{t+k-1,t+k} = \alpha + \beta \Delta n_t + \nu_{t,t+k}$ (Panel D). The regressions are estimated via OLS with Newey-West standard errors with $k-1$ lags and overlapping annual observations. The estimates from the model regression correspond to population estimates.

	Horizon (Years)				
	1	2	3	4	5
Panel A: Output Forecasts with R&D Intensity					
β (Data)	0.020	0.046	0.068	0.089	0.114
S.E. (Data)	0.013	0.022	0.029	0.041	0.051
R^2 (Data)	0.040	0.084	0.119	0.158	0.210
β (Growth cycle)	0.085	0.163	0.236	0.306	0.372
R^2 (Growth cycle)	0.105	0.161	0.195	0.217	0.231
Panel B: Output Forecasts with R&D Growth					
β (Data)	0.267	0.453	0.572	0.763	0.940
S.E. (Data)	0.130	0.261	0.387	0.457	0.499
R^2 (Data)	0.061	0.067	0.073	0.113	0.159
β (Growth cycle)	0.635	1.230	1.780	2.307	2.792
R^2 (Growth cycle)	0.120	0.159	0.193	0.212	0.222
Panel C: Productivity Forecasts with R&D Intensity					
β (Data)	0.014	0.031	0.049	0.069	0.091
S.E. (Data)	0.009	0.015	0.024	0.032	0.041
R^2 (Data)	0.031	0.080	0.120	0.174	0.232
β (Growth cycle)	0.075	0.142	0.204	0.261	0.314
R^2 (Growth cycle)	0.039	0.062	0.077	0.088	0.095
Panel D: Productivity Forecasts with R&D Growth					
β (Data)	0.431	0.820	1.230	1.707	2.092
S.E. (Data)	0.190	0.315	0.452	0.522	0.599
R^2 (Data)	0.113	0.192	0.262	0.376	0.444
β (Growth cycle)	0.560	1.070	1.533	1.948	2.322
R^2 (Growth cycle)	0.037	0.060	0.076	0.084	0.090

several years significantly, with R^2s that are increasing with horizon. Qualitatively, the model replicates these patterns well.

Taken together, we find evidence that aggregate growth rates, including consumption, are indeed time-varying and predictable by innovation-related measures over longer horizons, just as predicted by the benchmark model.

The model thus helps identify economic sources of long-run risks in the data.

IV. Conclusion

This paper provides a quantitative analysis of a production economy whose long-term growth prospects are endogenously determined by innovation and R&D. By integrating innovation and R&D into a real business cycle model with recursive preferences, our model constitutes a straightforward extension of the workhorse model of modern macroeconomics. In sharp contrast to the latter, however, our baseline model jointly rationalizes key features of asset returns and long-run macroeconomic performance in the data.

In the model, favorable economic conditions boost innovation and the development of new technologies. Since technological progress fosters long-run economic growth, endogenous innovation generates a powerful propagation mechanism for shocks reflected in persistent variation in long-term growth prospects. With recursive preferences, innovations to expected growth are priced and lead to high-risk premia in asset markets, as agents fear that persistent slowdowns in growth coincide with low asset valuations. Formally, we show that R&D drives an endogenous predictable component in measured productivity, which gives an innovation-based explanation of long-run productivity risk in the data.

Our model thus allows us to empirically identify economic sources of long-run risks. Indeed, we document novel empirical evidence that measures of innovation have significant predictive power for aggregate growth rates at longer horizons.

Initial submission: April 22, 2012; Final version received: September 1, 2014
Acting Editor: David Backus

Appendix A: Equilibrium

In this section, we collect all the equations that determine the symmetric equilibrium in our economy.

A symmetric equilibrium in the model is defined as an exogenous stochastic sequence, $\{\Omega_t = \exp(a_t)\}_{t=0}^{\infty}$, an initial condition $\{K_0, N_0\}$ for the endogenous state variables, a sequence of endogenous variables, $\{C_t, U_t, M_t, Y_t, W_t, q_t, I_t, \Lambda_t, X_t, \Pi_t, V_t, S_t\}_{t=0}^{\infty}$, and laws of motion $\{K_{t+1}, N_{t+1}\}_{t=0}^{\infty}$ such that

a. the state variables $\{K_t, N_t\}_{t=0}^{\infty}$ satisfy their laws of motions,
b. the endogenous variables solve the producers' and the consumers' problems, and
c. the aggregate resource constraint is satisfied.
d. prices are set such that markets clear.

