Is the IT revolution over? An asset pricing view

Colin Ward*

November 16, 2015

Abstract

I develop a method that structures financial market data to forecast the IT sector's remaining growth contribution to the US economy. Future average annual productivity growth is predicted to fall to 52bps from the 87bps recorded over 1974–2012, due to intensifying IT sector competition and decreasing returns to employing IT. My median estimate's two standard-error band indicates the transition ends between 2032 and 2038. I estimate these numbers with a model that links economy-wide growth to the IT sector's market valuation, and by calibrating it to match historical data on factor shares, price-dividend ratios, growth rates, and discount rates.

*Carlson School of Management, University of Minnesota. I thank Andy Abel, João Gomes, Christian Opp, and Nick Roussanov for their advice and Vincent Glode for many fruitful conversations. I also thank Jules van Binsbergen, Nicolae Gârleanu, Gunnar Grass, Brent Neiman, Cecilia Parlatore, Chris Parsons, Ivan Shaliastovich, Ali Shourideh, Robert Stambaugh, Mathieu Taschereau-Dumouchel, Luke Taylor, Jessica Wachter, Amir Yaron, and participants at the 2015 AFA meetings, ASU (Carey), Michigan (Ross), Minnesota (Carlson), Rochester (Simon), Stanford GSB, Toronto (Rotman), Western (Ivey), and Wharton for several helpful comments. Email: cward@umn.edu.

1 Introduction

Information technology continues to change the way firms do business. While the market valuations of star firms, like Apple and Google, make headlines, the broad economy's continued adoption of IT drives improvements in output and productivity, leading Jorgenson, Ho, and Samuels (2011) to conclude that "...information technology capital input was by far the most significant [contributor to US economic growth over the period 1995–2007]." Substantial controversy exists, however, over IT's future bearing on US growth, perhaps arising from existing analyses' heavy reliance on historical macroeconomic data.[1]

In this paper, I argue that we can learn more about IT's future bearing by better structuring our use of forward-looking financial market data. To my knowledge, the method I develop to do this is new. While I apply it to IT because of the sector's importance to growth and growth's first-order implications for pension fund financial health, government indebtedness, and corporate investment, it could in principle be applied to study other phenomena, such as peak oil, as well.

I begin by building an asset pricing model that endogenously links economy-wide growth to innovation in the IT sector, with an intensity governed by the sector's market valuation.[2] Consistent with this link, I empirically show that the IT sector's price-dividend ratio univariately explains nearly half of the variation in future productivity growth. I then estimate the model's transition paths to match historical data on factor shares, price-dividend ratios, growth rates, and discount rates. This estimated model allows me to study the IT sector's temporal evolution toward its long-run factor share and to predict its future growth contribution by comparing the model's historical and future distributions of productivity growth.

My median estimate's two standard-error confidence interval indicates the transition ends between 2032 and 2038, approximately six decades after its 1974 inception. Future average annual productivity growth for the next twenty years is predicted to approximately

[1]Moore (2003), Brynjolfsson and McAfee (2011), and Byrne, Oliner, and Sichel (2013) are optimistic, whereas Cowen (2011) and Gordon (2000, 2012, 2013) are not. Even *The Economist* held an internet debate over 4-15 June 2013 on whether technological progress is accelerating. The summary is listed here: `http://www.economist.com/debate/files/view/Techprogressartifact0.pdf`.

[2]I define the IT sector in Appendix B. I call the IT sector's complement (the "non-IT" sector) the industrial sector. Regarding factor shares, IT capital—the sum of hardware, software, and communications capital—is treated as a type of capital distinct from industrial capital; the former being referred to as "IT"; the latter, simply as "capital". Both refer to stocks of a quantity of "machines". Hence, the IT-capital ratio, which will be prominently featured in what follows, is analogous to a capital-labor ratio, both ratios measuring a relative intensity of factor use.

fall to 52bps from the 87bps recorded over 1974–2012. This is due to both an intensifying of competition in the IT sector, which reduces the marginal benefit of it innovating, and decreasing returns in the broad economy's employment of IT.

I make two further predictions about the IT revolution. First, the sector is more likely to reach its long-run share within the decade before the median estimate of 2035 than within the decade after: formally, the density of convergence times of when the sector's long-run share is reached is right-skewed. Because dear IT sector valuations lead to economy-wide growth and, importantly, vice versa, the model exhibits a salient general equilibrium effect that hastens the transition.

Second, the model not only makes predictions about changes in growth rates, but also about changes in covariances, and thus changes in the nature of systematic risk. More specifically, the information technology sector is initially exposed to negative shocks to expected growth, but later in the transition serves as a hedge against adverse innovations to expected growth in the long run. Indeed, in the long run bad news about expected growth raises IT's possible future contribution to growth; upon impact, the sector's price-dividend ratio encodes this news and rises. This study more broadly suggests because systematic risk changes slowly over time, studies of asset prices employing a particular sample period for empirical analysis might be prone to generating conclusions that are not fully representative of the secular pattern in the economy being investigated.

To elaborate on how we can use asset prices to forecast a sector's growth prospects and future relative size, consider the Gordon growth model for an economy populated by risk-neutral investors with common discount rate r:

$$\frac{P_0^{(i)}}{D_0^{(i)}} = \frac{1}{r - g^{(i)}},$$

where i denotes the sector; P, the sector's market capitalization; D, its aggregate payout; and g, its dividend growth rate. Specify two sectors, and endow the first sector with a slower growth rate, $g = g^{(1)} < g^{(2)} = g + \Delta$, where $\Delta > 0$ is a growth wedge, possibly reflecting an exceptional dividend growth rate or a growing mass of firms. If this endowment were permanent, the outcome would be trivial: sector one's dividend share tends to zero and sector two dominates in the long run. An interesting analysis emerges, however, if sector two's superior growth rate is transient.

Consider now a *convergence time* $T > 0$ when sector two's growth rate instantaneously converges to sector one's. Sector two's price-dividend ratio becomes[3]

$$\frac{P_0^{(2)}}{D_0^{(2)}} = \frac{1}{r-g-\Delta}\left[1 - \frac{e^{-(r-g-\Delta)T}}{r-g}\Delta\right]. \tag{1}$$

By estimating values of r, g, and Δ, and by observing sector two's price-dividend ratio at a point in time, we can back out an estimated value of T. A corollary of this exercise is that we can infer the future relative size of the sectors because both Δ and T are now known:

$$\frac{D_T^{(2)}}{D_T^{(1)}} = \frac{D_0^{(2)}}{D_0^{(1)}} \times e^{\Delta T},$$

the current dividend ratio scaled by the temporary relative growth factor.

This stylized example illustrates the paper's novelty in inferring a convergence time from asset prices and relating them to future shares and thus potential growth contributions in the economy. In the paper, I proceed to construct a more detailed model by introducing additional features such as stochastic growth, uncertainty and risk, investment, and sectoral interdependence. This fleshed-out model endogenizes the analysis's pertinent variables. It further allows me to match historical paths of macroeconomic and asset market data, and then to use its structure to infer both when the IT sector's transition ends and the associated gains to future productivity growth.

An alternative to model-building would be estimating a vector autoregression that employs present-value restrictions, like those developed in van Binsbergen and Koijen (2010). This study's chief endogenous state variable, IT's factor share, would pose two challenges to favoring this alternative, however. Factor shares are highly persistent processes and the temporal behavior of the unfolding of a transition path would likely be a non-linear and low-frequency movement. Model-building allows me to directly model this movement, thereby overcoming these two challenges.

Related literature

My paper builds on the work that relates financial market performance to shifts in the technological frontier.[4] Pástor and Veronesi (2006, 2009) develop models where learning about

[3] Equation (1) solves $\int_0^T D_0^{(2)} e^{(g+\Delta)s} e^{-rs} ds + e^{-rT} P_T^{(2)}$, where $P_T^{(2)} = D_T^{(2)}/(r-g) = D_0^{(2)} e^{(g+\Delta)T}/(r-g)$. Setting either T or Δ to zero reduces it to the Gordon growth model and to sector one's ratio.

[4] A partial list includes Jovanovic and MacDonald (1994), Greenwood and Jovanovic (1999), Hobijn and Jovanovic (2001), Cochrane (2003), Ofek and Richardson (2003), Abel and Eberly (2012), Kogan,

a firm's profitability or a technology's productivity coincides with periods of high volatility and bubble-like patterns in stock prices. Gârleanu, Panageas, and Yu (2012) study the asset pricing implications of large, infrequent technological innovations that require investment to be adopted. Because firms are heterogeneous, firm-specific adoption is staggered across time, generating economy-wide persistence and investment-driven cycles. I take the presence of the IT sector as given, and study the financial effects of a gradual shift in the technological frontier as the sector expands and transitions towards its long-run factor share. By studying this transition path, I'm able to use the model to predict growth.

That said, my paper adds to the literature linking asset prices to aggregate growth to innovation in the economy.[5] The model developed here extends the work done in Romer's (1990) seminal paper, in a similar direction to the one taken by Comin and Gertler (2006). Kung and Schmid (2015) build a growth model, one that is adopted and modified here, but focus on the quantitative difference of exogenous and endogenous growth. Indeed, because of their insight to couple recursive preferences with an endogenously persistent growth process, they show these types of models performs better in jointly matching macroeconomic and asset market data. While their paper features R&D, an economic quantity, as the chief predictive state variable, my paper places the IT sector's price-dividend ratio, a financial ratio, as the centerpiece. On top of this, I specifically model the R&D activity of IT firms as being subject to decreasing, and not constant, returns to scale. This specification grants the IT sector with growth options that are integral for embedding information about the sector's future size in today's sectoral market valuation.

Gârleanu, Kogan, and Panageas (2012) study a growth model with overlapping generations. They find innovation increases the competitive pressure of existing firms and a lack of intergenerational risk sharing introduces a new source of systematic risk in the economy, called displacement risk. I explicitly model the IT sector and map all features directly to public market data and investment data. I also focus on the model's transition paths by initializing the economy with a small IT sector and studying its evolution to its long-run share, triangulating the model's transition paths with macroeconomic and asset market data and using these cross-equation restrictions to forecast the potential gains to growth.

Finally, my paper fits into the large literature tying cross-sectional and time-series asset

Papanikolaou, Seru, and Stoffman (2012), and Kogan, Papanikolaou, and Stoffman (2013).

[5]For example, Grossman and Helpman (1991), Aghion and Howitt (1992), Beaudry and Portier (2004), Beaudry and Portier (2006), Atkeson and Kehoe (2007), Acemoglu (2009), and Opp (2014).

returns to production economies.[6] Gomes, Kogan, and Yogo (2009) develop a production economy with two types of firms that links heterogeneity in output to differences in average returns. While the firms' decisions are intertwined through a common variable factor of production and the representative household's choices, they otherwise operate independently. My model features two interdependent sectors where one sector's output is the other's input and also generates sectoral differences in average returns. Work on investment-specific shocks, originating with Greenwood, Hercowitz, and Krusell (1997) and later being linked to asset prices by Papanikolaou (2011), suggests that investment-good producers load more than do consumption-good producers on investment shocks, which carry a negative price of risk, thus earning lower returns like growth firms. In my model the IT sector is analogous to an investment-good production sector, but it earns lower returns due to relatively smaller, and eventually negative, loadings on factors with positive risk prices.

I structure the paper as follows. Section 2 describes the environment. Section 3 empirically tests the theory's predictions. Section 4 estimates and analyzes the model. Section 5 concludes.

2 Environment

2.1 Information technology sector

Market structure and product division

The information technology sector is subject to two critical forces: cost structures and network economies.[7] Cost structures characterized by large, fixed costs and tiny marginal costs cultivate supply-side economies of scale.[8] A good that exhibits network economies has a value determined by how many people use it and is often described as demand-side

[6]Contributions include Cochrane (1996), Jermann (1998), Berk, Green, and Naik (1999), Gomes, Kogan, and Zhang (2003), Carlson, Fisher, and Giammarino (2004), Kogan (2004), Zhang (2005), Cochrane (2006), Belo (2010), Kaltenbrunner and Lochstoer (2010), Kogan and Papanikolaou (2013), Kogan and Papanikolaou (2014), and Loualiche (2015).

[7]There are several other important forces that affect the IT sector. See Shapiro and Varian (1999) for an excellent overview.

[8]Bakos and Brynjolfsson (1999) study a strategy of bundling a large number of information goods and selling them for a fixed price, like Microsoft Office. They show empirical evidence that this strategy works better for and is used more widely by the IT sector because its marginal costs of production are low. Other industries rely on bundling less often because their marginal costs are higher, which reduces the net benefit of bundling.

economies of scale.[9] These forces coalesce to endow producers of IT goods with market power and I consequently model the sector as monopolistically competitive.

A fixed continuum of IT firms of unit measure comprise the IT sector. Each firm is identical, of zero measure, and indexed by f. Information technology goods produced by the whole sector are on a continuum of measure N_t, and are indexed by g which indexes a good's particular variety. Each firm monopolistically prices its good(s).[10] Information technology capital is notorious for depreciating quickly, so I make a simplifying assumption: it depreciates fully every period.

Consider the quantity demanded $X_t(g)$, which will be later modeled explicitly, for some IT good $g \in [0, N_t]$. The monopolist of the good takes $X_t(g)$ as given and sets the price to maximize profit, subject to a linear production function that is common to all monopolists. My assumptions imply that each monopolist sets the same price $P_t(g)$ in every period (see Appendix A):

$$P_t(g) = \mu, \text{ for every } g \text{ and } t.$$

In consequence, the quantity demanded $X_t(g)$ and profit earned $\Pi_t(g)$ for each IT good are equal across varieties:

$$X_t(g) = X_t, \text{ for each } g \qquad \text{and} \qquad \Pi_t(g) = \Pi_t = (\mu - 1)X_t, \text{ for each } g.$$

Profitability here is kept simple: every monopolist simply charges a constant markup over marginal cost, and earns the difference multiplied by the quantity demanded. At first pass a constant markup may seem counterfactual, but estimates for the IT sector's average markup shows no discernable trend over time and are remarkably stable (see Table V).

I introduce a parameter $1-\phi$ to denote the probability that an existing IT good becomes

[9] Goolsbee and Klenow (2002) examine the importance of network externalities in the diffusion of home computers. Controlling for many characteristics, they find that people are more likely to first-time buy a computer in areas where a high proportion of households already own one. Additional results suggest these patterns are unlikely to be explained by common unobserved traits or by features of the area.

[10] Thus, the measure of IT goods, N_t, reflects the entire sector's product line. Because any firm produces a zero-measure set of goods, there is no feedback from the price a single firm charges in relation to the prices charged by other firms, so the firm consequently prices its own goods independently of its other goods and of the rest of the sector. Since I focus on the sector as a whole, I abstract from intra-industry strategies to substantially simplify the analysis. You can think about this market structure as having the IT sector provide many differentiated products; for example, the goods could be smart phones, robots, consulting services, and even applications (apps)—any product that enhances productivity.

obsolete, is no longer demanded, and thus has zero value. An IT good's value is

$$V_t = \Pi_t + \phi \mathbb{E}_t[M_{t+1} V_{t+1}],$$

where M_{t+1} is the stochastic discount factor, and $\mathbb{E}_t[\cdot] \doteq \mathbb{E}[\cdot | \mathcal{F}_t]$ denotes the conditional expectation with respect to the filtration $\mathcal{F}_t$ that includes all information up to time t. Because all IT goods have identical values, newly developed goods are expected to share the same value. Information technology firms will thus spend more on research to create new goods when the future value of them is high.

Research division

The information technology sector as a whole spends a lot on research and development.[11] The division for research is contained within the IT sector and is distinguished by two conditions. First, any firm f can conduct research subject to a common, decreasing returns to scale technology parameterized by $\eta_s \in [0, 1)$. Second, each firm's research independently realizes success or failure, and the existence of a stock market allows IT firms' owners to diversify away these idiosyncratic risks.

