Technological Growth and Asset Pricing

NICOLAE GÂRLEANU, STAVROS PANAGEAS, and JIANFENG YU[*]

ABSTRACT

We study the asset-pricing implications of technological growth in a model with "small," disembodied productivity shocks and "large," infrequent technological innovations, which are embodied into new capital vintages. The technological-adoption process leads to endogenous cycles in output and asset valuations. This process can help explain stylized asset-valuation patterns around major technological innovations. More importantly, it can help provide a unified, investment-based theory for numerous well-documented facts related to excess-return predictability. To illustrate the distinguishing features of our theory, we highlight novel implications pertaining to the joint time-series properties of consumption and excess returns.

ECONOMIC HISTORIANS FREQUENTLY ASSOCIATE waves of economic activity with the arrival of major technological innovations. The profound changes to manufacturing during the industrial revolution, the expanding network of railroads in the late 19[th] century, electrification, telephony, television, and the Internet during the course of the last century are only a small number of well-known examples of a general pattern whereby a new technology arrives, slowly gets adopted, and eventually permeates and alters all aspects of production and distribution. The impact of technological waves on asset prices is the focus of this paper.

We build a tractable general equilibrium model within which we characterize the behavior of asset prices throughout the technology-adoption cycle. We argue that the model can help provide a unified, investment-based view of numerous facts related to time-series and cross-sectional properties of returns, complementing existing endowment-based approaches. To highlight novel and distinguishing features of our approach, we also derive and test new

implications of the model pertaining to the joint time-series properties of consumption growth and excess returns.

Our main point of departure from previous work on asset pricing is to explicitly allow for the joint presence of two types of technological shocks. Shocks of the first type are assumed to be technology "neutral" or "disembodied," in the sense that they affect the productivity of the entire capital stock irrespective of its vintage. However, such shocks do not fundamentally alter the technology used to produce consumption goods. Shocks of the second type correspond to (infrequent) arrivals of major technological or organizational innovations, like automobiles, the Internet, etc. These shocks do not affect the economy on impact, but only after firms have invested in new vintages of the capital stock that "embody" the technological improvements.

The investment in new capital vintages is assumed to involve a fixed (labor) cost that is irreversible. Firms choose the optimal time to invest in the new capital vintages, which leads to a lag between the arrival of embodied technological shocks and their eventual effects on output and consumption. This process of technological adoption generates endogenous persistence and *investment-driven cycles*, even though all shocks in the model arrive in an unpredictable i.i.d. fashion.

The link between the macroeconomy and asset pricing in our model revolves around the idea that growth options of firms exhibit a "life cycle" as technologies diffuse. On impact of a major technological shock, growth options emerge in the prices of all securities. These growth options are riskier than assets in place, and hence tend to increase the volatility of equity prices and the risk premia in the economy in the initial phases of the technological cycle (i.e., when consumption is below its stochastic trend line). As time passes, firms start to convert growth options into assets in place, which reduces the risk premium on their stock.

We show that the resulting slow and countercyclical movements in expected returns (high expected returns when consumption is below its stochastic trend and vice versa) can help provide a simple, unified explanation of numerous facts related to time variation in returns both in the time series and the cross-section (the value and size premia, dividend-yield predictability, etc.). To highlight the novel aspects of our approach, we focus attention on some new testable model implications that lie at the core of the proposed mechanism. When growth options have not yet been depleted, expected excess returns are high, and so is future consumption growth, as the economy has yet to absorb the gains from a new technology. Accordingly, our theory links *expected* excess returns with future expected consumption growth. We relate this key model prediction to empirical evidence showing an increasing covariance between excess returns and subsequent consumption growth, as the latter is aggregated over longer time intervals. Furthermore, we show empirically that this pattern is driven predominantly by the covariance of *expected* (rather than unexpected) excess returns with future consumption growth. Besides being supportive of the model, this evidence also helps distinguish our investment-based view of

predictability from the leading consumption-based approaches, which—as we show—face challenges in explaining these facts.

The paper is related to several strands of the literature. The paper by Carlson, Fisher, and Giammarino (2004) is the most closely related to ours. Carlson, Fisher, and Giammarino (2004) develop the intuition that the exercise of growth options can lead to variation in expected returns. Given their interest in firm-level decisions, they appropriately choose a partial equilibrium framework. In our paper, the focus is on aggregate quantities and excess returns, which necessitates that we take general equilibrium price-feedback effects into account.[1] Gomes, Kogan, and Zhang (2003) also analyze a general equilibrium production-based model and examine the time-series and cross-sectional properties of returns, as we do. The two most significant differences between their setup and ours are (1) the distinction between embodied and disembodied aggregate technological shocks,[2] and (2) the fact that we allow for an optimal timing decision concerning the exercise of growth options. Because all shocks in Gomes, Kogan, and Zhang (2003) are disembodied productivity shocks, they affect the economy on impact and afterwards their effects dissipate. Our model differentiates between technological shocks that affect the economy on impact (disembodied shocks) and shocks that affect the economy with a lag (embodied shocks). Cycles emerge *endogenously* as the economy responds to an embodied shock, and are directly related to the stock of undepleted growth options and hence expected excess returns. As a result, our model has a distinctive set of implications for the joint time-series properties of returns and macroeconomic aggregates, such as the covariance patterns we document in Section III.B. A final difference, which is of importance for our quantitative exercises, is that in Gomes, Kogan, and Zhang (2003) cycles are driven by an exogenous trend-stationary productivity process. This assumption leads to a trend-stationary consumption process. Our consumption preserves a strong random-walk component, which is a salient feature of consumption in the data.

This paper also relates to the theoretical literature on time variation in expected returns. We do not attempt to summarize this literature here;

---

[1] The need for an explicit, general equilibrium analysis is underscored by the investment literature showing the potentially very different qualitative and quantitative conclusions of real options models once price-feedback effects are included. For instance, Leahy (1993) shows that growth options have zero value in industry equilibrium, once price-feedback effects are taken into account. In our setup, real options have positive value, despite the presence of price-feedback effects. The reason is that firms have access to heterogeneous, nonscalable units of the capital stock. Thomas (2002) illustrates the different quantitative implications of models featuring lumpy investment in partial and general equilibrium.

[2] More generally, the literature on production-based asset-pricing routinely abstracts from this distinction. For contributions to this literature, see Cochrane (1996), Jermann (1998), Berk, Green, and Naik (1999, 2004), Kogan (2004), Kaltenbrunner and Lochstoer (2010), Gomes, Kogan, and Zhang (2003), Carlson, Fisher, and Giammarino (2004), Zhang (2005), Cooper (2006), Gourio (2004), Novy-Marx (2007), and Gala (2006) among others. Papanikolaou (2011) draws a distinction between productivity shocks and investment specific shocks, but does not discuss embodied shocks or different capital vintages. See also the related, intangible capital model of Ai, Croce, and Li (2010).

instead we refer to Cochrane (2005) for an overview. This literature typically uses an endowment-based framework to address variation in expected returns. (See, for example, Campbell and Cochrane (1999) and Bansal and Yaron (2004)). Although, we recognize the value and importance of the channels identified in this literature, our goal in this paper is to highlight the additional and distinguishing implications of an investment-based theory of return variation. To that end, in Section III.B we identify predictions of our theory that are new and distinctive, in the sense that the leading endowment-based approaches face challenges in explaining them. Furthermore, we show how an investment-based approach can help unify more aspects of the evidence on both cross-sectional and time-series predictability than is possible with endowment-based models that abstract from investment.[3,4]

Motivated by the events of the late 1990s, Pástor and Veronesi (2009) connect the arrival of technological growth with the bubble-type behavior of asset prices around the late 1990s.[5] Our model produces similar patterns. However, the focus of the two papers and the mechanisms are different. Our mechanism uses the endogenous exercise of growth options to produce variations in expected returns. Moreover, by considering recurrent arrivals of technological innovations we can discuss implications of the model for the joint stationary distributions of excess returns and macroeconomic aggregates and link technological growth with well documented time-series and cross-sectional patterns of returns.[6]

The structure of the paper is as follows: Section I presents a simplified, partial-equilibrium version of the model with the goal of building intuition. The full model is presented in Section II. Section III presents the empirical implications of the model in a calibrated framework. Section IV concludes. All proofs along with some ancillary results are included in the Internet Appendix.[7]

I. A Simplified Model

In this section, we present a simplified version of the model and discuss its basic intuitions. The next section develops the full model.

---

[3] See, for example, Hsu (2009), Lamont (2000), and Titman, Wei, and Xie (2004), who document various aspects of production-related predictability.

[4] See Santos and Veronesi (2010) for a discussion of the tensions faced by leading endowment-based general equilibrium models in matching simultaneously time-series and cross-sectional aspects of return predictability.

[5] See also Jermann and Quadrini (2007).

[6] There is a large literature in macroeconomics and economic growth that analyzes innovation, dissemination of new technologies, and the impact of the arrival of new capital vintages. A small sample of papers includes Jovanovic and MacDonald (1994), Jovanovic and Rousseau (2003), Greenwood and Jovanovic (1999), Atkeson and Kehoe (1999), and Helpman (1998). In contrast to our paper, this literature concentrates on innovation decisions in a typically deterministic environment, rather than the pricing of risk in a stochastic environment.

[7] An Internet Appendix for this article is available online in the "Supplements and Datasets" section at http://www.afajof.org/supplements.asp.