The equilibrium conditions of the model are summarized by the following 18 equations:

$$
\begin{aligned}
U_t &= \left\{(1-\beta)C_t^{\theta} + \beta\left(E_t\left[U_{t+1}^{1-\gamma}\right]\right)^{\frac{\theta}{1-\gamma}}\right\}^{\frac{1}{\theta}}, \\
M_t &= \beta\left(\frac{C_t}{C_{t-1}}\right)^{\theta-1}\left(\frac{U_t}{E_t(U_t^{1-\gamma})^{\frac{1}{1-\gamma}}}\right)^{1-\gamma-\theta}, \\
Y_t &= K_t^{\alpha} Z_t^{1-\alpha}, \\
Z_t &= \overline{A}e^{a_t}N_t, \\
a_t &= \rho a_{t-1} + \epsilon_t, \\
W_t &= (1-\alpha)(1-\xi)Y_t, \\
q_t &= \frac{1}{\Lambda_t'} \\
1 &= E_t\left[M_{t+1}\left\{\frac{1}{q_t}\left(\alpha(1-\xi)\frac{Y_{t+1}}{K_{t+1}} + q_{t+1}(1-\delta) - \frac{I_{t+1}}{K_{t+1}} + q_{t+1}\Lambda_{t+1}\right)\right\}\right], \\
K_{t+1} &= (1-\delta)K_t + \Lambda_t K_t \\
X_t &= \left(\xi\nu e^{a_t(1-\alpha)(1-\xi)}K_t^{\alpha(1-\xi)}N_t^{\frac{\xi}{\nu}-1}\right)^{\frac{1}{1-\xi}},
\end{aligned}
$$

$$
\begin{aligned}
\Lambda_t &= \frac{\alpha_1}{\zeta}\left(\frac{I_t}{K_t}\right)^{\zeta} + \alpha_2, \\
\Lambda_t' &= \alpha_1\left(\frac{I_t}{K_t}\right)^{\zeta-1}, \\
\Pi_t &= \left(\frac{1}{\nu}-1\right)X_t, \\
V_t &= \Pi_t + (1-\phi)E_t[M_{t+1}V_{t+1}], \\
N_{t+1} &= \vartheta_t S_t + (1-\phi)N_t, \\
\vartheta_t &= \frac{\chi\cdot N_t}{S_t^{1-\eta}N_t^{\eta}}, \\
S_t &= E_t[M_{t+1}V_{t+1}](N_{t+1}-(1-\phi)N_t), \\
C_t &= Y_t - I_t - N_tX_t - S_t.
\end{aligned}
$$

We solve the model in dynare++4.2.1 using a second-order approximation. The policies are centered about a fixed point that takes into account the effects of volatility on decision rules.

Appendix B: Data

Annual and quarterly data for consumption, capital investment, and GDP are from the Bureau of Economic Analysis. Annual data on private business R&D investment are from the survey conducted by The National Science Foundation. Annual data on the stock of private business R&D are from the BLS. Annual productivity data come from the BLS and are measured as multifactor productivity in the private nonfarm business sector. Quarterly data on dividends are obtained from the BLS. The sample period is 1953 to 2008, since R&D data are only available during that time period. Consumption is measured as expenditures on nondurable goods and services. Capital investment is measured as private fixed investment. Output is measured as GDP. Dividends are measured as net corporate dividends. The nominal variables are converted to real terms using the consumer price index (CPI), which is obtained from the Center for Research in Security Prices (CRSP).

Monthly nominal return and yield data are from CRSP. The real market return is constructed by taking the nominal value-weighted return on the New York Stock Exchange and American Stock Exchange and deflating it using the CPI. The real risk-free rate is constructed by using the nominal average one-month yields on Treasury bills and taking out expected inflation.[15] Aggregate market and book values of assets are from the Flow of Funds account. Price-dividend ratio data are from Robert Shiller's webpage:
http://www.econ.yale.edu/~shiller/data.htm.

REFERENCES

Ai, Hengjie, Massimiliano Croce, and Kai Li, 2013, Towards a quantitative equilibrium model of intangible capital, *Review of Financial Studies* 26 491–538.

Alvarez, Fernando, and Urban Jermann, 2005, Using asset prices to measure the persistence of the marginal utility of wealth, *Econometrica* 73, 1977–2016.

Backus, David, Mikhail Chernov, and Stanley Zin, 2014, Sources of entropy in representative agent models, *Journal of Finance* 69, 51–99.

Backus, David, Bryan Routledge, and Stanley Zin, 2007, Asset prices in business cycle analysis, Working paper, New York University.

Backus, David, Bryan Routledge, and Stanley Zin, 2010, Cyclical component of U.S. asset returns, Working paper, New York University.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2013, Risks for the long run: Estimation with time aggregation, Working paper, Duke University.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1639–1672.

Boldrin, Michele, Lawrence Christiano, and Jonas Fisher, 2001, Habit persistence, asset returns and the business cycle, *American Economic Review* 91, 149–166.

Campanale, Claudio, Rui Castro, and Gian Luca Clementi, 2008, Asset pricing in a production economy with Chew-Dekel preferences, *Review of Economic Dynamics* 13, 379–402.

Carlson, Murray, Adlai Fisher, and Ron Giammarino, 2004, Corporate investment and asset price dynamics: Implications for the cross-section of returns, *Journal of Finance* 59, 2577–2603.