Firm f's research problem is to choose its research expenditure $S_t(f)$ to maximize

$$\mathbb{E}_t[M_{t+1} R_{t+1}(f)] - S_t(f),$$

where revenue generated by research is $R_{t+1}(f) = \theta_t S_t(f)^{\eta_s} V_{t+1}$: the quantity of IT goods created $\theta_t S_t(f)^{\eta_s}$ multiplied by a good's value tomorrow V_{t+1}. Aggregate research revenue across the IT sector is

$$R_{t+1} \doteq \int_0^1 R_{t+1}(f) df.$$

In equilibrium the marginal benefit of research is equated with its marginal cost for every firm:

$$1 = \theta_t \times \eta_s S_t(f)^{\eta_s - 1} \times \mathbb{E}_t[M_{t+1} V_{t+1}], \text{ for all } f. \tag{2}$$

The left side is the marginal cost of research. The right side is the marginal benefit of research,

[11] Computing and electronics, and software and internet firms constituted 35 percent of total R&D expenditure worldwide in 2011 according to Jaruzelski, Loehr, and Holman (2012, p.6, Exhibit C). I also estimate in Compustat data that as a fraction of sales, research and development expenditure is also higher for IT firms on average: seven to ten percent of sales versus under two percent for industrial firms.

which is interacted with a time-varying externality θ_t, taken as given by an individual firm. I interpret the externality as a measure of research productivity and give it the form

$$\theta_t = \bar{\theta} N_t^{1-\eta_s-\eta_k} K_t^{\eta_k}.$$

The parameter $\bar{\theta}$ scales the model's long-term growth rate. The elasticity of new IT good development with respect to research is η_s, the curvature of the research production technology. The strength of a capital reallocation friction is indexed by η_k. I assume $-1 < \eta_k \leq 0$ to make the externality aid the model in generating an S-shaped diffusion curve by capturing two features:[12]

$$\frac{\partial \theta_t}{\partial N_t} > 0 \qquad \text{and} \qquad \frac{\partial \theta_t}{\partial (K_t/N_t)} < 0.$$

The left-hand derivative sets research productivity to be increasing in N_t, capturing the idea that a set of technologies with a rich set of components, like microprocessors, can be combined and recombined to produce new products. The right-hand derivative is a reduced-form capital reallocation friction, possibly arising from it being difficult to integrate IT into myriad capital, for each integration could require an ad hoc approach, or from the accumulation of knowledge on existing capital disincentivizing the learning about novel capital, as in Atkeson and Kehoe (2007).

Because the quantity of IT goods created is $\theta_t S_t(f)^{\eta_s}$ for each firm and the measure ϕN_t remains the following period, the law of motion is

$$N_{t+1} = \phi N_t + \int_0^1 \theta_t S_t(f)^{\eta_s} df = \phi N_t + \theta_t S_t^{\eta_s}, \text{ with } N_0 > 0. \tag{3}$$

Together, (2) and (3) imply an aggregate condition:

$$S_t = \eta_s (N_{t+1} - \phi N_t) \mathbb{E}_t[M_{t+1} V_{t+1}] = \eta_s \mathbb{E}_t[M_{t+1} R_{t+1}]. \tag{4}$$

Aggregate research expenditure S_t is equated to its aggregate benefit: the product of the net

[12] A S-shaped diffusion curve has three temporal states: initially, the adoption of the new technology is slow because its efficacy could be unclear; later, once the new technology becomes better understood, the pace of adoption rapidly picks up; finally, as the economy becomes saturated with the new technology, its rate of adoption slows and plateaus. Both Atkeson and Kehoe (2007) and Jovanovic and Rousseau (2005) provide empirical evidence of this curve for other economic revolutions, including the IT revolution.

increment in IT goods ($N_{t+1} - \phi N_t$), each good's discounted expected value $\mathbb{E}_t[M_{t+1}V_{t+1}]$, and the share of research revenue expensed during development η_s.

Today's financial market is linked to tomorrow's dynamics in the real economy. To see this, plug (4) into (3) (and temporarily hold $\eta_k = 0$ for clarity):[13]

$$\frac{N_{t+1}}{N_t} = \phi + \left(\bar{\theta}\eta_s^{\eta_s}\mathbb{E}_t\left[M_{t+1}V_{t+1}\right]^{\eta_s}\right)^{\frac{1}{1-\eta_s}}, \tag{5}$$

This equation shows how IT firms have a profit-driven motive to innovate: the greater an IT good's value, which is closely tied to its demand, the greater the IT sector's growth rate. The model thus ties innovation and growth to entry and varieties of goods, a result consistent with empirical evidence presented by Jovanovic and MacDonald (1994).

2.2 Industrial sector

Production function

The industrial sector is composed of identical competitive firms and therefore admits a representative firm. It produces a final good Y_t by combining capital K_t, a composite IT good G_t, and hours worked L_t subject to an exogenous productivity level A_t:

$$Y_t = \left(K_t^{\alpha}\left(A_t L_t\right)^{1-\alpha}\right)^{1-m} G_t^m, \tag{6}$$

where m denotes the share of IT goods in factor income, and α the capital share of non-IT good factor income.[14] I normalize the price of the final good to one. The production function specifies capital, IT goods, and labor as having positive cross-partial derivatives, so by renting more IT goods the marginal products of capital and labor increase.[15]

Composite IT good

At every date t, there is a varied continuum of measure N_t of IT goods. These information technology goods are bundled together into a composite good defined by a constant elasticity

[13] This observation was made previously by Kung and Schmid (2015) and is emphasized again here.

[14] The definition of output here is that of gross output, which includes the output of intermediate goods. Jones (2011) states that gross output is twice as large as net output.

[15] Plant-level evidence on valve manufacturers by Bartel, Ichnowski, and Shaw (2007) corroborates that the adoption of IT does more than simply replace factors: it enhances the productivity of them. Black and Lynch (2001) additionally discover that greater computer usage by nonmanagerial employees raises plant productivity.

of substitution aggregator

$$G_t = \left[\int_0^{N_t} X_t(g)^{\frac{1}{\mu}} dg \right]^{\mu}.$$

The parameter μ measures the degree of variety that each IT good possesses. As μ goes to one, all IT goods would become perfect substitutes and conducting research would become unprofitable. Thus, a restriction is $\mu > 1$. Therefore, the industrial sector is more productive if it equally uses two IT goods versus if it uses twice as much of one. Furthermore, when new IT goods are created, the industrial sector will diversify demand across the larger spectrum of goods, reducing the quantity demanded of every particular good. The variable G_t can be thought of as measuring the industrial sector's technological complexity.

Capital accumulation

The accumulation of capital faces Penrose-Uzawa adjustment costs and depreciates at rate δ:

$$K_{t+1} = (1-\delta)K_t + \Lambda\left(\frac{I_t}{K_t}\right)K_t, \text{ where } \Lambda\left(\frac{I_t}{K_t}\right) = c_0 + \frac{c_1}{1-\frac{1}{\zeta}}\left(\frac{I_t}{K_t}\right)^{1-\frac{1}{\zeta}}. \tag{7}$$

The free parameters c_0 and c_1 are chosen to eliminate adjustment costs in the deterministic steady state (Kaltenbrunner and Lochstoer (2010)).[16] The parameter $\zeta \in (0, \infty)$ sets the elasticity of the investment rate with respect to marginal q, the expected marginal value of an additional unit of capital. If ζ is low, marginal adjustment costs are high.

Stationary productivity

The exogenous source of total factor productivity (TFP) of the firm A_t follows an AR(1) process:

$$\log(A_{t+1}) = \rho \log(A_t) + \epsilon_{t+1},$$

where ϵ_{t+1} is an independently and identically distributed normal random variable with mean zero and constant variance σ^2. I set the autoregressive coefficient, ρ, near one, making exogenous productivity persistent.

[16]Explicitly, $c_0 = \frac{1}{1-\zeta}(g_N^* + \delta)$ and $c_1 = (g_N^* + \delta)^{\frac{1}{\zeta}}$, where the steady-state growth rate of IT goods is g_N^*. Note that $\Lambda'\left(\frac{I_t}{K_t}\right) > 0$ and $\Lambda''\left(\frac{I_t}{K_t}\right) < 0$ for $\zeta > 0$ and $\frac{I_t}{K_t} > 0$. Therefore the steady-state investment rate $\frac{I}{K}^* = \Lambda\left(\frac{I}{K}^*\right) = g_N + \delta$. Investment is always positive because $\Lambda'\left(\frac{I_t}{K_t}\right)$ goes to infinity as $\left(\frac{I_t}{K_t}\right)$ goes to zero.

Maximization

Given an initial capital level of K_0, the firm chooses stochastic sequences of investment, labor, and IT goods $\{I_t, L_t, \{X_t(g)\}_{g\in[0,N_t]}\}_{t\geq 0}$ to maximize the expected present value of dividends:

$$\mathbb{E}_t\left[\sum_{s=0}^{\infty} M_{t|t+s}D_{t+s}\right], \text{ where } D_t = Y_t - I_t - W_tL_t - \int_0^{N_t} P_t(g)X_t(g)dg,$$

where $M_{t|t+s} = M_t \cdot M_{t+1} \cdots M_{t+s}$ is the product of future stochastic discount factors from time t to $t+s$. The firm's optimality conditions are in Appendix A.

Output and balanced growth

Using equilibrium conditions, I rewrite (6) as

$$Y_t = \left(\frac{m}{\mu}\right)^{\frac{m}{1-m}} K_t^{\alpha}(A_tL_t)^{1-\alpha}N_t^{(\mu-1)\frac{m}{1-m}}. \tag{8}$$

To ensure balanced growth, the output equation must display constant returns to scale in reproducible factors (Rebelo (1991))—capital and the measure of IT goods. A required parameter restriction for balanced growth is thus

$$\alpha + (\mu-1)\frac{m}{1-m} = 1. \tag{9}$$

This restriction is interesting. Under the assumption that the long-run value of the IT-capital ratio has not yet been reached, it would be difficult to estimate what m is. The assumption of balanced growth, however, allows me to circumvent this issue by having m be pinned down by the estimates of α and μ, which are plausibly easier to measure today.

2.3 Resource constraint and households

Resource constraint

The final good is used for consumption C_t and for investment:

$$Y_t = C_t + I_t + N_tX_t + S_t.$$

Total investment in this economy is the sum of capital investment, I_t, the total investment of the IT sector, N_tX_t, including its aggregate expenditure on research, S_t.

Households

The economy is populated by a competitive representative household that derives utility from the flow of the single consumption good. It supplies labor perfectly inelastically, so $L_t = 1$ for all t, and has Epstein and Zin (1989) and Weil (1989) recursive preferences:

$$U_t = \left\{(1-\beta)C_t^{1-1/\psi} + \beta\left(\mathbb{E}_t[U_{t+1}^{1-\gamma}]\right)^{\frac{1-1/\psi}{1-\gamma}}\right\}^{\frac{1}{1-1/\psi}},$$

where γ is the coefficient of relative risk aversion, ψ is the elasticity of intertemporal substitution. I assume $\psi > \frac{1}{\gamma}$, so that the agent prefers the early resolution of uncertainty and dislikes shocks to long-run expected growth rates. The stochastic discount factor in the economy is

$$M_{t+1} = \beta\left(\frac{C_{t+1}}{C_t}\right)^{-\frac{1}{\psi}}\left(\frac{\mathbb{E}_t[U_{t+1}^{1-\gamma}]}{U_{t+1}^{1-\gamma}}\right)^{\frac{\gamma-1/\psi}{1-\gamma}}, \tag{10}$$

where the first term is the subjective discount factor, the second term reflects tomorrow's consumption growth, and the third term captures preferences concerning uncertainty about long-run growth prospects. The agent, therefore, requires compensation for these two sources of risk exposure. Indeed, sufficient exposure to persistent, long-run growth prospects creates a mechanism to generate large risk premia, as in Bansal and Yaron (2004).

The household maximizes utility by choosing consumption, earning wage income, and participating in financial markets, taking prices as given. The household participates in financial markets by taking positions in the bond market B_t and in the stock market $\mathcal{S}_t$, which pays an aggregate dividend $\mathcal{D}_t$. The budget constraint of the household is

$$C_t + \mathcal{S}_{t+1}Q_t + B_{t+1} = W_tL_t + (Q_t + \mathcal{D}_t)\mathcal{S}_t + (1 + r_{f,t})B_t,$$

where $\mathcal{S}_tQ_t$ is the aggregate market capitalization and $(1 + r_{f,t})$ is the gross real rate of interest. The variables $\mathcal{D}_t$ and Q_t are defined in (11) and (12). In equilibrium I normalize the aggregate supply of stock $\mathcal{S}_t$ to one and bonds to zero.

2.4 The valuation-productivity link

Stock market valuation

The stock market includes both the IT sector and industrial sector. The aggregate dividend is the sum of the dividends paid by the industrial firm plus the total profits of the IT sector in excess of its expenditure on research:

$$\mathcal{D}_t \doteq \underbrace{D_t}_{\text{Industrial dividend}} + \underbrace{D_t^{IT}}_{\text{IT dividend}}, \tag{11}$$

where $D_t^{IT} \doteq \Pi_t N_t - S_t$. This leads to the following observations, whose proofs are in Appendix A, and which are central to the paper.

Proposition (Stock market valuation)**.** *The aggregate value of the stock market in this economy is*

$$Q_t \doteq \underbrace{q_t K_{t+1}}_{\textit{Industrial sector}} + \underbrace{\underbrace{N_t(V_t - \Pi_t)}_{\textit{Assets-in-place}} + \underbrace{O_t}_{\textit{Growth options}}}_{\textit{IT sector}}, \tag{12}$$

where O_t is defined as the value of IT sector's growth options:

$$O_t \doteq \mathbb{E}_t \left[\sum_{s=1}^{\infty} M_{t|t+s} R_{t+s} \right] (1 - \eta_s).$$

Corollary. *If the returns to scale of doing research become constant ($\eta_s = 1$), growth options embedded in the value of the IT sector become nonexistent and the future IT capital stock is not priced into today's market value.*

The value of the stock market, therefore, incorporates three elements: the ex-dividend value of the industrial sector $q_t K_{t+1}$ and the IT sector, comprised of part assets-in-place and part growth options. The growth options are important. With them, the stock market contains information about the future size of the IT sector $N_{t+s} X_{t+s}$ (which is proportional to R_{t+s} and $N_{t+s} V_{t+s}$). Indeed, if the variety of IT goods is expected to grow, then a dear current market valuation can be justified even if the current IT-capital ratio is low. When research has constant returns to scale, there are no growth options because rents are not earned on research, and the value of the IT sector would be simply be the value of assets-in-place.

Productivity and price-dividend ratios

Given the restriction in (9), we can rewrite (8) as

$$Y_t = \left(\frac{m}{\mu}\right)^{\frac{m}{1-m}} K_t^{\alpha}(A_t N_t L_t)^{1-\alpha} \doteq K_t^{\alpha}(Z_t L_t)^{1-\alpha}. \tag{13}$$

Our usual measurement of productivity is the product of two components: the exogenous, stationary component, and the endogenous, expanding measure of IT goods; call this product $Z_t \doteq \left(\frac{m}{\mu}\right)^{\frac{m}{1-m}} A_t N_t$.[17]

A straightforward derivation for an arbitrary horizon h gives

$$\mathbb{E}_t\left[\log\left(\frac{Z_{t+h}}{Z_t}\right)\right] = (\rho^h - 1)\log(A_t) + \mathbb{E}_t\left[\log\left(\frac{N_{t+h}}{N_t}\right)\right]. \tag{14}$$

Because A_t is persistent, the first term on the right side of (14) is near zero. What does this equation say? It says that the conditional expectation of the IT sector's growth rate is the key for forecasting future productivity.

The link between future productivity growth and today's financial market data is particularly clear when looking at the IT sector's price-dividend ratio:

$$\begin{aligned} PD_t^{IT} \doteq \frac{N_t(V_t - \Pi_t) + O_t}{N_t \Pi_t - S_t} &= \frac{\mathbb{E}_t\left[\sum_{h=1}^{\infty} M_{t|t+h} D_{t+h}^{IT}\right]}{D_t^{IT}} \\ &= \mathbb{E}_t\left[\sum_{h=1}^{\infty} M_{t|t+h}\left(\frac{D_{t+h}^{IT}/N_{t+h}}{D_t^{IT}/N_t} \times \frac{N_{t+h}}{N_t}\right)\right]. \end{aligned} \tag{15}$$

The IT sector's price-dividend ratio therefore contains information not only about its dividend growth but also about the nature of economy-wide productivity growth.

3 Empirical analysis

Predictive regression

The tight relationship between the decomposition of productivity in (14) and the IT sector's price-dividend ratio in (15) is a strong prediction of the model. Does it hold in the data? I

[17] The adoption of IT goods showing up as immediate increases in measured productivity is consistent with the treatment of intermediate goods by Oberfield (2013), who writes a model where an entrepreneur's input choice of an intermediate good comes with an associated productivity-specific match.

test this prediction with the following regression:

$$TFP_{t\to t+h} = a + b \times PD_t^{IT} + e_{t\to t+h},$$

where $TFP_{t\to t+h} = TFP_{t+1} + \cdots + TFP_{t+h}$ is the cumulative growth of TFP over h periods (the TFP variable is measured as the first difference of logarithms and is in percent), and the independent variable PD_t^{IT} is the annualized price-dividend ratio of the IT sector, which has been adjusted for repurchases (see Appendix B). Standard errors need to reflect the error term's overlapping structure, which could potentially be serially correlated; for this reason, I use Hodrick (1992) standard errors, which should perform better than Newey and West's (1987) adjustment because the former sums variances and avoids the latter's summing of autocovariances which are poorly estimated in small samples.