A. *Trees, Firms, and the Arrival of a Technological Epoch*

There exists a continuum of firms indexed by $j \in [0, 1]$ that produce *consumption* goods. Each firm owns a unit of capital that produces dividends without requiring labor. We refer to each unit of capital as a "tree," adopting the terminology of the seminal Lucas (1978) article. Every existing tree produces a dividend equal to θ_t, where θ_t evolves as a geometric Brownian motion:

$$\frac{d\theta_t}{\theta_t} = \mu dt + \sigma dB_t, \quad \text{with } \mu > 0, \ \sigma > 0. \tag{1}$$

Now suppose that at time $t = 0$ a new technological epoch arrives. This technological epoch gives every firm the opportunity to plant a new tree, which "embodies" the new technological discoveries. In this section we assume that there is only one epoch; we postpone the treatment of multiple epochs to Section II.

We introduce firm heterogeneity by assuming that different firms can plant different trees. Specifically, at $t = 0$ each firm j draws a random number i_j from a uniform distribution on $[0, 1]$. The firm that drew number i_j can plant a new tree that produces dividends equal to

$$Y_{j,t} = \zeta(i_j)\theta_t \tag{2}$$

forever. The first term, $\zeta(i_j)$, is given by a positive, strictly decreasing function $\zeta : [0, 1] \to \mathbb{R}^+$. Accordingly, the smaller i_j, the more productive the tree that firm j can plant.

Any given firm determines the time at which it plants a tree in an optimal manner. Planting a tree at time t requires a fixed cost of q_t. For simplicity, we assume that the company finances these fixed payments by issuing new equity.[8]

We assume throughout that markets are complete, and that the firm's objective is to maximize shareholder value. Accordingly, the optimization problem of firm j amounts to choosing the stopping time τ_j that maximizes

$$P_{j,t}^o \equiv \sup_{\tau_j} E_t \left\{ \zeta(i_j) \int_{\tau_j}^{\infty} \frac{H_s}{H_t}\theta_s ds - \frac{H_{\tau_j}}{H_t}q_{\tau_j} \right\}, \tag{3}$$

where H_s is the stochastic discount factor and $P_{j,t}^o$ denotes the value of the (real) option of planting a new tree. Given the setup, a firm's value $P_{j,t}$ comprises the value of assets in place $P_t^A \equiv E_t \int_t^{\infty} \frac{H_s}{H_t}\theta_s ds$ and the value of the growth option $P_{j,t}^o$. (Naturally, the value of a firm that has already planted a tree is given by the value of its old and new assets in place, $P_{j,t} \equiv [1 + \zeta(i_j)] \times E_t \int_t^{\infty} \frac{H_s}{H_t}\theta_s ds$, although the value of its growth option is zero.)

Moreover, the total output of consumption goods C_t is given by the total output of all firms. Letting $1_{\{t > \tau_j\}}$ denote the indicator function that takes the

---

[8] This assumption is inessential, because the completeness of markets (which we assume shortly) ensures that the Modigliani–Miller theorem holds.

value one if $t > \tau_j$ and zero otherwise, total consumption is given by the output of preexisting and newly planted trees $C_t = \int_0^1 (\theta_t + 1_{\{t > \tau_j\}} Y_{j,t}) dj$. Letting $F(x) \equiv \int_0^x \zeta(j) dj$, and assuming that firms with more productive trees plant their trees first (we verify shortly that this is indeed the case), aggregate consumption can alternatively be expressed as

$$C_t = \theta_t [1 + F(K_t)], \tag{4}$$

where $K_t \in [0, 1]$ is the mass of firms that have planted a tree by time t. Before proceeding, it is useful to note that as $\zeta(\cdot)$ is positive and declining, we obtain $F_x \geq 0$ and $F_{xx} < 0$, so that $F(x)$ has two key properties of a production function, namely, it is increasing and concave.

Equation (4) implies that aggregate consumption is the product of two terms: (i) the nonstationary stochastic trend θ_t, which captures aggregate productivity growth, and (ii) the component $[1 + F(K_t)]$, which captures the contribution of technological adoption to total output.

It is convenient to specify $\zeta(i)$ as a power function, to obtain closed-form solutions. In particular, we let

$$\zeta(i) = bp(1 + bi)^{p-1}, \quad i \in [0, 1], \tag{5}$$

where $b > 0$ and $p \in (0, 1)$ are constants that control the level and curvature of $\zeta(i)$.

B. Stochastic Discount Factor and Cost of Planting Trees

To obtain intuition, it is easiest to start by considering an economy that faces an exogenous stochastic discount factor and costs of planting trees (for instance, a "small" open economy). In the next section we endogenize these quantities by considering a closed economy general equilibrium version of the model.

For now, we simply assume that the stochastic discount factor H_t is given by

$$\frac{dH_t}{H_t} = -r dt - \kappa dB_t, \tag{6}$$

where $r > 0$ is the interest rate and κ captures the market price of risk. Throughout we assume that $r + \kappa\sigma - \mu > 0$, to ensure that the value of the stock market is finite. For parsimony, we also assume that the costs of planting a tree are constant: $q_t = q$.

The following proposition describes firms' optimal investment strategies and the evolution of K_t in equilibrium.

PROPOSITION 1: *Let* $\beta \equiv (r + \kappa\sigma - \mu)^{-1}$, $\phi^+ \equiv (\frac{\kappa}{\sigma} + \frac{1}{2} - \frac{\mu}{\sigma^2}) + \sqrt{(\frac{\kappa}{\sigma} + \frac{1}{2} - \frac{\mu}{\sigma^2})^2 + \frac{2r}{\sigma^2}}$, *and* $\upsilon \equiv \frac{\phi^+}{\phi^+ - 1} \frac{q}{\beta}$. *The optimal stopping time for firm*

j is

$$\tau^* \left(i_j\right) = \inf \left\{ t \geq 0 : \theta_t \geq \frac{\upsilon}{\zeta(i_j)} \right\}, \tag{7}$$

and the resulting process for K_t is given by

$$K_t = \min \left\{ 1, \frac{1}{b} \left(\left(\frac{bp}{\upsilon} \max_{0 \leq s \leq t} \theta_s \right)^{\frac{1}{1-p}} - 1 \right)^+ \right\}, \tag{8}$$

where y^+ is shorthand for $y^+ \equiv \max(y, 0)$.

We make three observations about Proposition 1. First, the optimal firm policies take a threshold form. The firm that can plant the tree with productivity i_j invests once θ_t reaches the critical investment threshold $\frac{\upsilon}{\zeta(i_j)}$. This investment threshold is decreasing in $\zeta(i_j)$. Accordingly, firms with more productive trees invest earlier than firms with less productive trees.

Second, assuming that $\theta_0 < \frac{\upsilon}{\zeta(0)}$, there is some delay between the arrival of the new technology and its adoption by the first firm. However, the continuity of the investment threshold implies that once the first firm adopts, the firms with the next-most productive trees follow in close proximity. This leads to an investment-driven boom and high consumption growth in the initial phases of technological adoption, which slowly decays as the most profitable growth options get depleted.

Third, the variable K_t captures how much the technological cycle has advanced. When K_t is close to zero, the economy is at the early stages of technological adoption with high expected consumption growth going forward. By contrast, a value of K_t close to one indicates a lower expected growth rate. With a view toward later applications, it is useful to formalize this idea by following Beveridge and Nelson (1981) and decomposing log output into a permanent component d_t, defined as

$$d_t \equiv \lim_{T \to \infty} \left\{ E_t \log C_T - \left(\mu - \frac{\sigma^2}{2} \right) (T - t) \right\}, \tag{9}$$

and a cyclical component z_t, defined as $z_t \equiv \log C_t - d_t$. In words, the term d_t captures the level of consumption that will continue to exist in the long run, whereas z_t captures the transitory variation caused by the technological cycle. An immediate implication of definition (9) is that z_t is a predictor of future consumption growth. Specifically, fixing any time interval Δt and letting log-consumption growth be given as $\Delta \log C_{t+i\Delta t} \equiv \log C_{t+(i+1)\Delta t} - \log C_{t+i\Delta t}$, the definition of z_t implies

$$z_t = \log C_t - d_t = -E_t \sum_{i=0}^{\infty} \left(\Delta \log C_{t+i\Delta t} - \left(\mu - \frac{\sigma^2}{2} \right) \Delta t \right). \tag{10}$$

Noting that as $T \to \infty$ all firms will have planted a tree, and using (4), we obtain $d_t = \lim_{T \to \infty} \{E_t \log \theta_T - (\mu - \frac{\sigma^2}{2})(T - t)\} + \log(1 + F(1))$ and hence $d_t = \log \theta_t + \log(1 + F(1))$. Accordingly, z_t is a simple monotone function of K_t given by $z_t = \log(1 + F(K_t)) - \log(1 + F(1))$. Using this expression for z_t together with equation (10) formalizes the idea that low levels of K_t imply high anticipated future consumption growth, whereas higher values of K_t imply low expected consumption growth rates over the future.

C. *Expected Returns over the Technological Cycle*

Next we investigate how changes in the measure of companies that have planted a tree (K_t) affect expected excess returns.

To this end, we first note that, for a fixed θ_t, the share of growth options as a fraction of the aggregate stock market value is a declining function of K_t. Specifically, recalling the definition $P_t^A \equiv E_t \int_t^\infty \frac{H_s}{H_t} \theta_s ds$, we observe that the aggregate value of all assets in place is given by $(1 + F(K_t))P_t^A$, which increases with K_t. The value of all growth options is $P_t^o \equiv \int_{K_t}^1 P_{j,t}^o dj$, which decreases with K_t. Because the aggregate stock market value P_t is the sum of total assets in place and growth options, it follows that the share of growth options $w_t^o \equiv \frac{P_t^o}{P_t}$ is a declining function of K_t. This is intuitive: as the technological cycle progresses, growth options are converted into assets in place and hence account for a progressively smaller share of the aggregate stock market value.