Christiano, Lawrence, and Terry Fitzgerald, 2003, The band-pass filter, *International Economic Review* 44, 435–465.

[15] We model the monthly time series process for inflation using an AR(4).

Cogley, Tim, and James Nason, 1995, Output dynamics in real business cycle models, *American Economic Review* 85, 492–511.

Comin, Diego, and Mark Gertler, 2006, Medium term business cycles, *American Economic Review* 96, 523–551.

Comin, Diego, Mark Gertler, and Ana Maria Santacreu, 2009, Technology innovation and diffusion as sources of output and asset price fluctuations, Working paper, Harvard University.

Croce, Massimiliano, 2014, Long-run productivity risk: A new hope for production-based asset pricing, *Journal of Monetary Economics* 66, 13–31.

Dixit, Avinash, and Joseph Stiglitz, 1977, Monopolistic competition and optimum product diversity, *American Economic Review* 67, 297–308.

Ethier, Wilfried, 1982, National and international returns to scale in the modern theory of international trade, *American Economic Review* 72, 389–405.

Favilukis, Jack, and Xiaoji Lin, 2013a, Long run productivity risk and aggregate investment, *Journal of Monetary Economics* 60, 737–751.

Favilukis, Jack, and Xiaoji Lin, 2013b, Wage rigidity: A quantitative solution to several asset pricing puzzles, Working paper, London School of Economics.

Fernandez-Villaverde, Jesus, Ralph Koijen, Juan Rubio-Ramirez, and Jules van Binsbergen, 2012, The term structure of interest rates in a DSGE model with recursive preferences, *Journal of Monetary Economics* 59, 634–648.

Gala, Vito, 2010, Irreversible investment and the cross-section of stock returns in general equilibrium, Working paper, London Business School.

Garleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012, The demographics of innovation and asset returns, *Journal of Financial Economics* 105, 491–510.

Garleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, 2012, Technological growth and asset pricing, *Journal of Finance* 67, 1265–1292.

Gomes, Joao F., Leonid Kogan, and Motohiro Yogo, 2009, Durability of output and expected stock returns, *Journal of Political Economy* 117, 941–986.

Gomes, Joao F., Leonid Kogan, and Lu Zhang, 2003, Equilibrium cross-section of returns, *Journal of Political Economy*, 111, 693–731.

Gordon, Robert J., 2010, Revisiting U.S. productivity growth over the past century with a view of the future, Working paper, Northwestern University.

Gourio, Francois, 2012, Disaster risk and business cycles, *American Economic Review* 102, 2734–2766.

Gourio, Francois, 2013, Credit risk and disaster risk, *American Economic Journal: Macroeconomics* 5, 1–34.

Griliches, Zvi, 1990, Patent statistics as economic indicators: A survey, *Journal of Economic Literature* 28, 1661–1707.

Grossman, Gene, and Elhanan Helpman, 1991, *Innovation and Growth in the Global Economy* MIT Press, Cambridge, MA.

Helpman, Elhanan, 1998, *General Purpose Technologies and Economic Growth* MIT Press, Cambridge, MA.

Jermann, Urban, 1998, Asset pricing in production economies, *Journal of Monetary Economics* 41, 257–275.

Jermann, Urban, and Vincenzo Quadrini, 2007, Stock market boom and the productivity gains of the 1990s, *Journal of Monetary Economics* 54, 413–432.

Jovanovic, Boyan, and Peter Rousseau, 2005 Vol. 1, Part B, General purpose technologies in P. Aghion and S.N. Durlauf, eds.: *Handbook of Economic Growth* (Elsevier, North Holland, Amsterdam).

Kaltenbrunner, Georg, and Lars Lochstoer, 2010, Long run risk through consumption smoothing, *Review of Financial Studies* 23, 3190–3224.

Kogan, Leonid, and Dimitris Papanikolaou, 2010, Growth opportunities and technology shocks, *American Economic Review* 100, 532–536.

Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, 2012, Technological innovation: Winners and losers, Working paper, MIT.

Kuehn, Lars-Alexander, 2008, Asset pricing implications of real investment commitment, Working paper, Carnegie Mellon University.

Kung, Howard, 2015, Macroeconomic linkages between monetory policy and the term structure of interest rates, *Journal of Financial Economics* 115, 41–57.

Lin, Xiaoji, 2012, Endogenous technological progress and the cross-section of stock returns, *Journal of Financial Economics* 103, 411–428.

Papanikolaou, Dimitris, 2011, Investment shocks and asset prices, *Journal of Political Economy* 119, 639–685.

Pástor, Lŭboš, and Pietro Veronesi, 2009, Technological revolutions and stock prices, *American Economic Review* 99, 1451–1483.

Romer, Paul, 1990, Endogenous technological change, *Journal of Political Economy* 98, 71–102.

Tallarini, Thomas, 2000, Risk-sensitive real business cycles, *Journal of Monetary Economics* 45, 507–532.

Zhang, Lu, 2005, The value premium, *Journal of Finance* 60, 67–103.