I report the results in Table I in two panels. Panel A documents economic significance. Consistent with a research expenditure affecting the economy with a lag, the effects of the price-dividend ratio are stronger at longer horizons. A greater valuation of IT goods leads to more research expenditure and later the creation of new goods. The adoption of these new goods by the industrial sector coincides with gains in measured productivity growth. A unit change in the price-dividend ratio (from 50 to 51, for example) forecasts a 0.05 percentage point increase in TFP growth over the next four years.

The last column, which reports the expected change in TFP for a one standard-deviation change in the IT sector's price-dividend ratio, shows striking evidence of the mechanism at play. Focusing on the four-year result, a one standard-deviation move in the price-dividend ratio increases TFP growth over the next four years by two percent, or nearly a half percent per year. To put this in perspective, real GDP growth per person is around two percent on average. Also noteworthy are the adjusted R-squareds of the four- and five-year horizons: the price-dividend ratio explains effectively *half* of the variation in TFP growth over the period 1971–2012. I plot the time series for the price-dividend ratio and future five-year cumulative TFP growth in Figure I.

Robustness

Panel B checks robustness by calculating the standard errors via the Newey and West (1987) adjustment and a Monte Carlo method. Due to the persistence of the predictor variable, estimates of the significance of the slope coefficient can be biased (see Stambaugh (1999)). To address this, I compute bias-adjusted small sample (Hodrick (1992)) t-statistics and adjusted

R-squareds, generated by bootstrapping 10,000 samples of the long-horizon regression under the null of no predictability.[18]

I run horse races between the IT sector's and the industrial sector's price-dividend ratios in Table II. The IT sector's price-dividend ratio drives out the industrial sector's when comparing the two measures over all horizons. When adding the industrial sector's price-dividend ratio to the specification, the regression's adjusted R-squared does not increase by much. These results are consistent with the valuation and subsequent innovation in the IT sector driving the majority of future productivity growth in the industrial sector over this period.

One would expect the estimated coefficients of this regression to change over time, with perhaps a greater economic effect being attributed to the early days of the transition. Table III lists the one-, three-, and five-year statistics from regressions estimated over various subperiods: an early period (1971Q1–1986Q4), an intermediate period (1986Q4–2002Q3), and a later period (2002Q3–2012Q4). The first two periods are chosen to have 63 quarters, starting from the sample start date; the second period is chosen to capture the dynamics of the dot-com boom; and the last period estimates the regression on the period following the boom.

Economically, the coefficients are larger the earlier the subperiod and almost monotonically wane as the transition unfolds, as measured by the standard deviation of the fitted value. Initially, the marginal product of an additional IT good is high, so a small increase in the IT-capital ratio shows up as a large gain in productivity. But in later periods, this marginal product falls over time, and with it the economically large gains in measured productivity. The changing slope of adjusted R-squareds, which begins steep but flattens and falls in later subperiods, corroborates this view.

These regressions altogether capture a central fact: the price-dividend ratio of the IT sector contains significant information about the future productivity of the economy.

One concern with the analysis is that the variation of future productivity growth is being explained only by variation in the valuation of the *public* IT sector, and could potentially be biased and not be representative of the entire IT sector. I do not need this measure to be perfect and free of bias to accurately estimate growth expectations, however. The price-dividend ratio employed is constructed from value-weighted returns, as described in

[18] This bootstrapping procedure follows Kilian (1999) and Goyal and Welch (2008). It preserves the autocorrelation of the predictor variable and the contemporaneous correlation of the predictive regression's and the predictor variable's shocks.

Appendix B, and the size of the bias would depend on magnitudes of two factors: the market value of the private firms and the level of their price-dividend ratios. As long as the market value of private firms is small relative to public firms', the estimation error will be small.

Accurately measuring the value of the private IT sector is difficult, simply because good data on this, to my knowledge, does not exist. Instead, I proxy the value of mature private firms—those likely to be genuine challengers or displacers of incumbents—as the total value of technology IPOs in a given year (as similarly classified by NAICS code). Figure II plots the price-dividend ratio of the public IT sector versus a "true" estimate of the entire IT sector's ratio over the period 1974 until 2012. The "true" estimate is a value-weighted price-dividend ratio. While not proof that my employed ratio is becoming unbiased, it is at least comforting that the two ratios converge over time.

Having shown evidence that the IT sector's asset market data has predictive power for economy-wide productivity, I now quantitatively analyze the model, from which I infer the duration and size of future productivity gains from the IT sector's valuation.

4 Model analysis

I present the model analysis in two parts. I first calibrate the household's and industrial sector's parameters to common estimates in the literature. Second, I estimate structurally via direct inference (Gourieroux, Monfort, and Renault (1993)) the IT sector's parameters, completing the model's specification that governs its dynamics. Structural estimation allows me to run a test of overidentifying restrictions to see if the model fits the data well and also to calculate standard errors on the model's key estimates.

In second part of the model analysis, the model is simulated assuming it is on a transition path, which starts the economy at an initial value that is different from the long-run steady state of the model. To define the temporal aspects rigorously, consider a *convergence time* $\widetilde{T}$ adapted to the filtration $\mathcal{F}_t$ that is defined as the first moment when the economy's IT-capital ratio has crossed its unconditional expected value from below:

$$\begin{aligned}\widetilde{T} &\doteq \inf\left\{t : \frac{N_t X_t}{K_t} \geq \lim_{t\to\infty} \mathbb{E}\left[\frac{N_t X_t}{K_t}\right]\right\} \\ &= \inf\left\{t : \frac{N_t X_t}{K_t} \geq \left(\frac{m}{\mu}\right)^{\frac{1}{1-m}} \left[\left(\frac{K}{N}\right)^{*}\right]^{\alpha-1} \exp\left\{\frac{1}{2}(1-\alpha)^2 \frac{\sigma^2}{1-\rho^2}\right\}\right\}, \end{aligned} \tag{16}$$

where $(K/N)^*$ is the deterministic steady-state capital-IT good ratio, which is calculated numerically.[19]

My focus on the IT-capital ratio is consistent with Jorgenson, Ho, and Samuels's (2011) conclusion that the chief contributor to recent economic growth was IT's growing factor share. When the ratio nears its long-run value, the interpretation is that the industrial sector has effectively tapped the major productivity gains that can be exploited from adopting IT and adjusting work practices to best use it.

With this convergence time well-defined, I refer to the economy at time t as being on the transition path if $t < \widetilde{T}$ and as being in the (stochastic) steady state if $t \geq \widetilde{T}$.

Studying the transition path allows me to form predictions about the remaining duration of the IT revolution and the accompanying growth in productivity. More specifically, I can estimate expected future productivity growth by using the model to compare the simulated distributions of productivity growth from the period 1974 until 2013 with the period from 2013 until time $\widetilde{T}$.

The model is ran at a quarterly frequency. The equilibrium is computed numerically using a third-order perturbation method (Schmitt-Grohe and Uribe (2004)).

4.1 Calibration

Industrial sector

The ranges of these parameters have largely been agreed upon by the literature. A usual value for α is in the neighborhood of a third, and I use a value of 0.3. I set the quarterly rate of depreciation δ to 0.02, or around 8 percent at an annual rate, and the adjustment cost parameter ζ to 1.01. I choose ρ and σ to match the first-order autocorrelation and volatility of consumption growth in the steady state.

[19] For intuition, assuming a deterministic model with a risk-neutral household with constant discount rate r gives a steady-state IT-capital ratio of $\left(\frac{NX}{K}\right)^* = \frac{(r+\delta)m}{(1-m)\alpha\mu}$. The full model determines $(K/N)^*$ from the joint solution to the deterministic counterparts of (3), (4), and (22).

Households

Substantial empirical work has been done on these parameters (Bansal, Kiku, and Yaron (2012)). I set $\gamma = 10$, $\psi = 0.9$, and $\beta = 0.9915$ to produce reasonable levels for both sectors' price-dividend ratios. The elasticity of substitution is usually assumed to be greater than one in much of the long-run risks literature (Bansal and Yaron (2004), Croce (2014)) to have the substitution effect dominate the wealth effect, leading positive innovations in expected growth to increase the price-dividend ratio. Note, however, that *on the transition path* there is positive correlation between the sectors' price-dividend ratios and an eight-period moving average of measured TFP growth, $\Delta \log(Z_t)$, a proxy for expected growth. Additionally, the model requires an elasticity of substitution less than one to match the estimated time series of factor loadings, as will be later described.

4.2 Estimation

Information technology sector

The remaining six parameters governing the structure of the IT sector are calibrated and estimated here.

The IT share of factor income parameter m determines the importance of IT in the production. This is unknown under the assumption the steady state has not yet been reached. The choice is disciplined, however, by the balanced growth condition in (9) which specifies m given μ and α, two parameters that are plausibly easier to measure.

The five remaining parameters are estimated using a simulated method of moments procedure:

$$b = \{\bar{\theta}, \mu, \eta_s, \eta_k, \phi\}.$$

The method of moments estimator selects a vector of parameters that minimizes the distance between moments in the data and moments from simulated data produced from the model. Details on the data are in Appendix B.

Let M denote the $K \times 1$ vector of data moments. Given a parameter vector b, for each simulation, $s = 1, \ldots, S$, I simulate a time series of length T and compute a vector of moments from the simulated data $\widetilde{M}_s(b)$. The method of moments estimator for the parameter vector b is defined as

$$\hat{b} = \underset{b}{\operatorname{argmin}} \left(M - \frac{1}{S} \sum_{s=1}^{S} \widetilde{M}_s(b) \right)' W \left(M - \frac{1}{S} \sum_{s=1}^{S} \widetilde{M}_s(b) \right), \tag{17}$$

where W is a positive, semi-definite weighting matrix. Following Hennessy and Whited (2005), I choose $W = (T\Sigma_0)^{-1}$, where

$$\Sigma_0 = \sum_{j=\infty}^{\infty} \mathbb{E}\left([m_t - \mathbb{E}[m_t]][m_{t-j} - \mathbb{E}[m_t]]' \right),$$

and is approximated using the estimator in Newey and West (1987) and m_t denotes the date t observation in the data for each moment in the vector of moments M. Duffie and Singleton (1993) show that under appropriate conditions

$$\sqrt{T}(\hat{b} - b_0) \xrightarrow{d} \mathcal{N}(0, \Omega_0),$$

where $\Omega_0 = (1 + 1/S)\,(H_0' \Sigma_0^{-1} H_0)^{-1}$, b_0 is the true vector of parameters, and

$$H_0 = \left. \frac{\partial \left[\frac{1}{S} \sum_{s=1}^{S} \widetilde{M}_s(b) \right]}{\partial b} \right|_{b=\hat{b}}. \tag{18}$$

Tests of model fit

Following Hennessy and Whited (2005), the technique also provides a test of the overidentifying restrictions of the model with the criterion function

$$\frac{T}{(1 + 1/S)} \left(M - \frac{1}{S} \sum_{s=1}^{S} \widetilde{M}_s(\hat{b}) \right)' W \left(M - \frac{1}{S} \sum_{s=1}^{S} \widetilde{M}_s(\hat{b}) \right), \tag{19}$$

which converges to a Chi-squared distribution with degrees of freedom equal $K - 5$.

The median estimate of the convergence time is subject to parameter uncertainty. To quantify this, I use a variant of the delta method by estimating

$$\operatorname{var}(Median(\widetilde{T})) = \left. \frac{\partial Median(\widetilde{T})}{\partial b} \right|_{b=\hat{b}}' \times \Omega_0 \times \left. \frac{\partial Median(\widetilde{T})}{\partial b} \right|_{b=\hat{b}} \tag{20}$$

and taking its square root. The partial derivatives of $\widetilde{T}$ with respect to the parameter vector are calculated similarly to those calculated in (18) by numerically calculating the difference between the average simulation perturbed by the parameters around the point estimate $\hat{b}$:

$$\left.\frac{\partial Median(\widetilde{T})}{\partial b}\right|_{b=\hat{b}} = \left.\frac{\partial\left[\frac{1}{S}\sum_{s=1}^{S} Median(\widetilde{T})_s(b)\right]}{\partial b}\right|_{b=\hat{b}}.$$

Simulation procedure

For each parameter vector b, I solve the model and then simulate the economy 2,000 times (S = 2,000) for 500 quarters (T = 500) in the following two-step method:

1. Fix an initial IT-capital ratio $\frac{N_0 X_0}{K_0}$ near the 1974 data point
2. Simulate S times an entire shock sequence $\{A_t\}_{t=1,2,\ldots,T}$, for a given set of parameters b and compute model quantities and prices

I match the initial IT-capital ratio to as close to the 1974 data point as possible, but I also require the model to be consistent with initial asset market data as well. Changing the start date by plus-minus five years does not materially affect the results.

Notice the exercise is unorthodox because we do not observe the entire time series from which to estimate the parameters, because by assumption we are currently on the transition path and thus do not observe the data's ergodic distribution. To address this, in each of these simulations I compute the model's moments using only the first 38 years (152 quarters) of simulated data, consistent with the length of the time series used to form estimates in the actual data. I use the remaining 348 quarters to analyze the full dynamics of the system and to estimate the remaining duration of the IT revolution and the accompanying growth in productivity.

Selection of moments

The selection of moments used is important to ensure the five parameters are identified. The nature of the exercise is slightly different from that employed on a stationary model, where means, variances, and autoregressive coefficients usually characterize the stationary distribution. Here I calibrate to the entire transition paths of several variables, essentially generating a time series vector for each variable, leaving open the possibility of selecting a large number ($K \times T$) of moments.

Instead, I compute the averages of several variables that would be expected to be informative about the transition path of the IT sector and further demand that the simulated time series vectors of these variables be broadly consistent with the time series of these variables in the data. More specifically, I choose moments derived from the model's salient variables; namely, its factor shares, price-dividend ratios, growth rates, and discount rates.

The variables I choose are the average continuously compounded growth rates of the IT-capital ratio and the price-dividend ratios of the IT and the industrial sectors. I supplement these with the average net entry rate and sales growth rate of the IT sector, an estimate of IT sector markups, and average estimates of real returns for both sectors, thus focusing the model's estimation on a range of macroeconomic and financial data and reducing the number of moments to simply $K = 8$. I argue that because the analysis is primarily about long-run trends, focusing on levels of trending variables and on growth rates is the first-order concern.

4.3 Discussion of estimation results

Table IV summarizes the calibration and estimates of parameters. I compare these estimates with some that I estimate in the data using reduced-form techniques.

The rate of obsolescence of an IT good $1 - \phi$ in the model should capture two features: a economic obsolescence and default, as weaker firms would be expected to exit the marketplace. A BEA report by Li (2012) lists a 16.5 percent annual depreciation rate for computers and electronics in a two-step estimation procedure that includes an adjustment for obsolescence. This rate is higher than the 15 percent rate applied by the BEA to generic research and development goods. In addition, I estimate the unconditional probability of defaulting using two methods, which are described in Appendix C. Both methods produce results near 3 percent. Because ϕ is interpreted as a measure, I assume economic obsolescence and defaulting are independent and add the two to get $1 - \phi^{\text{Annual}} = 16.5 + 3 = 19.5$ percent, or nearly $\phi = 0.95$ at a quarterly frequency, which is close to my point estimate of 0.945.

To estimate η_s in the data, I approximate (5) to get

$$\log\left(\frac{N_{t+1}}{N_t}\right) \approx \frac{\eta_s}{1-\eta_s}\log\left(\mathbb{E}_t\left[M_{t+1}V_{t+1}\right]\right),$$

and then substitute this equation into (13) to yield what resembles a linear regression:

$$\log\left(\frac{Z_{t+1}}{Z_t}\right) = (\rho - 1)\log(A_t) + \frac{\eta_s}{1-\eta_s}\log\left(\mathbb{E}_t\left[M_{t+1}V_{t+1}\right]\right) + \epsilon_{t+1}.$$

I take this directly to the data to estimate η_s and provide estimates in Table VI. Because the (log) price-dividend ratio better explains TFP variation at a longer horizon, estimates of the four- and five-year horizon are considered. Estimates at these horizons range from 0.69 to 0.93. Griliches (1990) also provides some estimates, which range from 0.6 to 1.0, depending on the use of cross-sectional or panel data. My point estimate of 0.834 is within both ranges.

The average markup on IT goods is $(\mu - 1)$. Given the IT sector's heterogeneous product line, it is tough to measure accurately.[20] One study by Goeree (2008) finds that the median markups on personal computers across the total industry range from 5 to 15 percent, depending on the degree of information possessed by consumers in her limited-information model of consumer behavior. Moreover, direct estimates (see Table V) based on the IT sector's average EBITDA-to-Sales ratio, a measure of markups, are 9.5 and 14 percent, depending if cross-sectional medians or aggregate means are used. The estimates are very stable over time judging by their narrow confidence intervals. My point estimate of 13.4 percent is within these ranges.

Estimating the parameter that governs the capital readjustment cost η_k is difficult, and is one reason a simulated moments estimator was applied. My point estimate is -0.242. This estimate is disciplined by having this single parameter jointly match the S-shaped diffusion dynamic of the IT-capital ratio, the transitions of both sectors' price-dividend ratios, and the IT sector's net entry and sales growth rates.