This fact has important implications for the cyclical behavior of expected excess returns. To see why, we use the following lemma.

LEMMA 1:   *The expected excess return of a claim to the aggregate value of growth options is constant and exceeds the expected excess return of assets in place, which is also constant.*

Lemma 1 has an intuitive interpretation. Growth options can be viewed as perpetual American call options on a unit of assets in place. Because the replicating portfolio of a call option involves a levered position in the underlying security, it is immediate that growth options must have a higher expected return than assets in place.

Combining the facts that (i) the fraction of growth options in the aggregate stock market value is declining in K_t, and (ii) the difference between the expected returns on growth options and on assets in place is positive and constant implies that market expected excess returns are declining in K_t.

As we have discussed earlier, the technological cycle (z_t) is an increasing function of K_t. Hence, expected excess returns are countercyclical in the sense that they are high at the early stages of the technological cycle (low values of z_t) and low at the advanced stages of the technological cycle (high values of z_t). In light of (10), this implies that when expected returns are high, expected future consumption growth should be high as well and vice versa. In

Section III.B we investigate empirically this link between expected excess returns and subsequent consumption growth, and contrast the performance of our model with other models of countercyclical returns.

To generalize and quantify the qualitative insights of the present section, we now develop a stationary general equilibrium model.

II. The Complete Model

We generalize the simplified model in two ways. First, we close the model by introducing consumer–workers and deriving the stochastic discount factor and costs of planting a tree in general equilibrium. Second, we introduce multiple epochs, so that technological growth has stationary (rather than transient) effects.

A. Firms, Investment, and Aggregation

Similar to the simplified model, there exists a continuum of firms indexed by $j \in [0, 1]$, which produce *consumption* goods. However, because we now allow for multiple technological epochs, we assume that each firm owns a collection of trees that have been planted in different technological epochs. In analogy to equation (2), the output of a tree is given by

$$Y_{n,i,t} = \zeta(i)\theta_t A_n. \tag{11}$$

As in the simplified model, $\zeta(i)$ captures a tree-specific component and $\zeta : [0, 1] \to \mathbb{R}^+$ is a positive and strictly decreasing function. The random variable θ_t is the common productivity shock and has the same dynamics (equation (1)) as in the simplified model. Compared to equation (2), specification (11) includes the additional component A_n, where $n \in (-\infty, \dots, 0, 1, 2 \dots)$ denotes the technological epoch during which the tree was planted. Accordingly, A_n can be interpreted as a "vintage-specific" effect, because it is common to all trees that are planted in epoch n. We make two assumptions about the evolution of A_n. First, $A_{n+1} \geq A_n$, so that vintages of trees planted in epoch $n + 1$ are more productive than their predecessors (all else equal). Second, the ratio A_{n+1}/A_n is increasing in the extent of technological adoption that took place in epoch n. Specifically, letting $K_{n,t} \in [0, 1]$ denote the mass of trees planted in epoch n by time t, and τ_{n+1} the time of arrival of epoch $n + 1$, we postulate the following dynamics for A_{n+1}:

$$A_{n+1} = A_n \left(1 + \int_0^{K_{n,\tau_{n+1}}} \zeta(i)\,di\right). \tag{12}$$

Equation (12) reflects a standard assumption in the endogenous growth theory that is sometimes referred to as "standing on the shoulders of giants." The act of planting new trees produces knowledge and stimulates further

innovation in future epochs.[9] Thus, the increase between A_n and A_{n+1} depends on the investment activity in period n.

Technological epochs arrive exogenously at the Poisson rate $\lambda > 0$. Throughout, we denote the arrival time of epoch n as τ_n. Once a new epoch arrives, the index n becomes $n + 1$, and every firm gains the option to plant a single tree of the new vintage at a time of its choosing.

Firm heterogeneity is introduced in a manner identical to the simplified model: once epoch n arrives, each firm j draws a random number $i_{n,j}$ from a uniform distribution on $[0, 1]$, and this number informs the firm that it can plant a tree with tree-specific productivity $\zeta(i_{n,j})$ in epoch n. These numbers are drawn in an i.i.d. fashion across epochs. To simplify the setup, we assume that once an epoch changes, the firm loses the option to plant a tree that corresponds to any previous epoch;[10] it can only plant a tree corresponding to the technology of the current epoch.

As in the simplified model, any firm determines the time at which it plants a tree in an optimal manner, and planting a tree at time t requires a fixed cost of q_t. Because the productivity index $i_{n,j}$ is i.i.d. across epochs, there is no linkage between the decision to plant a tree in this epoch and any future epochs. Accordingly, every firm in epoch N solves a stopping problem analogous to equation (3) of the simplified model:

$$P^o_{N,j,t} \equiv \sup_{\tau_j} E_t \left\{ 1_{\{\tau_j < \tau_{N+1}\}} \left[\left(A_N \zeta(i_{N,j}) \int_{\tau_j}^{\infty} \frac{H_s}{H_t} \theta_s ds \right) - \frac{H_{\tau_j}}{H_t} q_{\tau_j} \right] \right\}.$$

Letting N denote the technological epoch at time t and $1_{\{\widetilde{\chi}_{n,j}=1\}}$ be an indicator function equal to one if firm j has already planted a tree in technological epoch n and zero otherwise, the overall value of a firm consists of three components: (i) the value of assets in place $P^A_{j,t} \equiv (\sum_{n<N} A_n \zeta(i_{n,j}) 1_{\{\widetilde{\chi}_{n,j}=1\}})(E_t \int_t^{\infty} \frac{H_s}{H_t} \theta_s ds)$, (ii) the value of the growth option in the current technological epoch $P^o_{N,j,t}$, and (iii) the value of the growth options in all subsequent epochs $P^f_{N,t} \equiv E_t(\sum_{n=N+1}^{\infty} \frac{H_{\tau_n}}{H_t} P^o_{n,j,\tau_n})$.

Aggregate consumption is given by $C_t \equiv \int_0^1 Y_{j,t} dj$. Assuming that firms with more productive trees plant their trees first (we verify that this is indeed the case in equilibrium), we obtain the following lemma.

---

[9] For some background on the "standing on the shoulders of giants" assumption, see, for example, the seminal paper of Romer (1990), who assumes that the arrival rate of new blueprints depends on the level of past blueprints. For a textbook treatment, see Jones (1997). Scotchmer (1991) contains a number of concrete examples of the positive effects of past innovation on new innovation ranging from the cotton gin to techniques of inserting genes into bacteria.

[10] As we explain in the Appendix (Section B.3), the assumption that a firm can plant a tree corresponding only to the current epoch can be relaxed if we modify equation (12) to $A_{n+1} = A_n \bar{A}(1 + \int_0^{K_{n,\tau_{n+1}}} \zeta(i)di)$, where $\bar{A} \geq \frac{\zeta(0)}{\zeta(1)}$. However, this extension adds complexity without additional insights, and we avoid it for parsimony.

LEMMA 2: *Assuming that there is a positive probability that $K_{n,t} > 0$ for all n, C_t is given by*

$$C_t = A_N \theta_t [1 + F(K_{N,t})] \tag{13}$$

with probability one.

The aggregate consumption equation (13) is reminiscent of the expression (4) in the simplified model, with two important differences: (1) θ_t is replaced by $A_N \theta_t$, and (2) $K_{N,t}$ is a process that gets reset to zero every time a new epoch arrives. As a result, $K_{N,t}$ introduces recurrent (rather than purely transient), cyclical components into consumption.

B. Consumer–Workers and Preferences

To endogeneize the stochastic discount factor and the costs of planting trees, we next assume that the economy is populated by a continuum of identical consumer–workers that can be aggregated into a single representative agent. The representative agent owns all the firms in the economy, and is also the (competitive) provider of labor services. Purely for simplicity, we assume that work is not directly useful in the production of *consumption* goods but is useful in the production of *investment* goods, that is, trees.

Specifically, following the external habit specification of Abel (1999), we assume that the consumer's utility depends on both her own consumption and her consumption relative to some benchmark level. To facilitate closed-form solutions, we take the benchmark level to be the running maximum of aggregate consumption, $M_t^C \equiv \max_{s \le t} \{C_s\}$, and specify the consumer's instantaneous utility as

$$U\left(c_t, M_t^C\right) = \frac{1}{1-\gamma} \left[\left(\frac{c_t}{M_t^C}\right)^{1-\alpha} c_t^\alpha \right]^{1-\gamma}, \quad \alpha \in [0, 1], \tag{14}$$

where c_t is the agent's own consumption and γ denotes risk aversion. This utility specification nests the commonly used constant-relative-risk-aversion preferences (when $\alpha = 1$) and the preferences considered by Abel (1999) and Chan and Kogan (2002) (when $\alpha = 0$) as special cases. The presence of external habit formation is useful for calibration, because it (i) allows one to match the low level of interest rates, without sacrificing the high equity premium, in the data, and (ii) mitigates the reaction of interest rates to an acceleration of anticipated consumption growth caused by the arrival of a new technological epoch. (As we show in Section B.4 of the Internet Appendix, external habit formation helps achieve these two goals by effectively increasing the intertemporal elasticity of substitution (IES) in a growing economy.) We also note that, unlike the specification in Campbell and Cochrane (1999), specification (14) implies constant relative risk aversion. Even though a certain degree of time-varying risk

aversion could be introduced into our framework,[11] the property of constant relative risk aversion in specification (14) helps us illustrate more clearly the new economic mechanisms that drive our results.

