University of Pennsylvania
ScholarlyCommons

Publicly Accessible Penn Dissertations

2014

Is the It Revolution Over? An Asset Pricing View

Colin Ward

University of Pennsylvania, wardcol81@gmail.com

Follow this and additional works at: <https://repository.upenn.edu/edissertations>

Part of the [Finance and Financial Management Commons](#)

Recommended Citation

Ward, Colin, "Is the It Revolution Over? An Asset Pricing View" (2014). *Publicly Accessible Penn Dissertations*. 1493.

<https://repository.upenn.edu/edissertations/1493>

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/edissertations/1493>
For more information, please contact repository@pobox.upenn.edu.

Is the It Revolution Over? An Asset Pricing View

Abstract

I develop a new method that puts structure on financial market data to forecast economic outcomes. I apply it to study the IT sector's transition to its long-run share in the US economy, along with its implications for future growth. 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 indicates the transition ends in 2033. I estimate these numbers by building an asset pricing model that endogenously links economy-wide growth to IT sector innovation governed by the sector's market valuation, and by calibrating it to match historical data on factor shares, price-dividend ratios, growth rates, and discount rates. 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.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Finance

First Advisor

Joao F. Gomes

Keywords

asset pricing, endogenous growth, information technology

Subject Categories

Finance and Financial Management

IS THE IT REVOLUTION OVER? AN ASSET PRICING VIEW

Colin Ward

A DISSERTATION

in

Finance

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2014

Supervisor of Dissertation

João F. Gomes, Howard Butcher III Professor of Finance

Graduate Group Chairperson

Eric T. Bradlow, The K.P. Chao Professor

Dissertation Committee

João F. Gomes, Howard Butcher III Professor of Finance

Andrew B. Abel, Ronald A. Rosenfeld Professor

Christian C. Opp, Assistant Professor of Finance

Nikolai Roussanov, Assistant Professor of Finance

Dedicated to Lumina and Michelle

ACKNOWLEDGEMENT

I would like to thank Andy Abel, João Gomes (chair), Christian Opp, and Nick Roussanov for their advice. I benefited from discussions with Gunnar Grass, Brent Neiman, Chris Parsons, Ivan Shaliastovich, Ali Shourideh, Robert Stambaugh, Cecilia Parlatore Siritto, Mathieu Taschereau-Dumouchel, Luke Taylor, Jessica Wachter, and Amir Yaron. I also thank Vincent Glode for many helpful conversations.

ABSTRACT

IS THE IT REVOLUTION OVER? AN ASSET PRICING VIEW

Colin Ward

João F. Gomes

I develop a new method that puts structure on financial market data to forecast economic outcomes. I apply it to study the IT sector's transition to its long-run share in the US economy, along with its implications for future growth. 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 indicates the transition ends in 2033. I estimate these numbers by building an asset pricing model that endogenously links economy-wide growth to IT sector innovation governed by the sector's market valuation, and by calibrating it to match historical data on factor shares, price-dividend ratios, growth rates, and discount rates. 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.

TABLE OF CONTENTS

ACKNOWLEDGEMENT	iii
ABSTRACT	iv
LIST OF TABLES	vi
LIST OF ILLUSTRATIONS	vii
CHAPTER 1 : Is the IT revolution over? An asset pricing view	1
1.1 Introduction	1
1.2 Environment	5
1.3 Deterministic model analysis	21
1.4 Calibration and quantitative analysis	26
1.5 Conclusion	48
APPENDIX	49
BIBLIOGRAPHY	64

LIST OF TABLES

TABLE 1 :	TFP-forecasting regressions I	20
TABLE 2 :	TFP-forecasting regressions II	22
TABLE 3 :	Calibration	29
TABLE 4 :	Estimates of η_s	31
TABLE 5 :	IT sector markups	32
TABLE 6 :	Price-dividend ratios	35
TABLE 7 :	Model and data macroeconomic statistics	41
TABLE 8 :	Asset pricing moments	42
TABLE 9 :	Estimates of return on wealth (long-run risk)	62

LIST OF ILLUSTRATIONS

FIGURE 1 :	Deterministic model: Price-dividend ratio plots	27
FIGURE 2 :	Model calibration I: Input ratio	34
FIGURE 3 :	Model calibration II: Price-dividend ratios	36
FIGURE 4 :	Model calibration III: IT sector's average sales growth rate	37
FIGURE 5 :	Model calibration IV: IT sector's firm-count growth rate	38
FIGURE 6 :	Model: Rolling risk exposures II	40
FIGURE 7 :	Model: Full transition paths	44
FIGURE 8 :	Model: Density of convergence times	46
FIGURE 9 :	Model: Densities of TFP growth per year	47
FIGURE 10 :	Model calibration V: Rolling risk exposures I	63

CHAPTER 1 : Is the IT revolution over? An asset pricing view

1.1. 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 firm 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, whose intensity is 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 calibrate the model's transition paths to match historical data of factor shares, price-dividend ratios,

¹Cowen (2011) and Gordon (2000, 2012, 2013) are pessimistic, whereas Moore (2003), Brynjolfsson and McAfee (2011), and Byrne, Oliner and Sichel (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>.

²I define the IT sector in Appendix A.2. 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 cases being a relative intensity of factor use.

growth rates, and discount rates. This calibrated model allows me to study the IT sector's temporal evolution toward its long-run factor share and to estimate its future bearing on growth.

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 indicates the IT sector's transition ends in 2033, six decades after its 1974 inception.

I further analyze the model to make two novel predictions about the IT sector's evolution. First, the sector is more likely to reach its long-run share within the decade before 2033 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, my model exhibits a salient equilibrium effect that hastens the transition. Second, the information technology sector serves as a hedge against adverse innovations to expected growth 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.