Finally, the parameter $\bar{\theta}$ is estimated to be to 1.66. This parameter pins down overall growth rate of the economy in the steady state which is estimated to be 1.79 percent (see Table VII).

Altogether, the estimates taken from the simulated moments procedure appear to match well with those estimated in reduced-form exercises, which is comforting. In addition, the Chi-squared test statistic of the overidentifying restrictions, calculated in (19), is 5.70 and with 8 - 5 = 3 degrees of freedom does not produce a rejection of the model at the 10 percent

[20] For example, software and hardware manufacturers abide by different standards. Hardware manufacturers of chipsets, motherboards, and processors abide by an open standard: many motherboards, for example, can take RAM, hard drives, and GPUs from several manufacturers. Software manufacturers, conversely, often times have a dominant player; and this is a symptom of software standards being proprietary. The markup across these two manufacturers could vary considerably.

level.

4.4 Analysis of transition paths

Rather than compare estimates of moments, which are simple averages of growth rates or returns, I demand a tougher test of the model and ask that its simulated sample paths closely match those of the data.

IT-capital ratio

Figure III plots the IT-capital ratio of the data versus that generated by the model. The IT-capital ratio data are only publically available up until 2006. Appendix B discusses in detail its construction. Overall, the model's average simulation does an excellent job of matching the time path of the IT-capital ratio, particularly the fast pace of expansion during the 1990s. An important disciplining device is the model's other transitions.

Price-dividend ratios

The price-dividend ratio of the IT sector is defined in (15). The price-dividend ratio of the industrial sector is $q_t K_{t+1}/D_t$. I plot the system's transitions in Figure IV. The model is able to match the industrial sector's price-dividend ratio to the time series data. Within its standard error bounds, the model can capture the run-up in prices during the dot-com boom and even the drop during Great Recession of 2008. The IT sector's price-dividend ratio of the model is able to capture the trend of the data, which is an important part of the analysis. The model's drop in the IT sector's ratio is consistent with that observed as well. The model has difficulty in generating the magnitude of the dot-com boom. Although this is not surprising because the model is not calibrated to match an episode of a "bubble". It is reassuring, however, that the data reverted back within the model's standard error bounds after the boom.

IT sector average sales growth rates

I plot the transition of the IT sector's average sales growth rate in Figure V. The model matches the fast, initial growth displayed by the data and then its drawn out path to convergence. This is consistent with competition ultimately driving down the sales generated per firm. In the model, initially few goods dominate the marketplace. Over time, as more IT goods enter the marketplace, the industrial firm reduces the quantity demanded of each good. Consequently, sales and profit earned per firm falls.

IT sector net entry rates

I depict in Figure VI the transition of the IT sector's net entry rate. The model matches the sharp decline displayed initially by the data and then its steady path to convergence. The model is unable to get its mean to be negative to match the data after the dot-com boom. The model, however, is able to match the negative data values because the model's lower confidence bound is negative. Note that a concentrating of firms is consistent with a "shake-out" period of an industry, which usually occur later in an industry's lifecycle.

Risk exposures

The stochastic discount factor specified in (10) implies two sources of risk: the first source relates to innovations in realized consumption growth; the second source, to innovations in expected consumption growth. While there is only one source of risk to the economy (ϵ_{t+1}), the IT sector's innovation endogenously generates a low frequency source of risk—variability of expected consumption growth rates.

Following the literature on equity risk premia, this first source is termed short-run risk; the latter, long-run risk. The measurement of short-run risk is standard and is taken to be the quarterly growth rate of real consumption of nondurables and services per capita in the data, and in the model it is simply the growth rate of C_t.

Long-run risk is measured in the model by the return on wealth ($r_{C,t}$), which is directly measurable in the model, but needs to be estimated as a latent variable in the data.[21] In the data, I use three methods to estimate the long-run risk factor, which are described in Appendix C.

Both the short-run risk and long-run risk factors are standardized (to mean zero and variance one) before running the following regression for each sector i over a rolling 50-quarter window:

$$r_{i,t} = a_i + \beta_{i,cg}\Delta c_t + \beta_{i,rc} r_{C,t} + \nu_{i,t}, \quad \nu_{i,t} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\nu^2).$$

Figure VII plots rolling estimates of betas for both short-run risk and long-run risk of the IT sector and compares the model's estimates versus the data's. The model is able to match the data's upward trend in short-run risk exposure and its downward trend in the long-run risk exposure. The model is able to generate the correct direction and sign of these trends

[21]The return on wealth in the model is defined recursively by the equation $W_t = C_t + \mathbb{E}_t[M_{t+1}W_{t+1}]$ or equivalently by $r_{C,t+1} \doteq \frac{W_{t+1}}{W_t - C_t}$.

because I calibrate $\psi = 0.9$. Specifying the elasticity to be greater than one changes the sign of each beta coefficient in the model from positive to negative at each date, thus producing trends that do not match the data. But I should again emphasize here that on the transition path there is a positive correlation between price-dividend ratios and a moving average of past TFP growth.

Moments and discount rates

I report in Table VII the consumption growth statistics of the model in the steady state. The model matches the mean, standard deviation, and first-order autocorrelation of the data in the sample period 1974 until 2012. Having the model generate consistent growth and variability of consumption in the long-run is imposed by the balanced growth condition in (9). This fit ensures the long-run behavior of the model is not counterfactual. It is important to get these consumption statistics right because they determine the properties of the representative household's stochastic discount factor, and thus the correct discounting for the sectors' price-dividend ratios. The model generates a moderate degree of consumption smoothing, which is measured by the relative standard deviation of consumption to output. Total investment volatility is slightly smaller than in the data, but the model nevertheless generates a substantial amount.

The model is able to generate a risk-free rate that has both a low volatility and a high persistence as shown in Table VIII. It is higher than commonly estimated because β is calibrated to match the levels of the sectors' price-dividend ratios.

In the same table, the average discount rates for industrial stocks, IT stocks, and the aggregate stock market are reasonable once leverage is accounted for. A standard estimate of the required adjustment for leverage is two times that of an unlevered claim following Rauh and Sufi (2012). I estimate the IT sector's leverage to be 1.36 from Compustat data.

The model here is calibrated to match the transition paths and first moments of the salient asset market data over the last forty years. I argue that these low-frequency movements are the first-order concerns when choosing targets. One shortcoming of the model is its inability to generate substantial variation in risk premia, deriving partially from the model having only one source of exogenous variation that is calibrated to macroeconomic quantities. The standard deviation of returns in the model falls short of reaching the magnitudes estimated in the data. While volatility plays an important role in leading asset pricing research (Ang, Hodrick, Xing, and Zhang (2006), Bansal, Kiku, Shaliastovich, and Yaron (2015)) and I acknowledge its importance in explaining macroeconomic variation at frequencies equal and greater than those of the business cycle, I find it less compelling that it drives variation in frequencies much lower than those at the business cycle and therefore regard the statistics as a second-order concern.

4.5 Analysis of full system dynamics

Risk exposures

Figure VIII plots the entire transition path of risk loadings for both sectors. There is little change in the estimates of the industrial sector. The information technology sector, however, experiences a dramatic shift in sensitivities over time. Why the change? The sector becomes more exposed to short-run risk as the transition unfolds. Because the initial size of IT sector is small, it absorbs a small amount of the risk to consumption. As the transition proceeds, the sector becomes a larger part of the economy, and therefore absorbs a greater amount of the risk in consumption.

The sector's sensitivity to long-run risks is interesting. It starts positive, and then becomes negative, suggesting it becomes a hedge. Initially, as the economy is growing rapidly because of investment in IT, any adverse shock to long-run growth is a risk to the IT sector, because research is conducted in anticipation of tomorrow's value. As the sector matures, however, its role in the economy changes. An adverse shock to the expected growth rate benefits the sector, because it increases the potential for IT to generate growth in the future. In consequence, the price-dividend ratio and the return of the sector increase upon a realization of an adverse shock to expected growth. Information technology in the steady state acts as a hedge to long-run risk.

In the model factor loadings and risk premia change for both sectors over the transition. This occurs without specifying stochastic volatility; in contrast, the traditional long-run risks framework requires stochastic volatility to generate time-varying risk premia. Here, because the size of the IT sector as well as its ability to generate growth in the economy change over time, the model generates time-varying risk premia with homoskedastic shocks.

More generally, this finding of changing risk exposures during a transition has implications for empirical studies of asset pricing which use a particular sample to estimate parameters governing the joint dynamics of consumption and asset returns. In particular, because the estimated loadings are changing slowly over time, it is probable that the parameter estimates obtained in previous studies might not be fully representative of the secular pattern of the economy being investigated.

Full dynamics

I discuss the average transition dynamics of four key variables in Figure IX. The top-left panel shows the transition path of the IT-capital ratio, the ratio of interest. It follows a S-shaped pattern because of the capital reallocation friction. Initially, the rate of expansion of the IT sector is slow, because the capital reallocation friction drags heavily on research productivity. As time passes, the strength of this friction wanes; consequently, innovation begins feeding on itself, as the sector finds it easier to build new innovations on the top of existing ones. In the later stages of the transition, the marginal returns to both the IT sector innovating and the industrial sector employing IT subside, thereby dampening the growth of information technology. Because of the restriction on balanced growth, eventually all economic quantities grow on average at the same rate.

The bottom-left panel displays the price-dividend ratios of the two sectors. Partly because the value of an IT good is initially very high, the price-dividend ratio of the IT sector is also high. The industrial firm's production function specifies positive cross-partial derivatives for capital and IT, so its value is also higher than its steady state value. The price-dividend ratios converge to their steady-state levels nearly 10 years before the IT-capital ratio, highlighting asset market's forward-looking information that are contained in the IT sector's embedded growth options. The exercising of the growth options over the transition path is analogous to Gârleanu, Panageas, and Yu's (2012) exposition on infrequent technological change.

The top-right panel plots annualized productivity growth $\log(Z_{t+1}/Z_t) * 4$. Growth is much higher, on average, in the first half of the transition than in the latter half. This is because productivity growth is intimately tied to the IT sector's growth rate, and the valuation of an IT good. As the returns to innovating fall, the productivity gains of the economy fall as well because of decreasing returns to scale in the employment of IT.

The bottom-right panel features a marked drop in the profitability of each IT good resulting from an increase in competition. This is consistent with a "lifecycle" view of competition within a sector increasing over the course of its transition to its steady state. This intensifying of competition lowers the returns to innovating, and puts a limit on the possible exceptional gains to growth from an expansion of the IT sector.

Density of convergence times

Figure X plots the density of convergence times. The distribution is skewed right, and because of this I focus on the median estimate. The skewness arises because of a salient equilibrium effect of the model. A dear IT sector price-dividend ratio encourages research to develop new IT goods. These goods are subsequently rented, raising the industrial sector's productivity. Importantly, greater industrial productivity increases its demand for for IT goods, which feeds back into IT-good valuations. When the model is started at a low IT-capital ratio, IT is particularly valuable, so the dynamic is initially strong and puts the bulk of the distribution to the left of the mean.

The calibrated parameters imply that the steady-state IT-capital ratio, $\lim_{t\to\infty} \mathbb{E}\left[\frac{N_t X_t}{K_t}\right]$, is 0.44; in comparison, as of 2006—the last available data point—the ratio is 0.18 and so the relative intensity of the use of IT capital is expected to more than double from 2006. The model's median convergence time is 2035, and puts the revolution's duration at 60 years.[22] The standard error around this median estimate is estimated from (20) and is 1.5 years. A two-standard error band puts the interval of the median convergence time between 2032 and 2038.

Distributions of productivity

Finally, Figure XI plots the model's distributions of historical and future productivity growth. There are two things to notice. The first is that the historical distribution is much more symmetric than the future distribution. This is because both distributions are normalized per year, but the future distribution's convergence time is random. Realizations of quick convergence times are coupled with high rates of TFP growth, generating a long, right tail.

The second is that the historical distribution's mean is higher, by about a factor of about five-thirds (the historical distribution's mean is 4.5 percent and the future distribution's mean is 2.7 percent). This reflects the net entry rate of the IT sector. Initially, it is fast, but later it slows as the competition lessens the profitability of researching new IT goods. This leads to a prediction.

Because of greater competition in the future than in the present, an IT good's value will continue to fall, lowering the incentive to conduct research and to produce new goods. As a result, future TFP growth of the economy is expected to be lower than before. To adjust the model's calculation to bring it more in line with the data, I take the ratio of means to adjust the forecasted TFP growth rate per year. The historical average of utilization-adjusted TFP over the period 1971 until 2012 from the San Francisco Federal Reserve is 87 basis points. Therefore, expected TFP growth per year is $0.87 \times 3/5 = 0.52$, a reduction of 35 basis

[22]The industrial revolution took 70 to 80 years, and the electrical revolution took around 40 years. Jovanovic and Rousseau (2005) show that IT has been diffusing across industries more slowly than electricity did. IT's convergence time, therefore, should be expected to take longer.

points.[23] While a simple adjustment, the qualitative dynamics of intensifying competition and declining returns to employing IT are likely to be featured in a more general class of model.

5 Conclusion

Long-run growth is relevant to today's decisions related to pension fund financial health, government indebtedness, and corporate investment. In this paper, I develop a method that puts economic structure on financial market data to forecast future economic outcomes. In particular, the method stresses the consistency of important long-run dynamics in making its predictions. I apply this method to the IT sector's transition towards it long-run share in the US economy, along with its implications for future growth.

The asset pricing model built endogenously links economy-wide growth to innovation in the IT sector, with an intensity is governed by the sector's market valuation. Consistent with this link, I show empirically that the IT sector's price-dividend ratio univariately explains nearly half of the variation in future productivity growth and that this empirical finding is robust and consistent with intuition. I then estimate the model's transition paths to match historical data on factor shares, price-dividend ratios, growth rates, and discount rates. The estimated model allows me to study the IT revolution in detail and in particular make predictions about its future growth contribution.

Future average annual productivity growth is predicted to fall to 52bps from the 87bps recorded over 1974–2012. This is due to both an intensifying of competition in the IT sector, which reduces the marginal benefit of it innovating, and decreasing returns in the broad economy's employment of IT. My median estimate has a two-standard error band of 2032 to 2038, predicting that the IT sector's transition six decades after its 1974 inception.

Future work could apply this methodology to revolutions past, such as the electricity revolution, to assess its predictive ability in a sample where we observe the entire transition path. It could also be applied in principle to study other phenomena, such as the discovery of a large, exhaustible energy resource.

[23]This 35 basis point adjustment is similar to Robert J. Gordon's "educational plateau adjustment" of 27 basis points he calculated in *The Economist*'s online debate, whose weblink is mentioned in the paper's first footnote.

A Proofs and derivations

Constant markups

The price, $P_t(g)$, of IT good g is chosen to maximize the individual monopolist's profit. Take the quantity demanded $X_t(g)$ for this particular good as given. Each monopolist solves the following static profit maximization problem each period:

$$\max_{P_t(g)} \Pi_t(g) \doteq P_t(g)X_t(g) - X_t(g)$$

Differentiating with respect to $P_t(g)$ and plugging in (24) gives

$$\begin{aligned} P_t(g) &= 1 - \frac{X_t(g)}{\frac{\partial X_t(g)}{\partial P_t(g)}} \\ &= 1 - \frac{P_t(g)(1-\mu)}{\mu}, \\ \Rightarrow P_t(g) &= \mu. \end{aligned}$$

Because the parameter μ is independent of both time and a particular good, it holds for all g and t.

First-order conditions

$$I_t: \qquad q_t = \frac{1}{\Lambda'\left(\frac{I_t}{K_t}\right)} \tag{21}$$

$$K_{t+1}: \qquad q_t = \mathbb{E}_t\left[M_{t+1}\left\{\alpha(1-m)\frac{Y_{t+1}}{K_{t+1}} + q_{t+1}\left((1-\delta) - \Lambda'\left(\frac{I_{t+1}}{K_{t+1}}\right)\cdot\left(\frac{I_{t+1}}{K_{t+1}}\right) + \Lambda\left(\frac{I_{t+1}}{K_{t+1}}\right)\right)\right\}\right]. \tag{22}$$

$$L_t: \qquad W_t = (1-\alpha)(1-m)\frac{Y_t}{L_t} \tag{23}$$

$$X_t(g): \quad P_t(g) = \left(K_t^{\alpha}(A_t L_t)^{1-\alpha}\right)^{1-m} \mu m \left[\int_0^{N_t} X_t(g)^{\frac{1}{\mu}} dg\right]^{\mu m - 1} \frac{1}{\mu} X_t(g)^{\frac{1}{\mu}-1}. \tag{24}$$

The first equation relates marginal q to the investment rate. The second equation is the usual Euler equation for capital policy. The third equation equates the marginal product of labor with the wage rate. The fourth is the derivative with respect to $X_t(g)$. Because the measure of IT goods, N_t, is a quantity not controlled by the activity of any single firm, it is treated as exogenous by the industrial firm.