Besides deriving utility from consumption, the representative agent also derives disutility from providing labor services that are necessary for the production of new trees. Specifically, we assume that the representative agent maximizes

$$V_t = \max_{c_s, dl_s} E_t \left[\int_t^\infty e^{-\rho(s-t)} U(c_s, M_s^C) ds - \int_t^\infty e^{-\rho(s-t)} \eta(s) dl_s \right], \quad (15)$$

where ρ is the subjective discount factor, $\eta(s)$ captures the disutility associated with planting an additional tree, and $dl(s) \geq 0$ denotes the increments in the number of trees that the representative agent plants. Letting q_s denote the wage for planting a new tree at time s, and noting that the representative consumer is the owner of all firms and the recipient of all wage payments, the consumer's intertemporal budget constraint is

$$E_t \left(\int_t^\infty \frac{H_s}{H_t} c_s ds \right) \leq \int_0^1 P_{j,t} dj + E_t \left(\int_t^\infty \frac{H_s}{H_t} q_s dl_s \right). \quad (16)$$

Because markets are complete, the consumer's maximization problem amounts to maximizing (15) subject to (16).

To complete the presentation of the model, we need to make functional form assumptions about η_t. To motivate our choice of η_t, we start with a few observations about the relationship between η_t and the equilibrium *reservation wage* to plant a tree q_t. Letting V_W denote the derivative of the value function with respect to wealth, an agent has an incentive to plant a tree if and only if

$$q_t V_W \geq \eta_t. \quad (17)$$

The left-hand side of (17) is the increment in the agent's value function from receiving the wage q_t, whereas the right-hand side is the associated utility cost. The envelope condition of dynamic programming (see, e.g., Øksendal (2003), chapter 11) implies $V_W = U_c$. Using $V_W = U_c$ inside (17) implies that an additional tree is planted only as long as $q_t \geq \frac{\eta_t}{U_c}$. Because there is a continuum of workers, perfect competition among them drives the price of planting a tree to $q_t = \frac{\eta_t}{U_c}$.[12] We choose η_t so as to make the model consistent with some salient facts in the data. First, we want to ensure that labor income is cointegrated with total output, as in the data. Because $q_t = \frac{\eta_t}{U_c}$, it is therefore necessary that

---

[11] This could be done by either building time-varying risk aversion in the preferences of the representative agent as in Campbell and Cochrane (1999) or by assuming investor heterogeneity as in Chan and Kogan (2002) or Gârleanu and Panageas (2007).

[12] Rogerson (1988) shows the existence of a representative agent in models with indivisible labor, assuming the existence of "labor lotteries" as part of the tradeable contingent claims. We discuss this issue further in the Appendix (section B.1), along with alternative ways to achieve the same goal.

$\frac{\eta_t}{U_c C_t}$ is stationary. Furthermore, it seems plausible to specify η_t so that the cost of planting a tree q_t is growing between epochs (reflecting the complexities of creating more advanced units of the capital stock). Finally, for tractability, we choose the disutility η_t of planting a tree so that it is independent of the number of trees ($l_{N,t}$) already planted by the representative agent in the current epoch. Motivated by these requirements, we specify[13]

$$\eta_t = \eta U_C M_{\tau_N}^C (1 + F(l_{N,t}))^v, \tag{18}$$

where $v \equiv \gamma - (\gamma - 1)(1 - \alpha)$.

C. Equilibrium

Subject to some technical assumptions, Proposition 2 in the Internet Appendix shows that there exists a constant $v^* > 0$ such that firm j in round N finds it optimal to plant a tree at time $\tau_{j,N}^*$ given by

$$\tau_{j,N}^* = \inf_{\tau_N \leq t < \tau_{N+1}} \left\{ t : \frac{\theta_t}{M_{\tau_N}} \geq \frac{v^*(1 + F(i_{N,j}))^v}{\zeta(i_{N,j})} \right\}, \tag{19}$$

where $M_t \equiv \max_{s \leq t} \theta_s$. The optimal investment policy in the complete model (equation (19)) has the same properties as the investment policy in the simplified model. Specifically, firms with more productive trees in the current epoch always plant them earlier than firms with less productive trees. A further implication of policy (19) is that no firm finds it optimal to plant a tree when $t = \tau_N$ (i.e., right at the beginning of the epoch), as long as[14]

$$v^* > \zeta(0), \tag{20}$$

which we shall assume throughout.

Letting $m_t \equiv M_t / M_{\tau_N}$ denote the ratio of the current running maximum of the common productivity process θ_t to the value of the same quantity at the beginning of the current epoch, and aggregating over the investment policies of equation (19), leads to

$$K_{N,t} = K(m_t) = \min \left\{ \left(\frac{1}{b} \left(\frac{bp}{v^*} m_t \right)^{\frac{1}{1-p+vp}} - \frac{1}{b} \right)^+, 1 \right\}. \tag{21}$$

The fact that m_t is a stationary quantity (as it gets reset to one every time a new epoch arrives) implies that $K_{N,t}$ is stationary, as is x_t defined by

$$x_t \equiv \log \left(1 + F\left(K_{N,t} \right) \right). \tag{22}$$

---

[13] Note that $l_{N,t} = K_{N,t}$, $M_t^C = A_N(\max_{s \leq t} \theta_s)(1 + F(l_{N,t}))$, (13), and (14) imply that $U_C(1 + F(l_{N,t}))^v$, and thus η_t, are independent of $l_{N,t}$.

[14] This condition is sufficient to induce waiting because of (19) and $\frac{\theta_t}{M_{\tau_N}} = \frac{\theta_{\tau_N}}{M_{\tau_N}} \leq 1$ at the beginning of epoch N. Hence, all firms (even the most productive one) are "below" their investment thresholds.

Log consumption (log C_t),
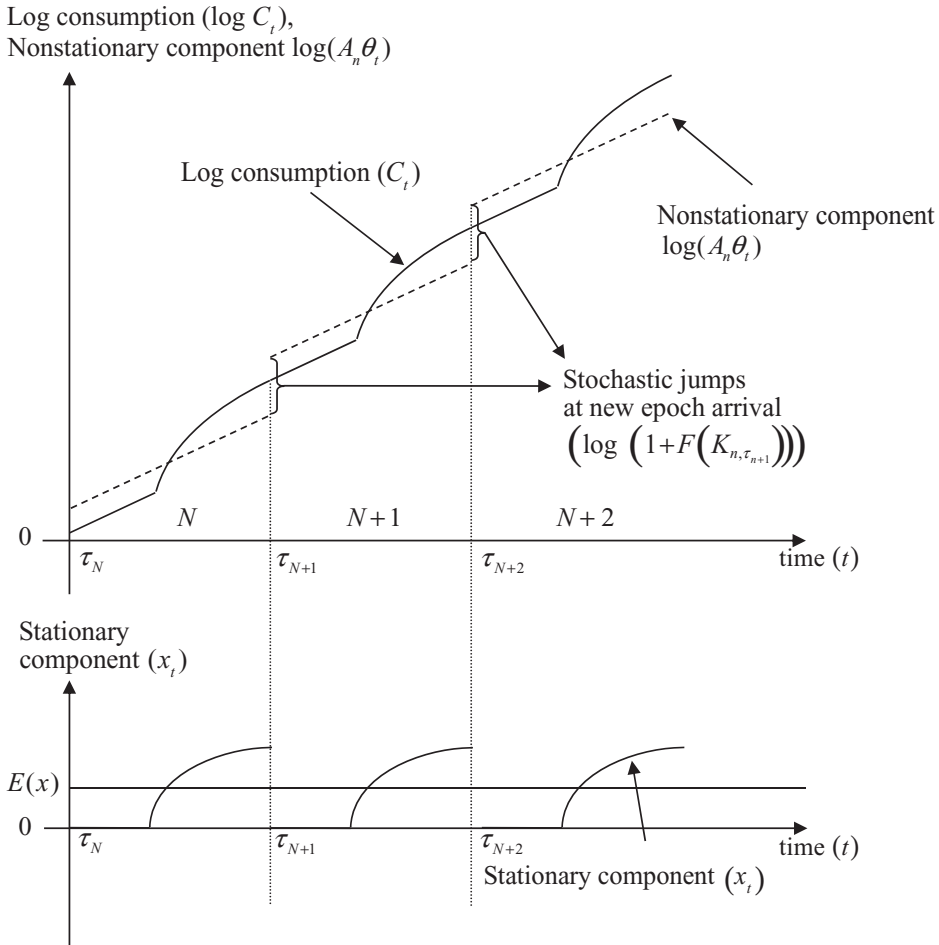Nonstationary component $\log(A_n\theta_t)$

Figure 1. Decomposition of log consumption. Log consumption $\log(C_t)$ (top figure, solid line) is decomposed into the nonstationary part $\log(A_N) + \log(\theta_t)$ (top figure, dashed line) and the stationary part x_t (bottom figure). The figure plots consumption for a path with $dB_t = 0$ and $K_{N,\tau_{N+1}} > 0$.

Taking logs on both sides of equation (13) gives

$$\log(C_t) = \log(A_N\theta_t) + x_t. \tag{23}$$

Equation (23) implies that aggregate log consumption can be decomposed into a stochastic nonstationary component $\log(A_N\theta_t)$ and an *endogenous investment-driven stationary component* given by x_t.

Figure 1 illustrates the decomposition of log consumption into its components. The figure illustrates how the arrival of a new epoch makes A_N jump upwards, consistent with equation (12). In the short run, this jump in the

stochastic trend line is not reflected in the level of consumption, because consumption itself does not jump. However, as time passes and firms start to invest, the consumption growth rate first increases as the most profitable firms exploit their investment opportunities, and then slowly decays thereafter. At some point a new epoch arrives and this pattern repeats itself.

Figure 1 implies that the model is broadly consistent with historical episodes of general purpose technological adoption.[15] Specifically, the onset of a new technological epoch leads to increased expectations of future consumption growth. In the short run, consumption growth is low because firms do not invest in new technologies initially. Interest rates rise, reflecting increased anticipation of consumption growth in the future, and aggregate expected excess returns also increase, reflecting the increased importance of growth options (see Section III.C). Higher interest rates lead to lower valuations of existing assets in place. In turn, this leads to lower valuations in firms that have little to benefit from the current technology (i.e., the firms that have drawn the least productive trees). However, valuations increase for the firms that can benefit the most, as their current-period growth options offset the effect of higher interest rates.[16] As technology starts to get adopted, the economy experiences an investment-driven boom; expected excess returns decline, reflecting the depletion of growth options; and interest rates decline, reflecting lower growth expectations. Cross-sectionally, the decline in expected returns is sharpest for firms that can benefit the most from the technology, as the fraction of their assets in place increases sharply at the time of investment.

In the next section we go beyond the qualitative evaluation of the model's behavior around specific historical episodes. Specifically, we investigate quantitatively the joint properties of macroeconomic and financial quantities as implied by the reaction of the economy to the recurrent arrival of new technologies.

III. Quantitative Implications for Expected Returns

A. Calibration

Table I presents our choice of the nine parameters for the baseline calibration exercise. Three parameters are related to the distribution of the exogenous shocks (μ, σ, and λ), four parameters pertain to preferences (ρ, γ, α, and η), and two parameters (p and b) control the function $\zeta(i)$, that is, the degree of heterogeneity across trees that can be planted in a given epoch. We choose μ to match the contribution of (neutral) total factor productivity to annual

---

[15] Jovanovic and Rousseau (2005) document various patterns for adoption, output growth, investment, valuations, etc., related to the adoption of electricity (1920s) and information technologies (1990s).

[16] Jovanovic and Rousseau (2005) document lower valuations for firms that had little to benefit from the new technologies. One channel that seems to lead to lower valuations of existing firms in the data and is not captured in the current model (for parsimony) is the increased competition created by the new firms adopting the new technologies. See Gârleanu, Kogan, and Panageas (2009) for a model that incorporates such a feature.

Table I
Parameters Used for the Calibration

Growth rate of total factor productivity (μ)	0.012
Instantaneous consumption volatility (σ)	0.03
New epoch arrival rate (λ)	0.1
Risk aversion (γ)	9
Subjective discount rate (ρ)	0.012
Parameter controlling external habits (α)	0.1
Parameters controlling heterogeneity of trees ($[b, p]$)	[0.8,0.6]
Parameter controlling labor supply (η)	22.6

aggregate growth. Hulten (1992) computes that number to be 1.17%, which motivates our choice of $\mu = 0.012$. The parameter σ controls the volatility of consumption growth. We set it to $\sigma = 0.03$ to match the volatility of annual time-integrated consumption data.[17]

The parameters λ, p, b, and η control the growth contribution of the quality and quantity increase in trees (capital goods), the speed of new tree adoption, and the time variation in consumption growth rates. We follow Comin and Gertler (2006), who estimate the frequency of technology-driven "medium-run" cycles, and set $\lambda = 0.1$. With this choice we also want to highlight that we think of our model as capturing medium-run fluctuations, rather than fluctuations associated with the business cycle. The parameters b and η control (respectively) the contribution of new capital vintages to aggregate growth per epoch and the time it takes until firms start planting trees. As a result, these parameters control the total consumption growth rate and the cyclical effects of technology adoption. We choose these parameters to approximately match (i) the total annual consumption growth rate in the data and (ii) the autocorrelation properties of consumption. Finally, the parameter p controls the curvature of the function $\zeta(i)$ and hence the acceleration in consumption growth once firms start adopting new technologies. To measure this acceleration in growth because of adoption of a new technology, we use the difference in annual consumption growth rates between 1980 and 1994 and between 1995 and 2000, which is about 1.1%. We choose p to approximately match such a difference in growth rates between the initial stages of the epoch (when no firm invests) and the latter stages of the epoch (when firms start investing).

In terms of the preference parameters ρ, γ, and α, we choose ρ and α so as to (i) match the low level of real interest rates in the data and (ii) obtain plausible degrees of the IES. Specifically, as we show in Section B.4 of the Internet Appendix, $\gamma + (\gamma - 1)(\alpha - 1)$ provides a measure of the inverse of the IES in a deterministically growing economy. With $\alpha = 0.1$, the implied IES with respect to such shocks is about 0.55, well within the reasonable range of values estimated in the literature (see, e.g., Attanasio and Vissing-Jorgensen (2003)).

---

[17] As is well understood, time integration makes the volatility of time-integrated consumption data lower than the instantaneous volatility of consumption.

Table II

Unconditional Moments in the Model and the Data
(Annualized Rates)

All data are from the long sample (1871–2005) in Campbell and Cochrane (1999), with the exception of the volatility of the 1-year zero coupon yield, which is from Chan and Kogan (2002). The unconditional moments for the model are computed from a Monte Carlo simulation involving 12,000 years of data, dropping the initial 1,000 years to ensure that initial quantities are drawn from their stationary distribution. The time increment dt is chosen to be 1/60.

	Data	Model
Mean of consumption growth	0.017	0.016
Volatility of consumption growth	0.033	0.027
Mean of 1-year zero coupon yield	0.029	0.010
Volatility of 1-year zero coupon yield	0.030	0.060
Mean of equity premium (logarithmic returns)	0.039	0.041
Volatility of equity premium	0.180	0.176

Finally, we choose $\gamma = 9$, which is sufficient to match the average equity premium.

Table II compares the model's performance to some unconditional moments in the data. To save space, we report some additional investment-related statistics in the Internet Appendix (Section B.6). The overall performance of the model in terms of unconditional asset-pricing moments is comparable to the pure endowment models of external habit formation, such as Abel (1990) and Chan and Kogan (2002).

Figure 2 plots the consumption autocorrelations implied by the model and in the data. Consistent with the data, the model-implied autocorrelations are small and decay rapidly. The intuition for this finding is that only a small fraction of the variability of consumption comes from the cyclical component x_t. We note in passing that the model produces similar autocorrelations to the data irrespective of whether we time-aggregate consumption at quarterly or annual frequencies.

B. Time-Series Predictability

In this section we evaluate the quantitative ability of the model to explain some well-documented predictability patterns in the data. More importantly, we discuss and test some new implications of our investment-based theory of predictability, and compare our theory to existing consumption-based approaches.

We start the discussion of the time-series properties of returns by performing the usual predictability regressions of aggregate excess returns on the aggregate log price-to-dividend ratio. Table III tabulates the results of these regressions, and compares them to the data. Because of well documented small-sample issues in return-predictability regressions, we simulate
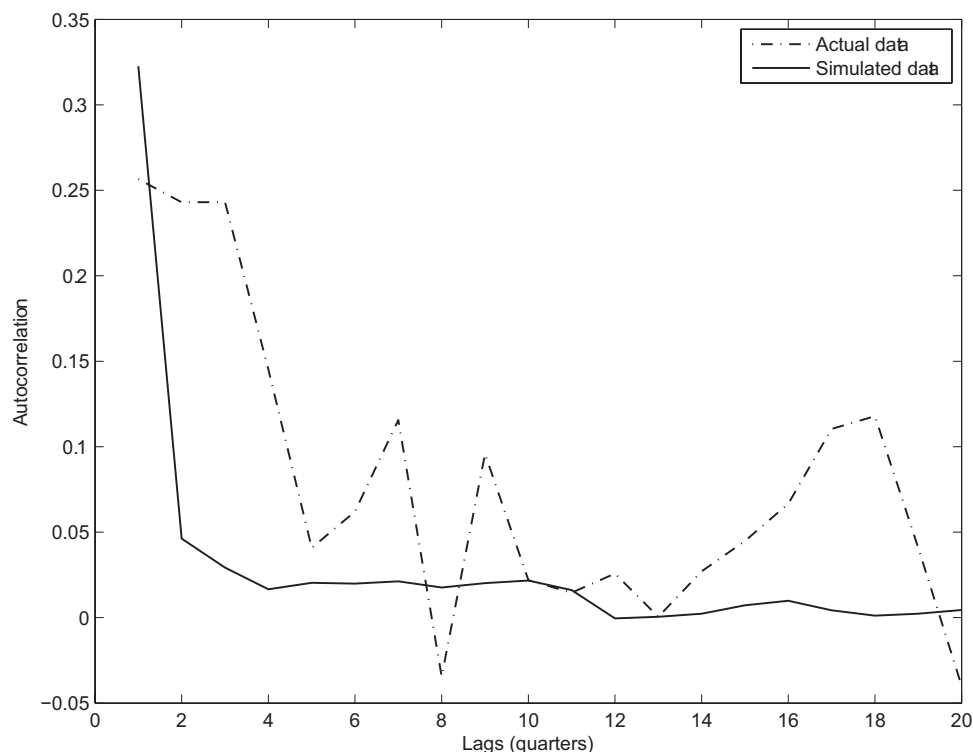
Figure 2. Autocorrelations of consumption growth. Autocorrelations of quarterly seasonally adjusted log-consumption growth of nondurables (except clothes and shoes) and services. Source: Bureau of Economic Analysis, 1947Q1–2009Q4. The dashed line refers to the data, whereas the solid line is obtained from simulating the model and time-aggregating consumption data.