Stock market valuation

Proof. First, multiply (22) by K_{t+1} to give

$$\begin{aligned}
q_t K_{t+1} &= \mathbb{E}_t \left[M_{t+1} \left\{ \alpha(1-m)Y_{t+1} + q_{t+1} \left((1-\delta)K_{t+1} - \Lambda' \left(\frac{I_{t+1}}{K_{t+1}} \right) \cdot I_{t+1} + \Lambda \left(\frac{I_{t+1}}{K_{t+1}} \right) K_{t+1} \right) \right\} \right] \\
&= \mathbb{E}_t \left[M_{t+1} \left\{ \alpha(1-m)Y_{t+1} + q_{t+1}K_{t+2} - q_{t+1}\Lambda' \left(\frac{I_{t+1}}{K_{t+1}} \right) \cdot I_{t+1} \right\} \right], \text{ by (7)} \\
&= \mathbb{E}_t \left[M_{t+1} \left\{ \alpha(1-m)Y_{t+1} - I_{t+1} + q_{t+1}K_{t+2} \right\} \right], \text{ by (21).}
\end{aligned}$$

The first expression in the parentheses is output times the industrial firm's share of factor income attributed to capital, and the second is the firm's investment expenditure; the sum of the two gives the current dividend D_t. Iterating on the equation and imposing the transversality condition, $\lim_{t\to\infty} \mathbb{E}_t M_{t|t+1} q_{t+1} K_{t+2} = 0$, gives the result.

Next, I show two things. First, as the parameter governing the decreasing returns to research expenditure goes to one, the rents earned on research go to zero, and consequently the growth options embedded in the IT sector go to zero and therefore the future IT capital stock is not priced in today's market value. Second, I show that my definition of the IT sector's ex-dividend values is consistent with it being the expected discounted present value of dividends.

Exploiting the fact that $S_t(f)$ equals S_t for all IT firms, a quick integration of firms across the IT sector shows that $R_{t+1} = \theta_t S_t^{\eta_s} V_{t+1}$. Expected discounted aggregate research profits then are

$$\begin{aligned}
\mathbb{E}_t \left[M_{t+1} R_{t+1} \right] - S_t &= (N_{t+1} - \phi N_t) \mathbb{E}_t \left[M_{t+1} V_{t+1} \right] - S_t \\
&= (N_{t+1} - \phi N_t) \mathbb{E}_t \left[M_{t+1} V_{t+1} \right] (1 - \eta_s),
\end{aligned}$$

where the first equality uses the law of motion for IT goods, equation (3), and the second equality uses the definition of aggregate research expenditure, equation (4).

With this observation we can define the growth options of the IT sector:

$$\begin{aligned}
O_t &\equiv \mathbb{E}_t \left[\sum_{s=1}^{\infty} M_{t|t+s} \left\{ V_{t+s} \left(N_{t+s} - \phi N_{t+s-1} \right) - S_{t+s} \right\} \right] = \mathbb{E}_t \left[\sum_{s=1}^{\infty} M_{t|t+s} \left\{ R_{t+s} - S_{t+s} \right\} \right] \\
&= \mathbb{E}_t \left[\sum_{s=1}^{\infty} M_{t|t+s} R_{t+s} \right] (1 - \eta_s)
\end{aligned}$$

where I used the law of iterated expectations (for $s > t$): $E_s[S_s] = (N_{s+1} - \phi N_s)\mathbb{E}_s[M_{s+1}V_{s+1}]\eta_s$. Thus as $\eta_s \to 1$, rents from research go to zero and the future IT capital stock, $N_{t+s}X_{t+s}$, which is proportional to $N_{t+s}V_{t+s}$, is not priced in today's market value. I'll now show the ex-dividend value of the IT sector is simply the expected discounted present value of dividends:

$$
\begin{aligned}
N_t(V_t - \Pi_t) + O_t &= \phi \mathbb{E}_t\left[M_{t+1}V_{t+1}\right] N_t + \mathbb{E}_t\left[M_{t+1}V_{t+1}(N_{t+1} - \phi N_t) - M_{t+1}S_{t+1} + M_{t+1}O_{t+1}\right] \\
&= \mathbb{E}_t\left[M_{t+1}V_{t+1}N_{t+1} - M_{t+1}S_{t+1} + M_{t+1}O_{t+1}\right] \\
&= \mathbb{E}_t\left[M_{t+1}\left\{\Pi_{t+1}N_{t+1} - S_{t+1} + (V_{t+1} - \Pi_{t+1})N_{t+1} + O_{t+1}\right\}\right].
\end{aligned}
$$

Iterating this equation forward and assuming the transversality conditions of $\lim_{t\to\infty} E_t[M_{t|t+1}(V_{t+1} - \Pi_{t+1})N_{t+1}] = 0$ and $\lim_{t\to\infty} E_t[M_{t|t+1}O_{t+1}] = 0$ proves the claim. □

B Data construction

Macroeconomic and financial data

The macroeconomic data I use begins in 1974, a year thought by leading growth economists to be the inception of the IT revolution (Greenwood and Yorukoglu (1997), Greenwood and Jovanovic (1999), Hobijn and Jovanovic (2001)). It is also the year after CRSP added the NASDAQ index to its database. Moreover, the first Intel microprocessor suitable for desktop use, the "4004", was commercialized 1971, spawning the PC industry. My financial market data begins earlier in 1971, allowing me to apply the Hodrick-Prescott filter before having this data share the same 1974 start date.

Data on US consumption of nondurables and services, gross domestic product, nonresidential, private, fixed investment, and population are from the National Income and Product Accounts of the Bureau of Economic Analysis. Data on the value-weighted market's price-dividend ratio are from CRSP. Risk-free returns are from Ken French's data library. Consumer price index inflation and the spread between Baa and Aaa corporate bonds are obtained from the St. Louis FRED. The TFP measure comes in two forms—adjusted and non-adjusted—from the San Francisco Fed. The adjusted measure adjusts for variation in capacity utilization and hours worked within a workweek.

Construction of IT-capital ratio

Consistent with Jorgenson, Ho, and Samuels's (2011) work, I measure information technology as the sum of hardware, software, and communications capital. I treat it as a type of capital that is distinct from industrial capital. I refer to the former as "IT", and to the latter as "capital". Both types refer to stocks of a quantity of "machines" and are measured in units. Hence, the IT-capital ratio is analogous to a capital-labor ratio, both ratios being a relative intensity of factor use.

Dale Jorgenson provides data from 1948 until 2006 through Harvard's DataVerse, a public database, for two of the four series of interest: the price and value of capital *services* of IT capital, and the price and value of the capital *stock* of tangible capital. The two remaining series required, however, are the capital *service* price and value series for tangible capital, which are not provided on DataVerse.

A capital service of an input measures the flow of services from a quality-adjusted index of the stock of the input. Jorgenson assumes a constant quality and thus the service flow from a stock for each asset within an input is a *constant*—so capital service flows match capital quantity stocks. Using a quality-adjusted service flow, especially for a quickly changing input such as IT, is the best estimate of an input's periodic factor income, which has a close analog to quantities employed in theoretical macroeconomic models, like the one presented in this paper.

Jorgensen estimates the capital service series from the capital stock series in great detail. While a genuine updated series would be preferred, the task requires a tremendous amount of work. Estimates of depreciation rates, price indices, quantity indices, investment tax credits, capital consumption allowances, corporate and personal tax rates, property taxes, and debt versus equity financing values are required for estimates of after-tax real rates of return for *each* of the 65 investment classifications of the Bureau of Economic Analysis. See Appendix B in Jorgenson and Stiroh (2000) for details on this. In place of this, the following was performed:

- Both price and value series of both capital stock and capital service series for tangible capital were taken from Jorgenson and Stiroh (2000, Table B2, p. 74), which covers the years 1959-1998.
- Linear regressions were run of the ratio of services to stock on a constant for both the price and value series. This estimate provided a measure of an average service flow that is derived from the stock. The fit in both regressions resulted in R-squareds of over 98 percent.
- The estimates for the value and price series were then multiplied by their respective tangible capital stock series supplied by DataVerse to get estimates of the longer time series capital services measure.

Finally, for both IT and capital, the value series was divided by the price series to compute a quantity series for both IT capital services and industrial capital services, which have close analogs to the quantities N_tX_t and K_t in the model. These two quantity series were divided to construct the data's counterpart to the ratio of interest, the IT-capital ratio.

IT sector definition

For financial market data, I use the term "information technology" to describe the collection of technologies related to computer software, computer hardware, communications equipment, and those employed by technical consultants hired either to incept or to enhance an adopting company's use of information technology. This latter qualification arises because large producers of IT, like IBM, sell consulting services along with IT itself. Earnings from this service would show up in IBM's financial market data.

Data are from Chicago's Center for Research in Security Prices and Compustat. Data are restricted to stocks trading on the NYSE, AMEX, and NASDAQ exchanges, having share codes 10 and 11, and being US-headquartered firms. A firm is classified as being in the IT sector if it has one of the following North American Industrial Classification System (NAICS) four-digit codes:

- 3341 - Computer and peripheral equipment manufacturing
- 3342 - Communications equipment manufacturing
- 3344 - Semiconductor and other electronic component manufacturing
- 5112 - Software publishers
- 5172 - Wireless telecommunications carriers (except satellite)

- 5174 - Satellite telecommunications
- 5182 - Data processing, hosting, and related services
- 5191 - Other information services (includes Internet publishing and broadcasting and web search portals)
- 5415 - Computer systems design and related services
- 5416 - Management, scientific, and technical consulting services

A firm's Compustat NAICS code is preferred over a firm's CRSP code, when codes conflict, because Compustat NAICS data are more complete and CRSP switched NAICS data sources from Mergent to Interactive Data Corporation in December 2009, possibly changing some firms' classifications. I use primary NAICS codes that are assigned to each firm and that matches its primary activity—generally the activity that generates the most revenue for the establishment.

Price-dividend ratios

Monthly price-dividend ratios are from the CRSP annual value-weighted return series with and without dividends. These series are defined as

$$R_{t+1} \doteq \frac{P_{t+1} + D_{t+1}}{P_t}, \qquad RX_{t+1} \doteq \frac{P_{t+1}}{P_t}.$$

Price-dividend ratios are then constructed as the inverse of

$$\frac{D_{t+1}}{P_{t+1}} = \frac{R_{t+1}}{RX_{t+1}} - 1.$$

By using an annual horizon, the strong seasonal component of dividends is attenuated, even when using monthly or quarterly observations. This definition reinvests dividends to the end of the year, consistent with the methodology of Cochrane (2011).

Because the incidence of firms which repurchase shares has increased, an alternative measure of payouts to equity shareholders is used.[24] Following Bansal, Dittmar, and Lundblad (2005), for every month, denote the number of shares at time t after adjusting for splits, stock dividends, et cetera (using the CRSP share adjustment factor) as n_t. An adjusted capital gain series is constructed for a given firm:

$$RX^*_{t+1} = \left[\frac{P_{t+1}}{P_t}\right] \max\left\{0.95, \min\left\{\frac{n_{t+1}}{n_t}, 1\right\}\right\}.$$

The construction differs from that of Bansal, Dittmar, and Lundblad's (2005) because of the additional maximum operator above, which trims the amount of a repurchase to a maximum of 5 percent of a stock's

[24]Fama and French (2001) document that the proportion of firms paying dividends falls after the introduction of the NASDAQ index in 1973; moreover, estimates from a logistic regression model suggest the propensity to pay dividends also declined. Grullon and Michaely (2002) provide evidence of a SEC regulatory change (Rule 10b-18) which occurred in 1982 granted a safe harbor for repurchasing firms against the previously considered manipulative practice. Repurchase activity, consequently, is much larger post-1983.

total shares outstanding. Without this additional operator, the price-dividend ratios are significantly affected by outliers, especially the IT sector's in the early 1970s when few firms are classified (only 134 firms by 1974). Bounds at 0.8 and 0.9 result in similar price-dividend ratio series. I select a value of 95 percent because the probability of observing a share repurchase greater than 5 percent in a month is 1 percent for both sectors, consistent with usual winsorization practices. As a case in point, FactSet, a data service, reports the largest share repurchaser, as a percentage of total shares outstanding, was Seagate Technology, a manufacturer of hard drives, which repurchased 35.2 percent of its shares outstanding in 2012, or about 3 percent per month.[25]

I also constructed another valuation measure similar to Grullon and Michaely's (2002), which is based on the actual repurchase dollar amounts of common shares in Compustat. For this measure, dividends were constructed as trailing twelve-month sums. The price-dividend ratio of this Compustat-based series was similiar to the CRSP-based series construction above.

I created and considered price-earnings ratios as well. The time series dynamics are similar to those of the dividend series. They are not the preferred series, however, because IT firms spend relatively more money on research than do non-IT firms (see footnote 13). Research can classified by management as either an expenditure before taxes and earnings or as investment in a capital asset and could potentially be used to manipulate earnings. Furthermore, the model has firms pay out all excess cash flows as dividends and hence do not make a dividend distribution decision based on earnings, and so using the price-dividend ratio in the data is a closer match to the model's assumptions.

C Estimation of default and the return on wealth

Estimation of default

Following Campbell, Hilscher, and Szilagyi (2008) and Boualam, Gomes, and Ward (2015), I proxy default with performance-related delisting events on CRSP.[26] I use two methods. First, I simply compute the frequency of delisting and divide it by the starting count of firms for each year for the IT sector. Second, I compute annual percentage change in the number of IT firms for each year, and take the minimum of this measure and a value of zero to only count observations that are negative:

$$\min\left[\frac{n_{t+1} - n_t}{n_t}, 0\right],$$

where n_t is the number of IT firms at year t, sampled at an annual frequency. I then temporally average both methods to estimate the unconditional probability of default. Both methods produce results near 3 percent.

Estimation of long-run risk

I use three methods to estimate the return on the wealth portfolio (long-run risk) in the data:

[25]Details are provided at "http://www.factset.com/" under the BuyBack Quarterly report, 2 April 2013.

[26]Delisting codes used are 500, 550, 552, 560, 561, 574, 580, and 584. They are defined at `http://www.crsp.com/products/documentation/delisting-codes`.

1. Kalman filter
2. Predictive regression
3. Vector autoregression

In what follows, all variables have been demeaned, and all errors below are assumed to be *iid* standard normal random variables. Estimates are detailed in Table A. The Kalman filter method follows Croce (2014). The long-run risk component is estimated via the following system:

$$\begin{aligned}\Delta c_{t+1} &= x_t + \sigma \nu_{t+1} \\ x_{t+1} &= \rho x_t + \sigma \eta_{t+1}.\end{aligned}$$

The Kalman filter estimates the latent state x_t and treats it as the long-run risk component $r_{C,t}$. It is estimated by maximum likelihood.

The predictive regression approach follows Colacito and Croce (2011) where tomorrow's consumption growth is regressed on the value-weighted market price-dividend ratio, the risk-free rate, lagged consumption growth, the consumption-output ratio, and a measure of default risk (the Baa-Aaa spread):

$$\Delta c_{t+1} = \beta X_t + \sigma \nu_{t+1}, \text{ where } X_t = \{\Delta c_t, pd_t, r_{f,t}, cy_t, def_t\}.$$

The long-run risk component can be extracted by projecting tomorrow's consumption growth onto today's state variables X_t: $r_{C,t} = \text{proj}[\Delta c_{t+1}|X_t] = \widehat{\beta} X_t$.

Finally, specifying a vector autoregression using the same state vector as above

$$X_{t+1} = AX_t + \Sigma \nu_{t+1}$$

can be used to extract the long-run risk component, the expected discounted value of consumption growth over the infinite horizon:

$$r_{C,t} = \mathbb{E}_t \left[\sum_{s=0}^{\infty} \kappa^s \Delta c_{t+s} \right] = (1 - \kappa A)^{-1} X_t \mathbf{1}_{\Delta c_t},$$

where $\mathbf{1}_{\Delta c_t}$ is an indicator that picks out the vector associated with consumption growth. The discount factor κ is related to the unconditional mean of the price-consumption ratio as in Campbell and Shiller (1988). I set it to 0.965, a value consistent with Lustig, Nieuwerburgh, and Verdelhan's (2013) work. Results are not dependent on this setting.

Table A: Estimates of return on wealth (long-run risk)

Panel A reports the maximum likelihood estimates of a Kalman filter to extract the return on wealth using only consumption data. Panel B reports the VAR coefficients of the matrix A. The results from the predictive regression follow from using the predicted values of the top row in the A matrix. Data are quarterly and cover the years 1971–2012. Data construction is described in Appendix B.