1,000 independent samples of 100-year-long paths of artificial data. We run predictability regressions for each of these samples and report the average co-efficient along with a 95% distribution band. The coefficients in the simulations have the right sign but are about one-third of their empirical counterparts. Most of the empirical point estimates, however, are within the 95% distribution band according to the model.

In interpreting these results, we note that the price-to-dividend ratio—conditional on a given level of habit C_t/M_t^C—is increasing (rather than de-creasing) in K_t. The depletion of growth options associated with a higher value of K_t has two effects: it lowers the anticipated cash flow growth, and also the interest rates. Given an IES below one, the latter effect dominates and the price-to-dividend ratio increases.

Needless to say, our theory of predictability is not the first to account for the patterns documented in Table III. We therefore devote the remainder of this section to deriving some novel implications of our investment-based theory and comparing them with the data.

Table III
Results of Predictive Regressions

Excess returns in the aggregate stock market between t and $t + T$ for $T = 1, 2, 3, 5$, and 7 are regressed on the aggregate price-to-dividend ratio at time t. A constant is included but not reported. The data column is from Chan and Kogan (2002). The simulations are performed by drawing 1,000 time series of a length equal to 100 years. We report the means of these simulations next to the respective point estimates in the data. The numbers in parentheses are the 95% confidence interval of the estimates obtained in the simulations.

	Data		Model	
Horizon(years)	Coefficient	R^2	Coefficient	R^2
1	-0.120	0.040	-0.051	0.005
			$(-0.363, 0.049)$	$(0.000, 0.112)$
2	-0.300	0.100	-0.105	0.013
			$(-0.431, 0.108)$	$(0.000, 0.100)$
3	-0.350	0.110	-0.156	0.017
			$(-0.500, 0.152)$	$(0.000, 0.114)$
5	-0.640	0.230	-0.208	0.022
			$(-0.720, 0.225)$	$(0.000, 0.172)$
7	-0.730	0.250	-0.247	0.023
			$(-0.894, 0.279)$	$(0.000, 0.240)$

What lies at the core of our investment-based theory of predictability is that the variation in the amount of growth options implies a positive covariation between expected excess returns and subsequent anticipated consumption growth over the long run.[18] Figure 3 shows such a pattern in the data. The figure reports the covariance between the market excess return R_{t+1}^e from t to $t + 1$ and consumption growth cumulated from $t + 2$ to $t + 2 + T$, where $T = 1, 2, 3 \ldots$ quarters.[19] We compute 95% confidence bands by employing the bootstrap method under the null hypothesis that consumption growth and returns are contemporaneously correlated but i.i.d. over time. Figure 3 shows that, both in the data and in the model, the covariances increase with T.[20]

In the Internet Appendix we investigate whether such a pattern can be explained by two of the leading consumption-based models in the literature,

[18] We discuss this feature carefully at the end of Section III.C, and the mechanism is robust to the determination of the pricing kernel in general equilibrium and the arrival of future epochs.

[19] Because our theory is about the covariance between current excess returns and "long-run," rather than contemporaneous, consumption growth, we cumulate consumption growth from time $t + 2$ onward, so as to ensure that our results are not affected by contemporaneous covariances and time-aggregation in consumption data. In addition, we drop consumption growth between dates t and $t + 2$ because time integration in consumption implies that the same θ_t shocks that affect R_{t+1}^e affect consumption growth not only between t and $t + 1$, but also between $t + 1$ and $t + 2$. Hence, to fully exclude any mechanical, contemporaneous covariance, we compute cumulative consumption growth from date $t + 2$ onward.

[20] Daniel and Marshall (1999), Piazzesi (2002), and Parker and Julliard (2005) obtain similar results.
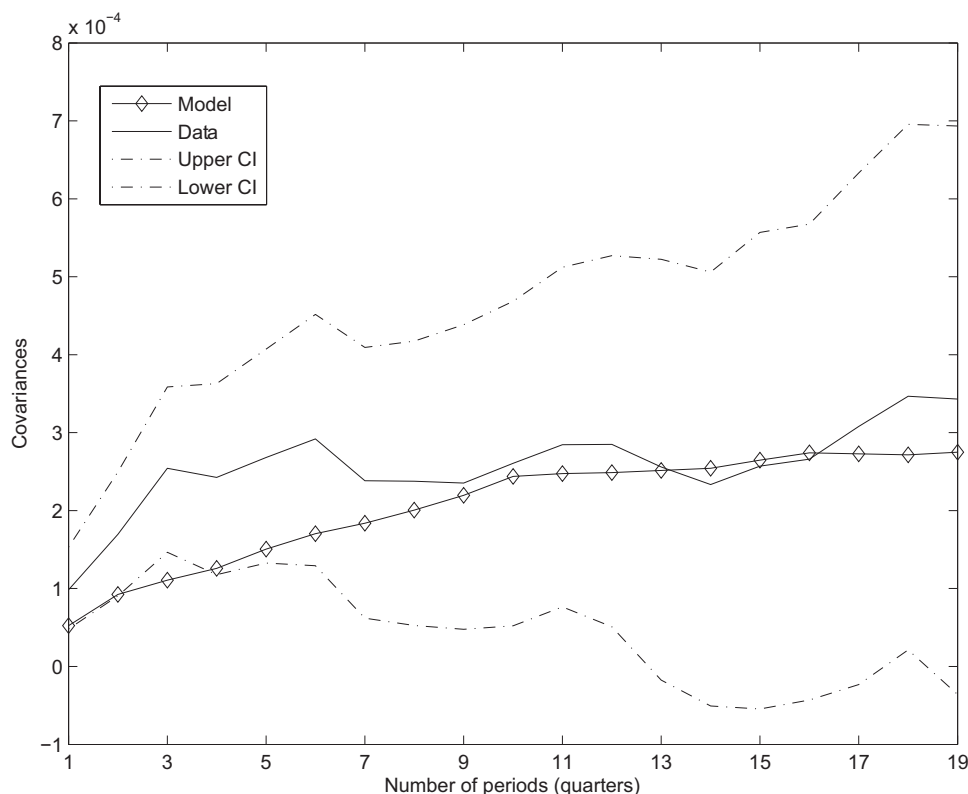
Figure 3. Covariances between excess return and consumption growth. Covariances between the excess return R_{t+1}^e and cumulative consumption growth between dates $t+2$ and $t+2+T$ for $T = 1, 2, 3 \ldots$. The line denoted "Data" corresponds to the point estimate in the data. The lines labeled "Upper CI" and "Lower CI" refer to the 2.5% and 97.5% confidence bounds, computed from bootstrap samples, which are "drawn" (with replacement) under the null hypothesis that excess returns and consumption growth are i.i.d. over time. Finally, the line "Model" refers to the respective covariances in model simulations.

namely, the models of Campbell and Cochrane (1999) and Bansal and Yaron (2004). To save space, here we simply summarize our findings; we refer the reader to the Internet Appendix (Section C) for details. The consumption-based model of Campbell and Cochrane (1999) cannot explain an increasing pattern of covariances, because log-consumption growth is i.i.d. and hence $\mathrm{cov}(R_{t+1}^e, \Delta \log c_{t+2+T}) = 0$ for $T \geq 1$. In the Internet Appendix we show that even if we introduce predictable consumption growth in such a model (say, by introducing first-order autocorrelation in expected consumption growth, as in Bansal and Yaron (2004)), the covariance between the current excess return and long-run consumption growth is downward-sloping rather than upward-sloping in T. We also show that the model of Bansal and Yaron (2004) is consistent with an upward-sloping pattern of covariance. However,

there is an important difference in the source of this upward-sloping pattern in our paper and Bansal and Yaron (2004). Specifically, for a given $T \geq 1$ we can decompose the covariance between excess returns and cumulative consumption growth as

$$\text{cov}\left(\sum_{i=t+3}^{t+2+T} \Delta \log c_i, R_{t+1}^e\right) = \text{cov}\left(\sum_{i=t+3}^{t+2+T} \Delta \log c_i, E_t R_{t+1}^e\right)$$
$$+ \text{cov}\left(\sum_{i=t+3}^{t+2+T} \Delta \log c_i, R_{t+1}^e - E_t R_{t+1}^e\right). \quad (24)$$

In words, the covariance between the cumulative consumption growth and the excess return R_{t+1}^e comprises a first component because of the *expected* excess return and a second component because of the *innovations* to the excess return. In the model of Bansal and Yaron (2004), time variation in expected returns is independent of time variation in expected consumption growth,[21] and accordingly the first component of (24) is zero.