Panel A: Kalman filter estimates

Parameter	Estimate
ρ	0.78***
σ	0.0024***

$p < 0.01$ - ***

Panel B: VAR estimates

	Δc_t	PD_t^{MKT}	r_t^f	cy_t	def_t
Δc_{t+1}	0.448***	0.000	-0.0003***	-0.049***	-0.001
PD_{t+1}^{MKT}	128.03***	0.974***	0.1239*	19.58***	-1.119***
r_{t+1}^f	6.978	-0.008	0.907***	-7.022*	-0.299*
cy_{t+1}	-0.184***	0.00005***	0.0003***	0.99***	0.000
def_{t+1}	-8.91**	0.000	0.019***	0.610	0.812***

OLS standard errors $p < 0.01$ - ***, $p < 0.05$ - **, $p < 0.1$ - *

References

Abel, Andrew B., and Janice C. Eberly, 2012, Investment, valuation, and growth options, *Quarterly Journal of Finance* 2.

Acemoglu, Daron, 2009, *Introduction to Modern Economic Growth* (Princeton University Press).

Aghion, Phillipe, and Peter Howitt, 1992, A model of growth through creative destruction, *Econometrica* 60, 323–351.

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.

Atkeson, Andrew, and Patrick J. Kehoe, 2007, Modeling the transition to a new economy: Lessons from two technological revolutions, *American Economic Review* 97, 64–88.

Bakos, Yannis, and Erik Brynjolfsson, 1999, Bundling information goods: Pricing, profits, and efficiency, *Management Science* 45, 1613–1630.

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad, 2005, Consumption, dividends, and the cross section of equity returns, *Journal of Finance* 60, 1639–1672.

Bansal, Ravi, Dana Kiku, Ivan Shaliastovich, and Amir Yaron, 2015, Volatility, the macroeconomy and asset prices, *Journal of Finance* forthcoming.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An empirical evaluation of the long run risks model for asset prices, *Critical Finance Review* 1, 183–221.

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

Bartel, Ann, Casey Ichnowski, and Kathryn Shaw, 2007, How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills, *Quarterly Journal of Economics* 122, 1721–1758.

Beaudry, Paul, and Franck Portier, 2004, An exploration into pigou's theory of cycles, *Journal of Monetary Economics* 51, 1183–1216.

——— , 2006, Stock prices, news, and economic fluctuations, *American Economic Review* 96, 1293–1307.

Belo, Frederico, 2010, Production-based measures of risk for asset pricing, *Journal of Monetary Economics* 57, 146–163.

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

Black, Sandra E., and Lisa M. Lynch, 2001, How to compete: The impact of workplace practices and information technology on productivity, *Review of Economics and Statistics* 83, 434–445.

Boualam, Yasser, Joăo Gomes, and Colin Ward, 2015, Understanding the behavior of distressed stocks, working paper.

Brynjolfsson, Erik, and Andrew McAfee, 2011, *Race Against The Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy* (Digital Frontier Press).

Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel, 2013, Is the information technology revolution over?, Discussion Paper, 25 Centre for the Study of Living Standards International Productivity Monitor.

Campbell, John Y., Jens Hilscher, and Jan Szilagyi, 2008, In search of distress risk, *Journal of Finance* 63, 2899–2939.

Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

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.

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

——— , 2003, Stocks as money: Convenience yield and the tech-stock bubble, in William C. Hunter, George G. Kaufman, and Michael Pomerleano, ed.: *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies* (MIT Press: Cambridge and London).

——— , 2006, The dog that did not bark, *Review of Financial Studies* 21, 1533–1575.

——— , 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

Colacito, Riccardo, and Mariano M. Croce, 2011, Risks for the long run and the real exchange rate, *Journal of Political Economy* 119, 153–181.

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

Cowen, Tyler, 2011, *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better Again* (Dutton).

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

Duffie, Darrell, and Kenneth J. Singleton, 1993, Simulated moments estimation of markov models of asset prices, *Econometrica* 61, 929–952.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Journal of Finance* 57, 937–969.

Fama, Eugene F., and Kenneth R. French, 2001, Disappearing dividends: Changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 3–43.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012, Displacement risk and asset returns, *Journal of Financial Economics* 105, 491–510.

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

Goeree, Michelle Sovinsky, 2008, Limited information and advertising in the us personal computer industry, *Econometrica* 76, 1017–1074.

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

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

Goolsbee, Austan, and Peter J. Klenow, 2002, Evidence on learning and network externalities in the diffusion of home computers, *Journal of Law and Economics* 45, 317–343.

Gordon, Robert J., 2000, Does the new economy measure of to the great inventions of the past?, *Journal of Economic Perspectives* 14, 49–74.

———, 2012, Is U.S. economic growth over? Faltering innovation confronts the six headwinds, NBER working paper No. 18315.

———, 2013, Is U.S. economic growth over? Faltering innovation confronts the six headwinds, Discussion Paper, 25 Centre for the Study of Living Standards International Productivity Monitor.

Gourieroux, Christian, Alain Monfort, and Eric Renault, 1993, Indirect inference, *Journal of Applied Econometrics* 8, S85–S118.

Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, 1997, Long-run implications of investment-specific technological change, *American Economic Review* 87, 342–362.

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

Greenwood, Jeremy, and Mehmet Yorukoglu, 1997, 1974, *Carnegie-Rochester Conference Series on Public Policy* 46, 49–95.

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

Grossman, Gene M., and Elhanan Helpman, 1991, *Innovation and Growth in the Global Economy* (MIT Press).

Grullon, Gustavo, and Roni Michaely, 2002, Dividends, share repurchases, and the substitution hypothesis, *Journal of Finance* 57, 1649–1684.

Hennessy, Christopher A., and Toni M. Whited, 2005, Debt dynamics, *Journal of Finance* 60, 1129–1165.

Hobijn, Bart, and Boyan Jovanovic, 2001, The information-technology revolution and the stock market: Evidence, *American Economic Review* 91, 1203–1220.

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

Jaruzelski, Barry, John Loehr, and Richard Holman, 2012, The 2012 global innovation 1000 study: Making ideas work, Discussion paper, Booz and Company.

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

Jones, Charles I., 2011, Intermediate goods and weak links in the theory of economic development, *American Economic Journal: Macroeconomics* 3, 1–28.

Jorgenson, Dale W., Mun S. Ho, and Jon D. Samuels, 2011, Information technology and us productivity growth: Evidence from a prototype industry production account, *Journal of Productivity Analysis* 36, 159–175.

Jorgenson, Dale W., and Kevin J. Stiroh, 2000, Raising the speed limit: Us economic growth in the information age, *Brookings Papers on Economic Activity* 31, 126–236.

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, 2005, General purpose technologies, in Philippe Aghion, and Steven N. Durlauf, ed.: *Handbook of Economic Growth*vol. 1 . chap. 18, pp. 1182–1221 (Elsevier B.V.).

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

Kilian, Lutz, 1999, Exchange rates and monetary fundamentals: What do we learn from long-horizon regressions?, *Journal of Applied Econometrics* 14, 491–510.

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

——— , and Dimitris Papanikolaou, 2013, Firm characteristics and stock returns: The role of investment-specific shocks, *Review of Financial Studies* 26, 2718–2759.

——— , 2014, Growth opportunities, technology shocks and asset prices, *Journal of Finance* 69, 675–718.

——— , Amit Seru, and Noah Stoffman, 2012, Technological innovation, resource allocation, and growth, working paper.

Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, 2013, Technological innovation: Winners and losers, working paper.

Kung, Howard, and Lukas Schmid, 2015, Innovation, growth, and asset prices, *Journal of Finance* forthcoming.

Li, Wendy C. Y., 2012, Depreciation of business r&d capital, Discussion paper, Bureau of Economic Analysis R&D Satellite Account Paper.

Loualiche, Erik, 2015, Asset pricing with entry and imperfect competition, working paper.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2013, The wealth-consumption ratio, *Review of Asset Pricing Studies* 3, 38–94.

Moore, Gordon E., 2003, No exponential is forever: But forever can be delayed!, in *Solid-State Circuits Conference* vol. 1 of *Digest of Technical Papers* pp. 20–23.

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

Oberfield, Ezra, 2013, Business networks, production chains, and productivity: A theory of input-output architecture, working paper.

Ofek, Eli, and Matthew Richardson, 2003, Dotcom mania: The rise and fall of internet stock prices, *Journal of Finance* 58, 1113–1138.

Opp, Christian C., 2014, Venture capital cycles, working paper.

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

Pástor, Ľuboš, and Pietro Veronesi, 2006, Was there a nasdaq bubble in the late 1990s?, *Journal of Financial Economics* 81, 61–100.

——— , 2009, Technological revolutions and stock prices, *American Economic Review* 99, 1451–1483.

Rauh, Joshua D., and Amir Sufi, 2012, Explaining corporate capital structure: Product markets, leases, and asset similiarity, *Review of Finance* 16, 115–155.

Rebelo, Sergio, 1991, Long-run policy analysis and long-run growth, *Journal of Political Economy* 99, 500–521.

Romer, Paul M., 1990, Endogenous technological change, *Journal of Political Economy* 98, S71–S102.

Schmitt-Grohe, S., and M. Uribe, 2004, Solving dynamic general equilibrium models using a second-order approximation to the policy function, *Journal of Economic Dynamics and Control* 28, 755–775.

Shapiro, Carl, and Hal R. Varian, 1999, *Information Rules: A Strategic Guide to the Network Economy* (Harvard Business School Press).

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

van Binsbergen, Jules H., and Ralph S. J. Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439–1471.

Weil, Phillipe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–421.

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

Table I: TFP-forecasting regressions I

The regression equation is $TFP_{t \to t+h} = a + b \times PD_t^{IT} + \epsilon_{t \to t+h}$. The dependent variable $TFP_{t \to t+h}$ is the utilization-adjusted TFP measure provided by the San Francisco Federal Reserve, which is a percentage change (quarterly log change times 100). The independent variable is the IT sector's price-dividend ratio adjusted for repurchases. Data are quarterly over 1971Q1–2012Q4. Panel A's t-statistics use the Hodrick (1992) correction equal to the forecast horizon length. $\sigma(\mathbb{E}[TFP])$ is the standard deviation of the fitted value: $\sigma(\hat{b} \times PD_t^{IT})$. Panel B reports the t-statistics calculated under Newey and West (1987) (t_{NW}), and a Monte Carlo bootstrap method (t_{MC}), developed by Kilian (1999) and used in Goyal and Welch (2008) (with Hodrick (1992) standard errors). I also report an adjusted R-squared for the Monte Carlo bootstrap under $\bar{R}^2_{MC}$. Data sources and definitions for the IT sector are detailed in Appendix B.

Panel A				
Horizon h	b	$t(b)$	$\bar{R}^2$	$\sigma(\mathbb{E}[TFP])$
1 year	0.01	1.0	0.08	0.45
2 year	0.03	2.1	0.21	0.97
3 year	0.04	3.4	0.39	1.56
4 year	0.05	4.7	0.48	1.97
5 year	0.06	4.9	0.48	2.23
Panel B				
Horizon h	b	$t_{NW}(b)$	$t_{MC}(b)$	$\bar{R}^2_{MC}$
1 year	0.01	2.5	0.7	0.03
2 year	0.03	4.1	1.5	0.17
3 year	0.04	3.4	2.3	0.34
4 year	0.05	6.7	2.8	0.43
5 year	0.06	6.7	3.3	0.43

Table II: TFP-forecasting regressions II

The regression equation is $TFP_{t\rightarrow t+h} = a + b_{IT} \times PD_t^{IT} + b_{IND} \times PD_t^{IND} + \epsilon_{t\rightarrow t+h}$. The dependent variables are the standard TFP measure ("TFP") and the utilization-adjusted TFP measure ("Adjusted TFP"), both of which are provided by the San Francisco Federal Reserve and are in percentage change (quarterly log change times 100). The independent variables are the repurchase-adjusted price-dividend ratios for the IT sector and the industrial sector. Data are quarterly over 1971Q1–2012Q4. Standard errors use the Hodrick (1992) correction equal to the forecast horizon length. t-statistics are in parentheses. Data sources and definitions are detailed in Appendix B.

		TFP			Adjusted TFP		
Horizon h	Statistic	(1)	(2)	(3)	(1)	(2)	(3)
1 year	b_{IT}	0.01		0.00	0.01		0.01
	$t(b_{IT})$	(0.30)		-(0.03)	(1.00)		(0.44)
	b_{IND}		0.04	0.04		0.06	0.04
	$t(b_{IND})$		(0.50)	(0.56)		(1.20)	(0.50)
	$\bar{R}^2$	*0.01*	*0.02*	*0.02*	*0.08*	*0.08*	*0.10*
2 year	b_{IT}	0.01		0.01	0.03		0.02
	$t(b_{IT})$	(0.60)		(0.23)	(2.10)**		(0.89)
	b_{IND}		0.08	0.06		0.12	0.07
	$t(b_{IND})$		(0.80)	(0.54)		(1.80)*	(0.59)
	$\bar{R}^2$	*0.03*	*0.05*	*0.04*	*0.21*	*0.19*	*0.24*
3 year	b_{IT}	0.03		0.02	0.04		0.03
	$t(b_{IT})$	(1.40)		(0.83)	(3.40)***		(1.89)*
	b_{IND}		0.12	0.05		0.18	0.07
	$t(b_{IND})$		(1.20)	(0.38)		(2.20)**	(0.65)
	$\bar{R}^2$	*0.11*	*0.08*	*0.11*	*0.39*	*0.28*	*0.42*
4 year	b_{IT}	0.05		0.04	0.05		0.04
	$t(b_{IT})$	(2.90)***		(2.34)***	(4.70)***		(3.46)***
	b_{IND}		0.17	0.03		0.21	0.07
	$t(b_{IND})$		(1.70)*	(0.26)		(2.20)**	(0.76)
	$\bar{R}^2$	*0.30*	*0.15*	*0.30*	*0.48*	*0.31*	*0.50*
5 year	b_{IT}	0.06		0.05	0.06		0.06
	$t(b_{IT})$	(4.30)***		(4.16)***	(4.90)***		(4.10)***
	b_{IND}		0.20	0.05		0.21	0.05
	$t(b_{IND})$		(1.80)*	(0.43)		(2.10)**	(0.52)
	$\bar{R}^2$	*0.38*	*0.20*	*0.38*	*0.48*	*0.25*	*0.48*

*** - $p < 0.01$, ** - $p < 0.05$, * - $p < 0.1$

Table III: TFP-forecasting regressions III

The regression equation is $TFP_{t\to t+h} = a + b_{IT} \times PD_t^{IT} + b_{IND} \times PD_t^{IND} + \epsilon_{t\to t+h}$. The dependent variable $TFP_{t\to t+h}$ is the utilization-adjusted TFP measure provided by the San Francisco Federal Reserve, which is a percentage change (quarterly log change times 100). The independent variables are the repurchase-adjusted price-dividend ratios for the IT sector and the industrial sector. Data are quarterly, and over various subperiods that are marked. Standard errors use the Hodrick (1992) correction equal to the forecast horizon length. t-statistics are in parentheses. $\sigma(\mathbb{E}[TFP])$ is the standard deviation of the fitted value: $\sigma(\hat{b} \times PD_t^{IT})$. Data sources and definitions are detailed in Appendix B.

Panel A: 1971Q1 − 1986Q4				
	b	$t(b)$	$\bar{R}^2$	$\sigma(\mathbb{E}[TFP])$
1 year	0.01	1.3	0.06	0.47
3 year	0.05	3.0	0.37	1.71
5 year	0.05	4.5	0.39	2.05
Panel B: 1986Q4 − 2002Q3				
	b	$t(b)$	$\bar{R}^2$	$\sigma(\mathbb{E}[TFP])$
1 year	0.01	2.2	0.09	0.38
3 year	0.04	4.0	0.30	1.02
5 year	0.05	5.7	0.41	1.52
Panel C: 2002Q3 − 2012Q4				
	b	$t(b)$	$\bar{R}^2$	$\sigma(\mathbb{E}[TFP])$
1 year	0.05	2.3	0.15	0.66
3 year	0.06	2.2	0.14	0.88
5 year	0.05	6.4	0.35	0.81

Table IV: Calibration and point estimates (quarterly)

Parameter	Value (std. err.)	Description
Household		
β	0.9915	Subjective discount factor
ψ	0.9	Elasticity of intertemporal substitution
γ	10	Coefficient of relative risk aversion
Industrial		
α	0.3	Non-IT capital share of factor income
δ	0.02	Capital depreciation rate
ζ	1.01	Adjustment cost parameter
ρ	0.9945	Exogenous TFP persistence
σ	0.0175	Productivity volatility
IT		
$\widehat{\overline{\theta}}$	1.66 (3.72)	Scale parameter
$\widehat{\mu} - 1$	0.134 (0.011)	IT net markup
$\widehat{\phi}$	0.945 (0.085)	IT good survival rate
$\widehat{\eta_s}$	0.834 (0.76)	Elasticity of new IT goods wrt R&D
$\widehat{\eta_k}$	-0.242 (0.33)	Capital reallocation friction

Chi-squared statistic for the test of overidentifying restrictions in (19) = 5.70 (p-value = 12.67%)

Table V: IT sector markups

This table reports average markups of IT firms over the annual period 1974–2012. Markups are estimated by $\mu = \frac{1}{1-x} - 1$, where x is the EBITDA-sales ratio, defined below. The row "Aggregate" refers to the sum of EBITDA divided by the sum of sales, and then temporally estimates the average value. The row "Cross section" takes the cross-sectional median of all firms in every year, and then temporally estimates the average value. Standard errors have the Newey-West (1987) adjustment with three lags. Data are defined in Appendix B.

	Lower 95%	Estimate	Upper 95%
Aggregate	0.138	0.142	0.145
Cross-section	0.089	0.093	0.098

Table VI: Estimates of η_s

This table estimates η_s from the regression $TFP_{t\to t+h} = a + b\log PD_t^{IT} + \epsilon_{t+h}$. The dependent variable $TFP_{t\to t+h}$ is the utilization-adjusted San Francisco Fed's utilization-adjusted TFP in percent. The independent variable is the (log) price-dividend ratio for the IT sector, adjusted for repurchases. Data are quarterly, from 1974Q1–2012Q4. The model counterpart is $\log\left(Z_{t+1}/Z_t\right) = (\rho-1)\log A_t + \frac{\eta_s}{1-\eta_s}\log \mathbb{E}_t\left[M_{t+1}V_{t+1}\right] + \epsilon_{t+1}$. The parameter $\widehat{\eta_s}$ is retrieved from the estimate of $\widehat{b}$ by $\eta_s(b) = \frac{b}{1+b}$. Ninety-five percent confidence intervals are constructed using the delta method: $se\left(\widehat{\eta}\right) = \eta_s'(\widehat{b})\left[se(\widehat{b})\right]^2 \eta_s'(\widehat{b})$, where $se(\widehat{b})$ is computed with the Newey-West (1987) adjustment with three lags. Data are described in Appendix B.

	Lower 95%	$\widehat{\eta_s}$	Upper 95%	$\bar{R}^2$
1 year	0.35	0.46	0.56	0.08
2 year	0.53	0.65	0.78	0.21
3 year	0.64	0.75	0.86	0.39
4 year	0.69	0.79	0.90	0.48
5 year	0.72	0.82	0.93	0.50

Table VII: Model and data macroeconomic statistics

This table reports statistics generated by the model's steady state and those found in the the data. The model is simulated quarterly and time-aggregated to an annual frequency. The steady state covers time period $s > \widetilde{T}$, where $\widetilde{T}$ is defined in (16). Data statistics are calculated from Bureau of Economic Analysis's annual data over 1974–2012. The model's total investment, INV, is the sum of capital investment, IT investment, and research expenditure; in the data, total investment is real, nonresidential, fixed, private investment. Δx is the log-difference of the variable x. Data sources are discussed in Appendix B.

Statistic	Data	Model
Consumption growth		
mean(Δc)	2.00	1.79
std(Δc)	2.27	2.38
AC1(Δc)	0.39	0.40
Business cycle		
$\sigma_{\Delta c}/\sigma_{\Delta y}$	0.61	0.96
$\sigma_{\Delta INV}/\sigma_{\Delta c}$	4.38	3.07

Table VIII: Asset pricing moments

Panel A reports the model's annualized moments. Equity premia include a leverage adjustment: with constant financial leverage, the levered equity premium is $\mathbb{E}[r_i^{\text{LEV}} - r_f] \doteq \mathbb{E}[r_i - r_f](1 + \overline{D/E})$, where $\overline{D/E}$ is the average debt-equity ratio, which is set to one to be consistent with firm-level data (Rauh and Sufi (2012)) for industrial firms, and is set to 0.36 for IT firms to match my estimate of the sector's average debt-equity ratio in Compustat. Volatility is also scaled by the same leverage factors. Panel B reports my estimates of the data's annual moments. Returns are value-weighted, monthly from the period 1971 until 2012, and deflated by the consumer price index. The portfolio strategy would be to buy firms at their post-IPO price at month-end and sell them at the delisting price, if occurring. All numbers are in percent except the Sharpe ratios.

Panel A: Model			
	Mean	Stdev	Sharpe ratio
$\mathbb{E}[r_f]$	3.86	0.8	0
$\mathbb{E}[r_{IND}]$	7.46	4.36	0.88
$\mathbb{E}[r_{IT}]$	4.39	3.46	0.21
$\mathbb{E}[r_{MKT}]$	6.46	5.40	0.51
Panel B: Data			
	Mean	Stdev	Sharpe ratio
$\mathbb{E}[r_f]$	1.49	0.9	0
$\mathbb{E}[r_{IND}]$	6.44	15.5	0.38
$\mathbb{E}[r_{IT}]$	5.40	28.3	0.27
$\mathbb{E}[r_{MKT}]$	6.00	16.1	0.34

Figure I: Five-year TFP growth vs. $\mathbf{PD_t^{IT}}$

The process is blue is $TFP_{t\rightarrow t+20}$, the five-year cumulative growth of utilization-adjusted TFP measure provided by the San Francisco Federal Reserve, which is a percentage change (quarterly log change times 100). The process in dashed-green is the IT sector's price-dividend ratio adjusted for repurchases. Data are quarterly, from 1974Q1–2012Q4. Data sources and definitions are detailed in Appendix B.

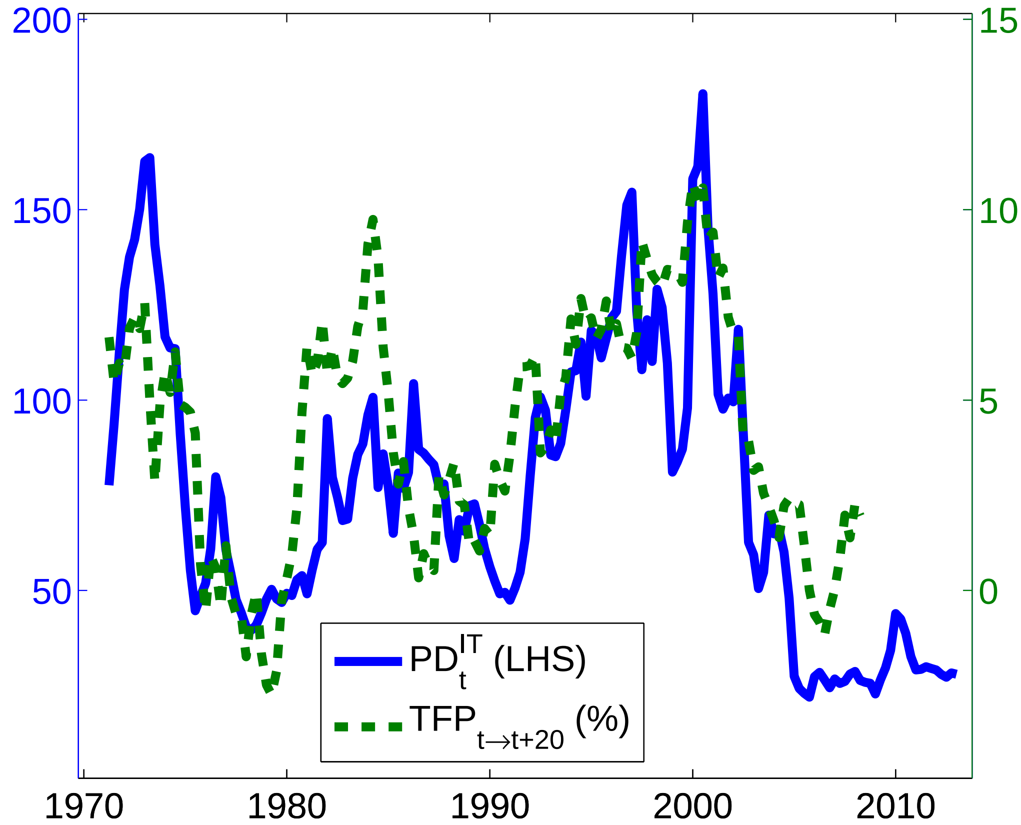

Figure II: Price-dividend ratios of IPO firms and incumbents

This figure plots the relative market values of private firms, proxied by the value of IPOs, to that of incumbents. "Tech P/D" is the annualized price-dividend ratio of the IT sector at year-end, which has been adjusted for repurchases. "True P/D" comes from the formula: $P/D_{True} = \frac{MV_{public}}{MV_{public}+MV_{private}} \times P/D_{Tech} + \frac{MV_{private}}{MV_{public}+MV_{private}} \times P/D_{Private}$. $P/D_{Private}$ comes from the value-weighted price-dividend ratios of all new IPO firms in the IT sector in the year, which has been adjusted for repurchases. Data sources and definitions are detailed in Appendix B.

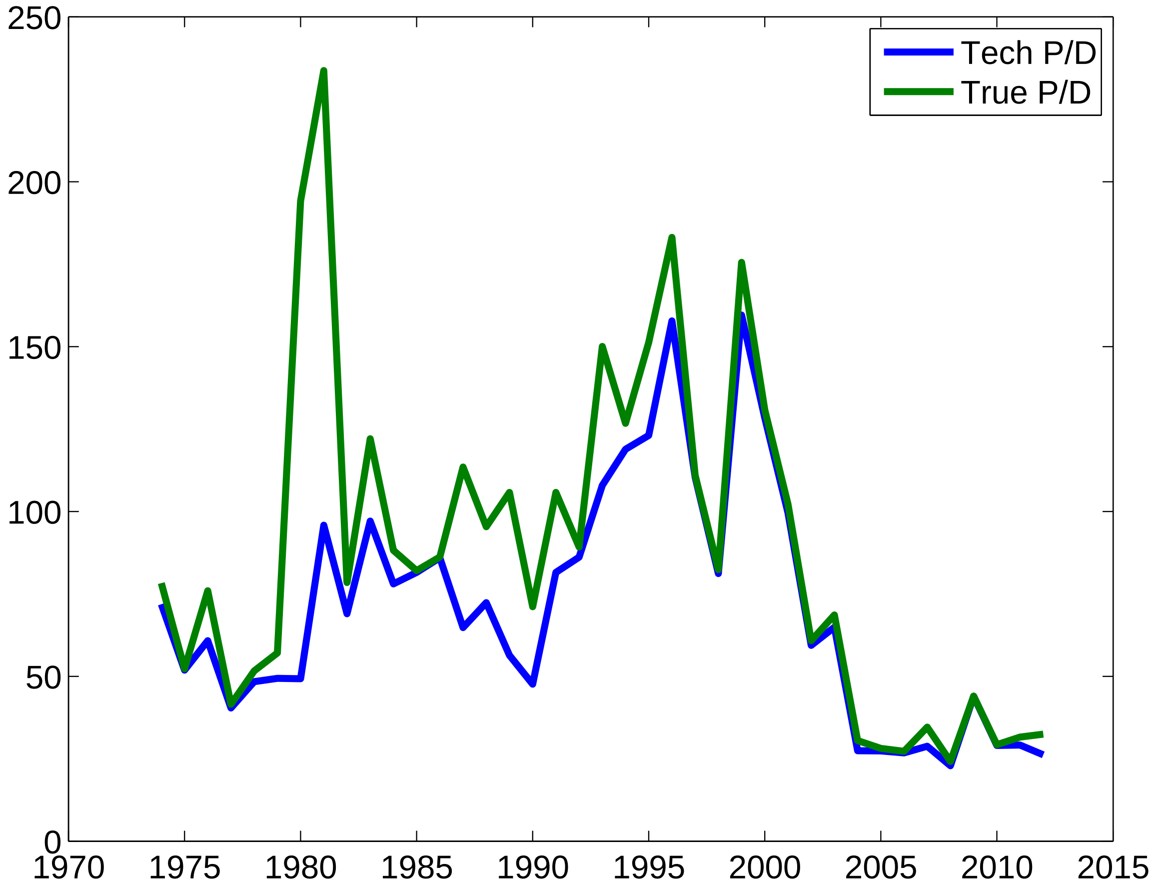

Figure III: IT-capital ratio

This figure plots the IT-capital ratio, the ratio of the IT sector's quantity of capital services to the industrial sector's in solid blue. Capital services are direct estimates of factor income which are based on flows derived from constructed constant-quality capital stock indices. The IT sector is defined as the sum of software, hardware, and communications as listed in the Bureau of Economic Analysis; the industrial sector comprises the remaining 62 asset classes. See Jorgenson and Stiroh (2000) for details. A detailed description of the origin of the data for this figure is in Appendix B. I fix an initial IT-capital ratio $\frac{N_0 X_0}{K_0} < \lim_{t \to \infty} \mathbb{E}\left[\frac{N_t X_t}{K_t}\right]$ and calibrate the model's average simulation path (in dashed red) to match the length and curve of the data.

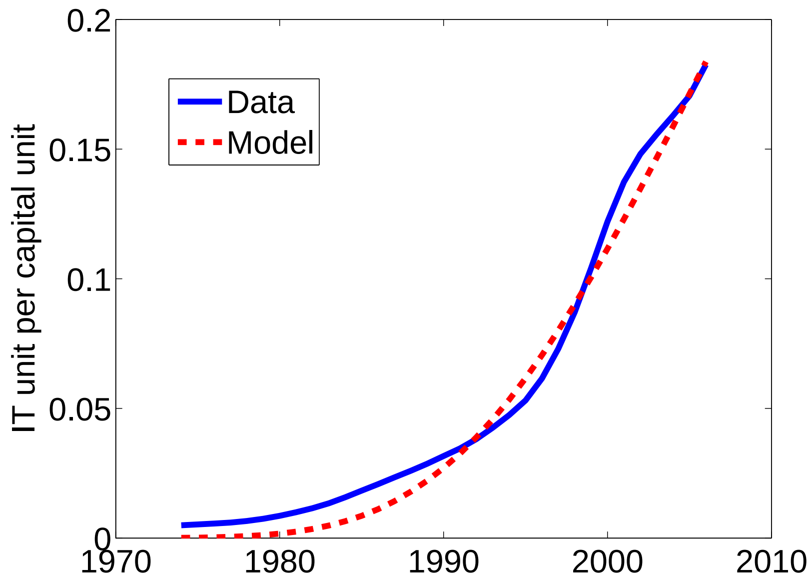

Figure IV: Price-dividend ratios

This figure plots time series of the sectors' price-dividend ratios. The dashed line is the model's average simulation path. The dotted lines are two times the model's standard deviation of point estimates across simulations. The solid blue line is data. The top figure is the IT sector; the bottom is the industrial sector. In the data, I calculate repurchase-adjusted price-dividend ratios, as described in Appendix B. Data are quarterly and are smoothed with a Hodrick-Prescott filter with a smoothing parameter equal to 1600.

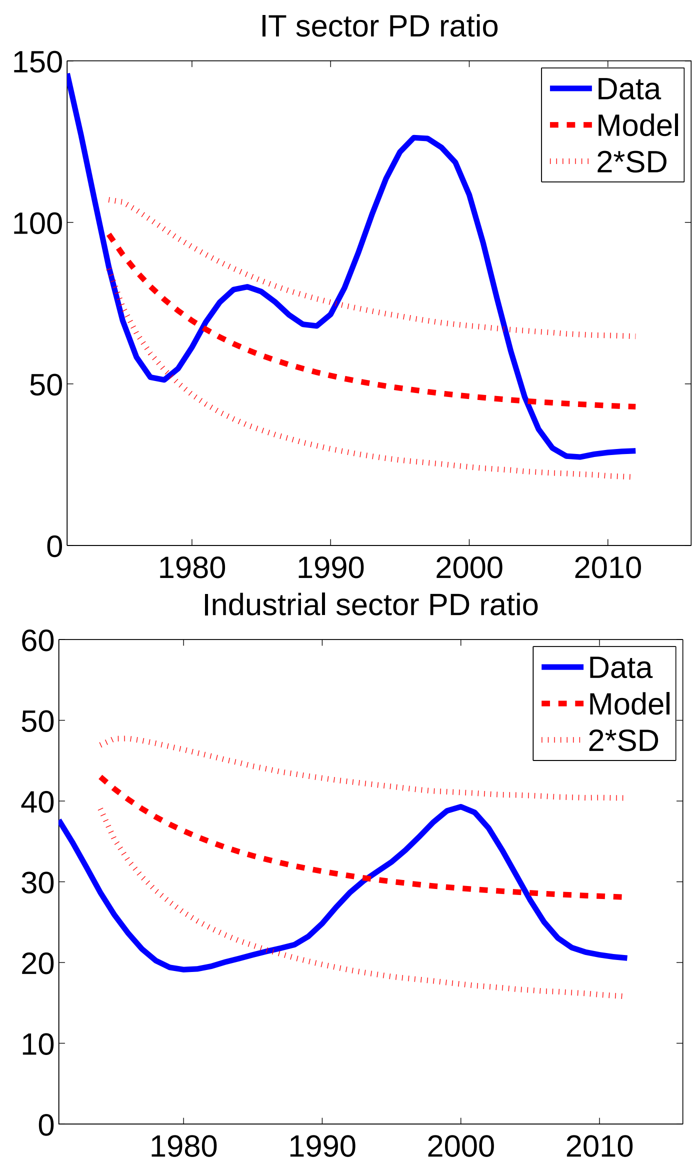

Figure V: IT sector's sales growth rate

This figure plots the average sales growth rate per firm of the IT sector. The model's average simulation path is the dashed line; the variable is $\log\left(\frac{\Pi_{t+1}}{\Pi_t}\right)$. The dotted lines are two times the model's standard deviation of point estimates across simulations. The solid line is data from Compustat: the aggregate sales growth rates per public IT firm (N_t): $\log\left(\frac{y_{t+1}}{y_t}\right)$, where $y_t = \frac{\sum_i^{N_t} Sales_{i,t}}{N_t}$. Data are quarterly and are Hodrick-Prescott filtered with a smoothing parameter of 1600. IT firms are identified by NAICS codes in Appendix B.