Table IV shows that (1) the first term in equation (24) is statistically different from zero and (2) especially at longer horizons, the first term accounts for the majority of covariance between returns and subsequent "long-run" consumption growth. Specifically, Table IV presents the results of two sets of regressions (in Panels A and B), both of which regress cumulative consumption growth between times $t + 2$ and $t + 2 + T$ on time-t expected excess returns. To determine expected excess returns, in Panel A we regress realized excess returns on a set of instruments (known at the beginning of period t) that have been shown in the literature to be good predictors of excess returns (term premium, default premium, interest rate, inflation rate, and dividend yield). According to the model of Bansal and Yaron (2004), this estimated expected return should only reflect stochastic volatility, which in their model is unrelated to subsequent long-run consumption growth. Hence, the regressions in Table IV should result in zero slope coefficients. By contrast, in our model these regressions should result in nonzero slope coefficients, as the expected returns predict subsequent consumption growth. Statistically, the hypothesis of a zero slope coefficient can be rejected. Moreover, as the row "Exp. component" shows, the fraction of the covariance between excess returns and subsequent consumption growth that is because of the covariance between *expected* excess returns and subsequent consumption growth is economically significant, especially at longer horizons. For instance, this fraction is 85%–90% at the 5-year horizon.[22] For robustness,

---

[21] Specifically, Bansal and Yaron (2004) use a first-order approximation to solve their model, and find that the expected component of consumption growth does not affect the expected excess return to the first order. (They also confirm their conclusions by solving the model exactly.) To obtain time variation in expected excess returns, Bansal and Yaron (2004) introduce stochastic volatility.

[22] In the Appendix (Section D), we perform a similar covariance decomposition inside the model. Similar to the data, the fraction of the covariance because of the expected component of returns is economically significant inside the model, especially at longer horizons. For instance, at the 5-year horizon this fraction is slightly above 50%.

<div align="center">

Table IV

Consumption Growth and (Lagged) Expected Returns

</div>

This table reports the coefficients, t-statistics (both Newey and West (1987) and Hodrick (1992)), and adjusted R^2s for the regression of cumulative consumption growth on (lagged) expected returns, $c_{t+2+T} - c_{t+2} = \alpha + \beta E_t[R^e_{t+1}] + \epsilon_{t+1,T}$. In the row "Exp. Component" we report the fraction of the covariance between excess returns and consumption growth because of the expected component of excess returns. Expected returns are estimated by regressing realized quarterly excess returns onto a set of instruments. In Panel A, the instruments are the term premium, default premium, interest rate, inflation rate, and dividend yield. In Panel B, the instrument is the output gap, defined as in Cooper and Priestley (2009). The default premium is defined as the yield spread between *BAA* and *AAA* bonds and the term premium is defined as the difference between the 20-year Treasury bond yield and the 1-year Treasury yield. Source: Federal Reserve Bank of St. Louis.

Horizon (T)	1	2	3	7	11	15	19
				Panel A			
β	0.064	0.137	0.214	0.458	0.547	0.566	0.655
t-stat (Newey-West)	3.149	3.634	3.964	4.195	3.680	2.850	2.728
t-stat (Hodrick)	3.277	3.519	3.652	3.567	3.682	2.948	3.011
Adjusted R^2	0.066	0.123	0.166	0.229	0.181	0.134	0.140
Exp. Component	0.296	0.367	0.379	0.878	0.864	0.974	0.853
				Panel B			
β	0.056	0.130	0.209	0.435	0.608	0.703	0.865
t-stat (Newey-West)	2.611	3.000	3.138	2.977	3.021	3.256	3.815
t-stat (Hodrick)	3.174	3.634	3.895	3.931	4.323	3.972	4.019
Adjusted R^2	0.038	0.083	0.120	0.157	0.175	0.168	0.199
Exp. Component	0.199	0.266	0.284	0.639	0.752	0.971	0.906

Panel B recalculates this fraction using as an instrument for predicting excess returns only the output gap of Cooper and Priestley (2009), which is a predictor more closely related to our theory.[23]

Parenthetically, we note that the model of Bansal and Yaron (2004) can be extended to account for the evidence of Table IV if one were to assume that stochastic volatility in consumption can predict subsequent long-run consumption growth. (See, e.g., Backus, Routledge, and Zin (2008)). However, it can be challenging to detect the empirical evidence behind this assumption. For instance, Boguth and Kuehn (2010) use filtering methods to infer time

[23] Daniel and Marshall (1999) and Parker and Julliard (2005) reach related conclusions. Daniel and Marshall (1999) show that removing the predictable components from one-quarter returns and consumption growth reduces significantly their business-cycle-frequency correlations. Our analysis does not focus on low-frequency correlations, but rather on the covariance between current excess returns and subsequent consumption growth. The patterns of this covariance rely not only on low-frequency comovements between the two series, but more importantly on the fact that expected excess returns lead consumption growth. Parker and Julliard (2005) also show predictability of consumption growth at longer horizons, but using cross-sectional (the Fama–French size and value factors SMB, respectively HML) rather than time-series instruments. Neither paper performs a covariance decomposition along the lines of equation (24), which is important for our purposes.

variation in the first two moments of consumption growth and find no evidence of a link between long-run growth and stochastic volatility in consumption. An advantage of our theory is that it does not require such a link to explain the patterns documented in Table IV.

We conclude this section by noting that, besides explaining some salient patterns of consumption and excess returns in the time-series dimension, our investment-based theory also helps explain why certain macroeconomic variables are helpful in predicting returns. Particularly related to our theory is the evidence in Cooper and Priestley (2009), Lamont (2000), and Hsu (2009) that the output gap, investment plans, and technological innovations, respectively, can help predict excess returns.

C. Cross-Sectional Implications

The cross-sectional differences in the fraction of assets in place and growth options across firms, along with the idiosyncratic variation in the productivity of trees planted in different epochs, have implications for the cross-sectional properties of returns. This section explains why the model is consistent with well-documented cross-sectional patterns of returns, such as the size and value premia. We also highlight some additional, investment-related, cross-sectional implications of the model.

To see why the model is able to produce a size premium, it is easiest to consider a firm j that has a higher market value of equity (size) than firm j', so that $P_{N,j,t} > P_{N,j',t}$. To simplify the analysis, assume further that both of these firms have exercised their current-epoch growth option, so that $P_{N,j,t}^o = P_{N,j',t}^o = 0$. Because the future growth options are the same for both firms, the relative importance of growth options for firm j must be smaller, and hence firm j must therefore have a lower expected return. Hence, abstracting from current-epoch growth options,[24] a sorting of companies based on size would produce a size premium.

The model is also consistent with the value premium. This may seem counterintuitive at first, because one would expect that firms with a high market-to-book ratio should have a substantial fraction of their value tied up in growth options and hence should be riskier. The resolution of the puzzle is linked to the fact that assets in place are heterogeneous. The easiest way to see how tree heterogeneity helps account for the value premium is to consider two firms j and j' that have planted a tree in every single epoch, including the current one. As a result, the two firms have identical book values and identical growth options. However, suppose that firm j has always been "luckier" than firm j' in

---

[24] The presence of current epoch growth options distorts the perfect ranking of expected returns implied by size. Intuitively, high market values may be associated with a valuable current-period growth option (in which case expected returns should be high) instead of numerous assets in place (in which case expected returns should be low). For the calibrations that we consider, however, we find that current-epoch growth options are not quantitatively important enough to affect the size effect.

terms of the productivity of the trees it has planted in the past. Then the market value of firm j is higher than the market value of firm j'. Accordingly, firm j has a lower book-to-market ratio than firm j' and a lower expected return (as the fraction of its market value that is because of assets in place is higher).[25] Hence, similar to Gomes, Kogan, and Zhang (2003), our model produces a value premium despite the presence of risky growth options.

Even though not at the core of our analysis, we note that the model is also consistent with several additional cross-sectional properties of the data. Thus, because high-size (and high-growth) firms typically have trees with higher productivity on average, the model is consistent with the empirical evidence reported in Fama and French (1995) that sorting on size and value produces predictability for a firm's profitability (earnings-to-book ratio). Because the relative size of firms is mean-reverting in the model, the model is also consistent with the evidence that small firms tend to grow faster than large firms. Furthermore, as the exercise of growth options leads to lower expected returns, the model is consistent with the empirical evidence in Titman, Wei, and Xie (2004), who show that increased investment activity at the firm level leads to lower subsequent returns. Finally, consistent with the data, the model predicts that firms with a low book-to-market ratio (high Tobin's Q) tend to exhibit stronger investment activity (as measured by the growth in the book value of assets).[26]

Quantitatively, Table V reports results on the cross-sectional predictability of returns. To match more accurately the cross-sectional distribution of size and book-to-market dispersion, we introduce idiosyncratic (disembodied) tree-specific shocks. To save space, we motivate and give a detailed description of these shocks in the Internet Appendix (Section B.5). Here, we simply note that we construct these shocks so that they do not affect a firm's optimal stopping problem, the stochastic discount factor, or any aggregate quantity. They simply add more variability to the stationary cross-sectional distribution of the size and book-to-market ratios, so as to allow us to match these distributions more accurately.

Table V shows that returns sorted by book-to-market and size replicate the qualitative patterns in the data. The magnitudes, however, are smaller.

We conclude this section by highlighting a limitation of the model in terms of explaining the cross-section of returns. In simulations, sorting on the size effect drives out the value effect, and vice versa. This is linked to the fact that within the model only one source of risk is reflected in the stochastic

---

[25] The presence of current-period growth options distorts the ranking of expected returns implied by the above argument. As we show below, in a calibrated version of the model this distortion is not powerful enough to substantially affect the value effect.

[26] The intuition for this fact is simple: A high Tobin's Q (low book-to-market) reflects (i) the productivity of existing trees, but also (ii) the magnitude of growth options compared to the current capital stock of the firm. The first component drives expected returns down as we showed above, but is irrelevant for predicting the growth rate in the capital stock. However, the second component predicts the growth in the capital stock. The interplay of these two forces can help explain the joint presence of a value premium and a weak positive correlation between Tobin's Q and the investment-to-capital ratio. See, for example, the evidence in Abel and Eberly (2002).