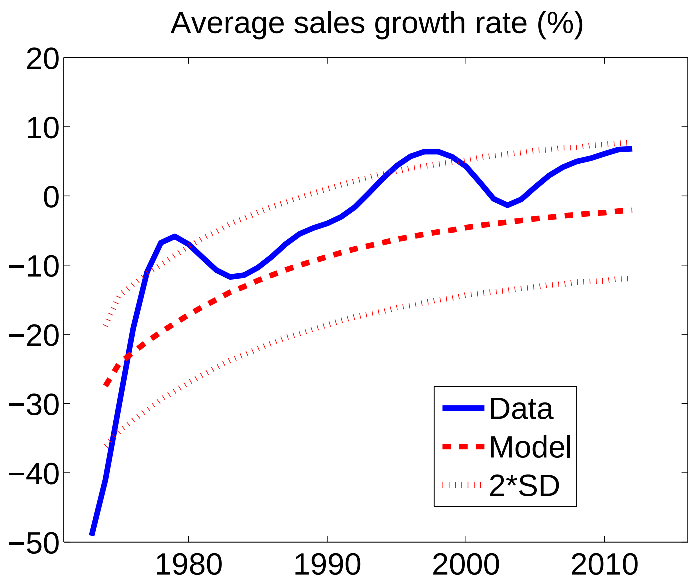

Figure VI: IT sector's net entry rate

This figure plots the net entry rate of the IT sector. The model's average simulation path is the dashed line; the variable is $\log\left(\frac{N_{t+1}}{N_t}\right)$. The dotted lines are two times the model's standard deviation of point estimates across simulations. The solid line is the two-year compound annual growth rate of the growth rate of public IT firms. Data are quarterly and are Hodrick-Prescott filtered with a smoothing parameter of 1600. IT firms are identified by NAICS codes described in Appendix B.

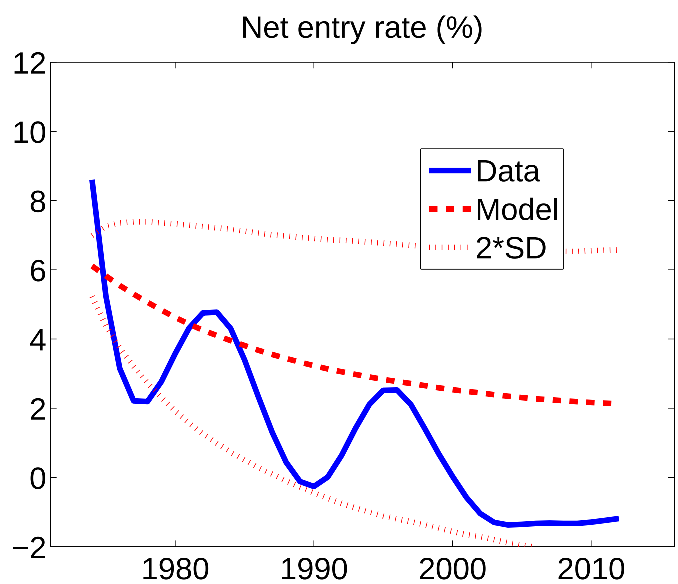

Figure VII: Rolling risk exposures I: 1974-2012 sample

The form of the regression is $r_{IT,t} = a_{IT} + \beta_{IT,cg} \times g_{C,t} + \beta_{IT,rc} \times r_{C,t} + \nu_t$. The risk factors, $g_{C,t}$ and $r_{C,t}$, are standardized. Returns are real. The regressions are rolling and each regression includes 50 quarters of data. The dotted lines are two times the model's standard deviation of point estimates across simulations. Three data series are plotted, depending on how the long-run expected consumption growth of the data was estimated: "Kalman", uses a Kalman filter; "Pred reg" uses a predictive regression; "VAR" uses a vector autoregression. Data estimation details are provided in Section C. Model estimation details are provided in 4.

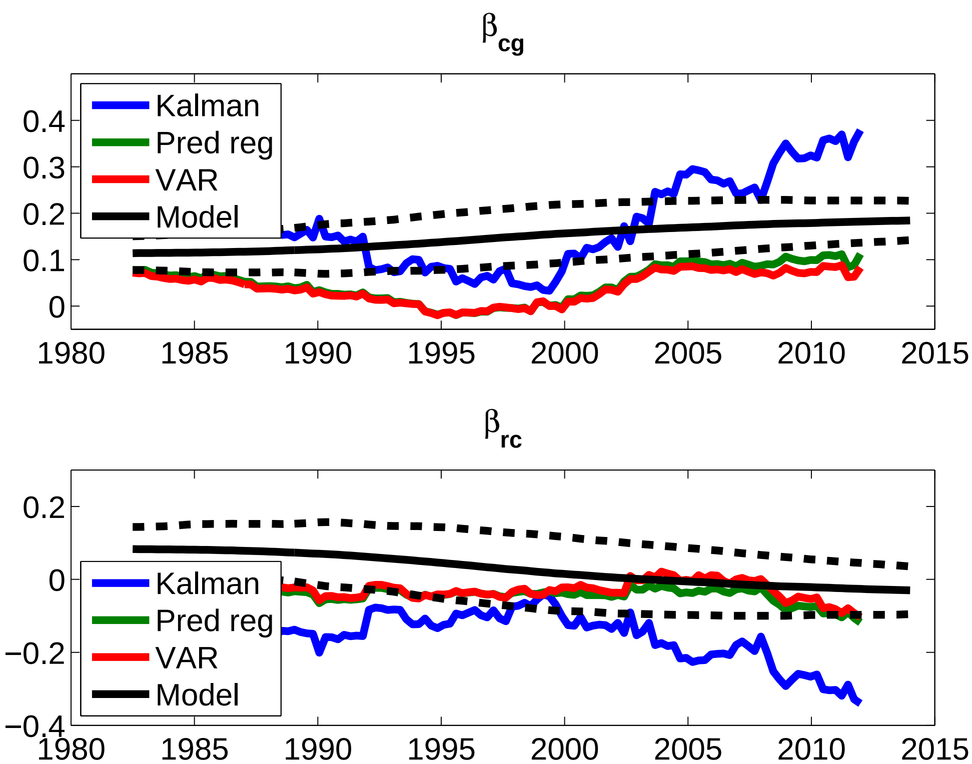

Figure VIII: Rolling risk exposures II: Full sample paths

The form of the regression is $r_{i,t} = a_i + \beta_{i,cg} \times g_{C,t} + \beta_{i,rc} \times r_{C,t} + \nu_t$, where i indexes the IT and industrial sector. The regressions are rolling and each regression includes 50 quarters of data. The risk factors, $g_{C,t}$ and $r_{C,t}$, are standardized in both the model and the data. Returns are real. Model estimation details are provided in Section 4.

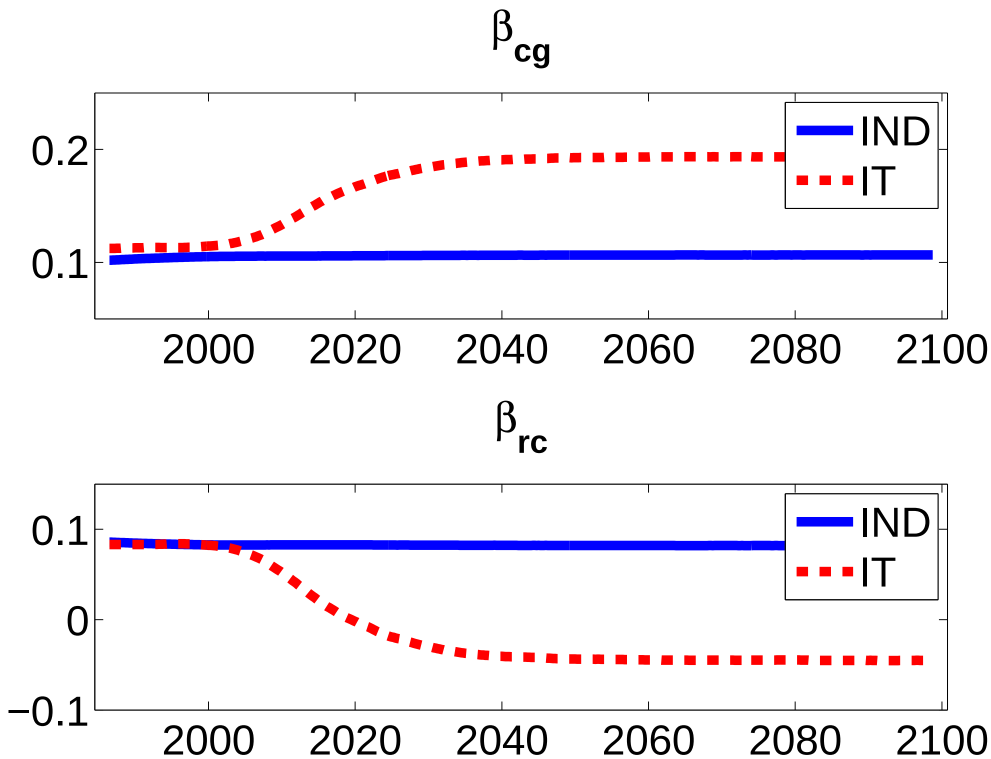

Figure IX: Full transition sample paths

The top-left panel plots the IT-capital ratio. The top-right panel plots the annualized growth rate in measured TFP. The bottom-right panel plots profits made by each IT good producer. The bottom-left panel plots the price-dividend ratios of the IT sector and the industrial sector. The figures below show the average across simulations.

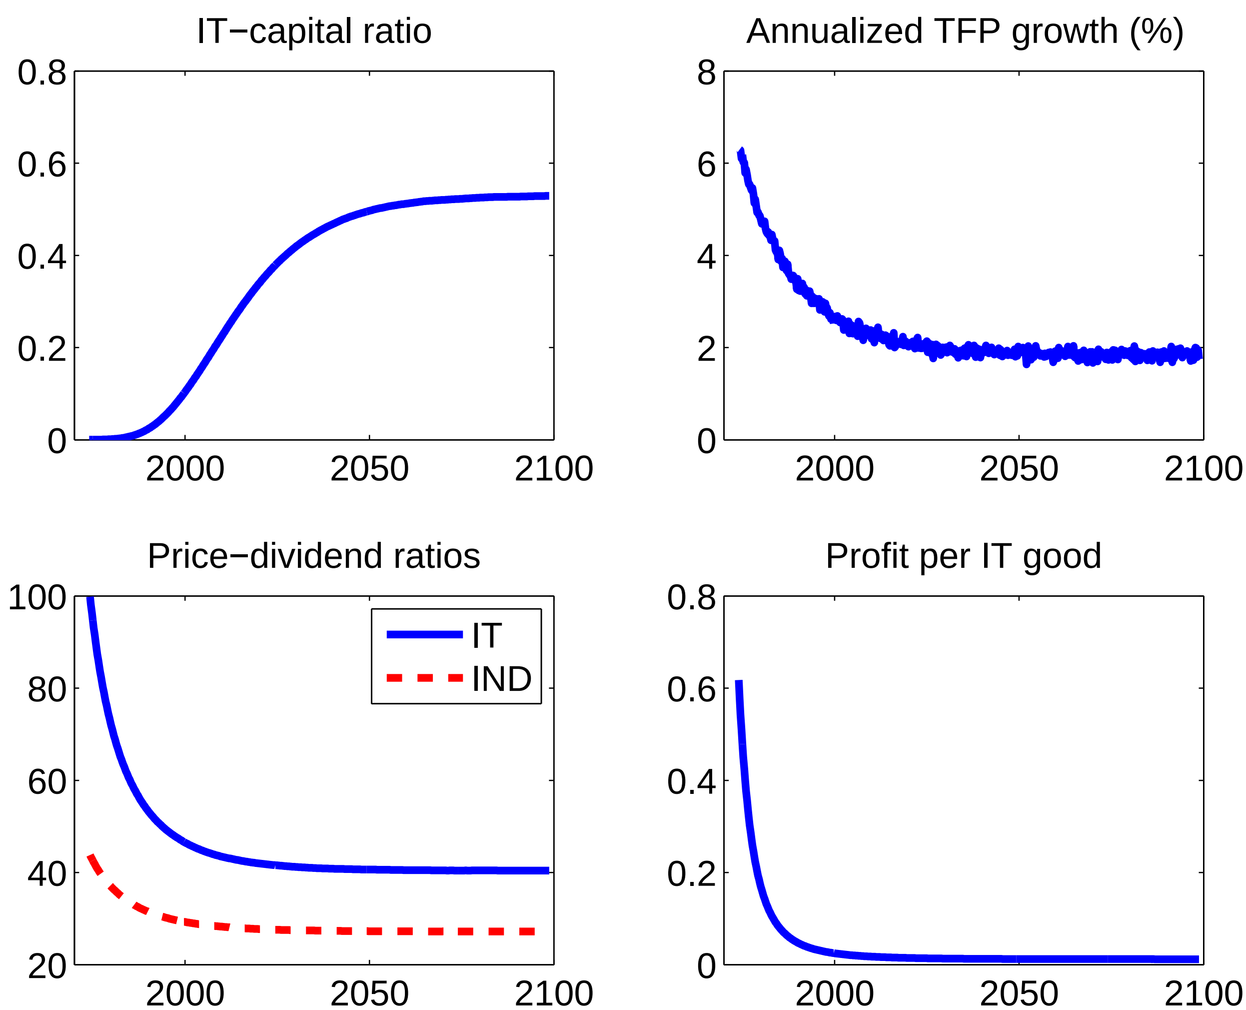

Figure X: Frequency of convergence times

This figure plots the frequency of convergence times, as described in Section 4 under the model being simulated with the parameter vector $\hat{b}$, as described in (17). The convergence time is defined as $\widetilde{T} \doteq \inf\left\{t : \frac{N_t X_t}{K_t} \geq \lim_{t\to\infty} \mathbb{E}\left[\frac{N_t X_t}{K_t}\right]\right\}$.

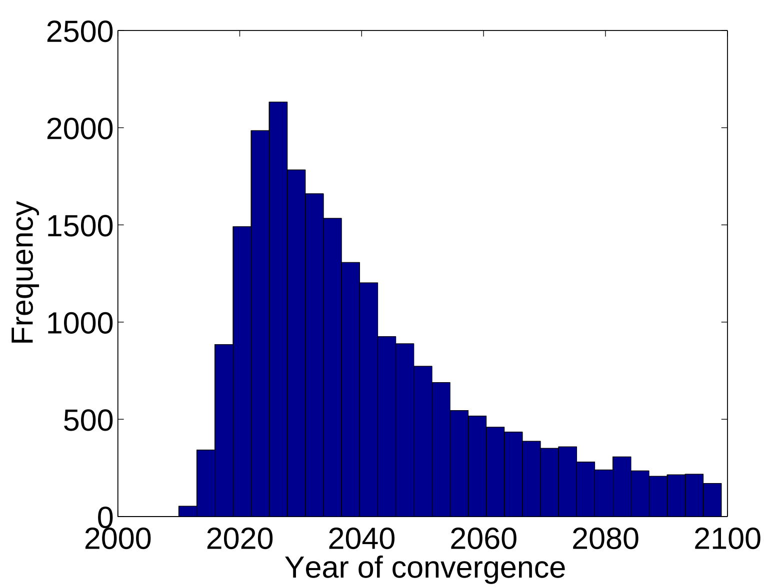

Figure XI: Frequencies of TFP growth per year

This figure plots the frequency of productivity growth per year under the estimated parameter vector $\hat{b}$. The top figure plots the historical amount in the model over the simulated period 1974–2013. The actual amount observed in the data was 0.87 percent TFP growth per year. The bottom figure generates the future TFP growth from 2013 until the convergence time $\widetilde{T}$ is reached. Both cumulative growth rates are divided by the number of years. Details are in Section 4.