<div align="center">

Table V

Portfolios Sorted by Size and Book-to-Market—Model and Data

</div>

The data are from the Web site of Kenneth French. Time period: 1927–2009. Average returns per decile are based on monthly data, which are converted to annualized rates. We subtract 3.09% from all returns to account for the average CPI inflation between 1927 and 2009. The median (log) firm size is normalized to zero.

Portfolios Formed on Size (Stationary Distribution)										
Deciles	1	2	3	4	5	6	7	8	9	10
log(Size) – Data	−2.45	−1.36	−0.82	−0.40	0.00	0.40	0.82	1.34	1.99	3.51
log(Size) – Sim.	−2.04	−1.28	−0.79	−0.39	0.00	0.38	0.77	1.16	1.72	3.61
Returns(Size) – Data	13.91	11.72	11.63	11.07	10.53	10.44	9.88	9.13	8.53	7.00
Returns(Size) – Sim.	7.96	6.55	5.79	5.72	5.79	5.85	6.00	6.17	5.77	5.63

Portfolios Formed on Book-to-Market (Stationary Distribution)										
Deciles	1	2	3	4	5	6	7	8	9	10
log(BM) – Data	−1.47	−0.88	−0.59	−0.38	−0.20	−0.03	0.14	0.34	0.60	1.22
log(BM) – Sim.	−2.94	−1.68	−0.98	−0.47	−0.04	0.32	0.66	1.03	1.52	2.50
Returns(BM) – Data	6.65	7.86	8.05	7.73	8.53	9.01	9.21	11.05	12.00	12.67
Returns(BM) – Sim.	5.71	5.77	5.77	5.87	6.04	6.28	6.19	6.15	6.30	7.17

discount factor. Therefore, as long as one of the two sorting procedures leads to a satisfactory ranking of the conditional betas, the other sorting procedure adds little.[27] Gârleanu, Kogan, and Panageas (2009) propose a model in which the stochastic discount factor rewards multiple sources of risk because of a lack of intergenerational risk sharing and rivalry between technological innovations. Within such a model value and size premia can be obtained as independent effects, but such an extension is beyond the scope of the current paper.

IV. Conclusion

We propose a model of technological change that posits, in addition to the usual small, embodied shocks, major disembodied ones that affect output only following new investment. Because it takes a while for the investment in the new technologies to become viable and thus translate into higher output, during the early stages of the adoption cycle consumption growth is low. At the same time, the presence of relatively numerous real options leads to high expected returns. The pattern reverses once investment increases the growth rate of consumption and the ratio between the values of assets in place and growth options.

The model's implications are consistent with the stylized facts concerning macroeconomic aggregates and asset valuations surrounding major technological innovations. In addition, we find that this simple, investment-driven theory

---

[27] In this connection we also note that (unconditional) market betas cannot explain the dispersion in excess returns, because they do not generate sufficient conditional-beta variation.

of discount-rate countercyclicality can account for several well-documented predictability relationships in the time series and the cross-section of returns (e.g., dividend-yield predictability and size and value premia). Furthermore, our theory has novel implications for the joint time-series properties of consumption and excess returns. We focus on the following core implication: *expected* excess returns should exhibit positive covariation with subsequent consumption growth. The reason is that when growth options are relatively abundant (depleted), both the expected consumption growth and the expected excess returns are higher (lower). We provide empirical evidence supporting such a positive covariation, and argue that the leading endowment-based approaches face limitations in terms of explaining some aspects of this covariation.

REFERENCES

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38–42.

Abel, Andrew B., 1999, Risk premia and term premia in general equilibrium, *Journal of Monetary Economics* 43, 3–33.

Abel, Andrew B., and Janice C. Eberly, 2002, Investment and q with fixed costs: An empirical analysis, Working paper, Wharton and Kellogg.

Ai, Hengjie, Mariano M. Croce, and Kai Li, 2010, Toward a quantitative general equilibrium asset pricing model with intangible capital, Working paper, Duke and UNC.

Atkeson, Andrew, and Patrick J. Kehoe, 1999, Models of energy use: Putty-putty versus putty-clay, *American Economic Review* 89, 1028–1043.

Attanasio, Orazio P., and Annette Vissing-Jorgensen, 2003, Stock market participation, intertemporal substitution and risk aversion, *American Economic Review, Papers and Proceedings* 93, 383–391.

Backus, David K., Bryan R. Routledge, and Stanley E. Zin, 2008, The cyclical component of U.S. asset returns, Working paper, NYU Stern and Tepper School of Business.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 2004, Valuation and return dynamics of new ventures, *Review of Financial Studies* 17, 1–35.

Beveridge, Stephen, and Charles R. Nelson, 1981, A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the "business cycle", *Journal of Monetary Economics* 7, 151–174.

Boguth, Oliver, and Lars Kuehn, 2010, Consumption volatility risk, Working paper, Carnegie Mellon and Arizona State.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

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.

Chan, Yeung L., and Leonid Kogan, 2002, Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices, *Journal of Political Economy* 110, 1255–1285.

Cochrane, John H., 1996, A cross-sectional test of an investment-based asset pricing model, *Journal of Political Economy* 104, 572–621.

Cochrane, John H., 2005, *Asset Pricing* (Princeton University Press, Princeton and Oxford).

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

Cooper, Ilan, 2006, Asset pricing implications of non-convex adjustment costs and irreversibility of investment, *Journal of Finance* 61, 139–170.

Cooper, Ilan, and Richard Priestley, 2009, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22, 2801–2833.

Daniel, Kent D., and David A. Marshall, 1999, Consumption based modeling of long horizon returns, Working paper, Northwestern University.

Fama, Eugene F., and Kenneth R. French, 1995, Size and book-to-market factors in earnings and returns, *Journal of Finance* 50, 131–155.

Gala, Vito D., 2006, *Investment and Returns*, PhD Dissertation, University of Chicago GSB.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2009, Displacement risk and asset returns, *Journal of Financial Economics* (forthcoming).

Gârleanu, Nicolae, and Stavros Panageas, 2007, Young, old, conservative and bold: The implications of heterogeneity and finite lives for asset pricing, Working paper, Berkeley Haas and Chicago Booth.

Gomes, Joao, Leonid Kogan, and Lu Zhang, 2003, Equilibrium cross section of returns, *Journal of Political Economy* 111, 693–732.

Gourio, Francois, 2004, Operating leverage, stock market cyclicality, and the cross-section of returns, Working paper, Chicago GSB.

Greenwood, Jeremy, and Boyan Jovanovic, 1999, The information-technology revolution and the stock market, *American Economic Review* 89, 116–122.

Helpman, Elhanan, 1998, *General Purpose Technologies and Economic Growth* (MIT Press, Cambridge and London).

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.

Hsu, Po-Hsuan, 2009, Technological innovations and aggregate risk premiums, *Journal of Financial Economics* 94, 264–279.

Hulten, Charles R., 1992, Growth accounting when technical change is embodied in capital, *American Economic Review* 82, 964–980.

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.

Jones, Charles I., 1997, *Introduction to Economic Growth*. (W.W. Norton, New York and London).

Jovanovic, Boyan, and Glenn M. MacDonald, 1994, The life cycle of a competitive industry, *Journal of Political Economy* 102, 322–347.

Jovanovic, Boyan, and Peter L. Rousseau, 2003, Two technological revolutions, *Journal of the European Economic Association* 1, 419–428.

Jovanovic, Boyan, and Peter L. Rousseau, 2005, General purpose technologies, in Philippe Aghion, and Steven N. Durlauf, eds.: *Handbook of Economic Growth*, Volume 1B (Elsevier B.V., Amsterdam).

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

Kogan, Leonid, 2004, Asset prices and real investment, *Journal of Financial Economics* 73, 411–431.

Lamont, Owen A., 2000, Investment plans and stock returns, *Journal of Finance* 55, 2719–2745.

Leahy, John V., 1993, Investment in competitive equilibrium: The optimality of myopic behavior, *Quarterly Journal of Economics* 108, 1105–1133.

Lucas, Robert E., Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Novy-Marx, Robert, 2007, An equilibrium model of investment under uncertainty, *Review of Financial Studies* 20, 1461–1502.

Øksendal, Bernt, 2003, *Stochastic Differential Equations* (Springer-Verlag, Berlin).

Papanikolaou, Dimitris, 2011, Investment shocks and asset prices, *Journal of Political Economy*, 119, 630–685.

Parker, Jonathan A., and Christian Julliard, 2005, Consumption risk and the cross section of expected returns, *Journal of Political Economy* 113, 185–222.

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

Piazzesi, Monika, 2002, The 6d bias and the equity-premium puzzle: Comment, in B. S. Bernanke and K. Rogoff, eds.: *NBER Macroeconomics Annual 2001*, Volume 16 (MIT Press, Cambridge and London).

Rogerson, Richard, 1988, Indivisible labor, lotteries and equilibrium, *Journal of Monetary Economics* 21, 3–16.

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

Santos, Tano, and Pietro Veronesi, 2010, Habit formation, the cross section of stock returns and the cash-flow risk puzzle, *Journal of Financial Economics* 98, 385–413.

Scotchmer, Suzanne, 1991, Standing on the shoulders of giants: Cumulative research and the patent law, *The Journal of Economic Perspectives* 5, 29–41.

Thomas, Julia K., 2002, Is lumpy investment relevant for the business cycle?, *Journal of Political Economy* 110, 508–534.

Titman, Sheridan, John K. Wei, and Feixue Xie, 2004, Capital investments and stock returns, *Journal of Financial and Quantitative Analysis* 39, 677–700.

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