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:

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

where i denotes the sector; P , the sector's market capitalization; D , its aggregate payout; and g , its dividend growth rate.³ 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

³The growth rate of the sector encompasses not only the growth rate of dividends of individual firms, but also the net increase in the number of firms in the sector: that is, $g^{(i)} = d^{(i)} + n^{(i)}$, with d denoting the sector's per firm dividend growth rate, and n denoting the sector's net entry rate.

reflecting an exceptional dividend growth rate or a growing mass of industry constituents. 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 becomes⁴

$$\frac{P_0^{(2)}}{D_0^{(2)}} = \frac{1}{r - g - \Delta} \left[1 - \frac{e^{-(r-g-\Delta)T}}{r - g} \Delta \right]. \quad (1.1)$$

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 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 allows me to match historical paths of pertinent 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 growth.

Related literature

My paper builds on the work that relates financial market performance to shifts in the

⁴Equation (1.1) 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.

technological frontier.⁵ Pástor and Veronesi (2006, 2009) develop models where learning about 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 (2012b) study the asset pricing implications of large, infrequent technological innovations that require firm-specific 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 market effects of a gradual shift in the technological frontier as the sector expands and transitions towards its long-run factor share. Furthermore, I ultimately use the model to forecast growth.

That said, my paper adds to the literature linking asset prices to aggregate growth to innovations in the economy.⁶ 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 (2012) build a growth model similar to the one used here but focus on the quantitative difference implied by assuming exogenous or endogenous growth; they show the latter performs better in matching macroeconomic and asset market data. Their insight of asset prices reflecting anticipated future growth is one shared in this paper. While their paper features R&D as the chief state variable, my paper places the IT sector’s price-dividend ratio as the centerpiece. The papers can thus be viewed as complementary. Gârleanu, Kogan and Panageas (2012a) study a growth model of innovation in an overlapping-generations economy. They find that innovation increases the competitive pressure of existing firms, similarly to this study, and that a lack of intergenerational risk sharing introduces a new source of “displacement” risk in the economy. The novelty in my work is in the context and the application. I explicitly model the “innovation” sector as the IT sector, and map all model features directly to readily available public market data and investment data. I also focus on the model’s transition paths: I initialize the economy with a small IT sector and

⁵A 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, Papanikolaou, Seru and Stoffman (2012), and Kogan, Papanikolaou and Stoffman (2013).

⁶For 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 (2010).

study its evolution to its larger, long-run share, triangulating the model's transition paths with macroeconomic and asset market data.

Finally, my paper fits into the large literature tying cross-sectional and time series asset returns to production economies.⁷ 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 in 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 1.2 describes the model environment. Section 1.3 builds a deterministic model to highlight its qualitative features being consistent with broad movements in the data. Section 1.4 calibrates, simulates, and analyzes the full model. Section 1.5 concludes.

1.2. Environment

There are two sectors: the industrial sector and the IT sector. The information technology sector houses both a production division and a research division. The industrial sector rents IT goods from the IT sector; these goods enhance the productivity of the industrial sector,

⁷Contributions include Cochrane (1996), Jermann (1998), Berk, Green and Naik (1999), Gomes, Kogan and Zhang (2003), Carlson, Fisher and Giammarino (2004), Kogan (2004), Zhang (2005), Cooper (2006), Novy-Marx (2007), Belo (2010), Gala (2010), Kaltenbrunner and Lochstoer (2010), Gourio (2012), Kogan and Papanikolaou (2013a), Kogan and Papanikolaou (2013b), and Loualiche (2013).

and the greater the variety of goods, the greater the enhancement. Sustained demand for these goods increases the value of them and incentivizes the IT sector to create more of them. The information technology sector conducts research today in anticipation of creating new IT goods tomorrow. The created goods are subsequently rented by the industrial sector, and through this process, growth is endogenized.

The ratio of interest

Consistent with Jorgenson et al.'s (2011) conclusion that the chief contributor to recent economic growth was IT's growing factor share, I focus this paper's analysis on the IT-capital ratio.

$$\textbf{Ratio of interest: the IT-capital ratio} = \frac{N_t X_t}{K_t},$$

where K_t is the quantity of capital and X_t is the industrial sector's quantity of demand for each IT good, of which there is a varied continuum of measure N_t . All quantities will be explicitly defined in what follows.

The basic method to analyze this ratio's transition follows: I start the economy at a small IT-capital ratio relative to its larger, model-implied, long-run value. I then run the model and analyze the ratio's transition, which is governed by the model's dynamics, towards its long-run value. When the ratio nears this 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. I explicitly define this method for the full model in Section 1.4.

1.2.1. Information technology sector

Market structure and product division

The information technology sector is subject to the same market forces as every other

industry, but two critical forces affecting it are cost structures and network economies.⁸ Fixed costs and tiny marginal costs are rarely observed in the industrial economy, but for the IT sector, they are common.⁹ This is not just true for pure information goods, such as ebooks and other media, but even for physical goods such as silicon chips. Constructing a chip fabrication plant can cost several billion dollars, but producing an incremental chip only costs a few dollars. This cost structure cultivates supply-side economies of scale.

The distinguishing feature of a good exhibiting network economies is that the demand for the good depends on how many other people use it.¹⁰ Purchasing a word processor with the largest market share is natural, as it allows you to more easily transfer files, resolve problems online, and work on multi-authored documents. These economies also contribute to another market-power-granting effect: lock-in. Consider learning software. Becoming proficient with a piece of software takes time. Switching to a new piece of software is costly because you will have to relearn computing commands or functions; switching shoes from Nike to Adidas, on the other hand (foot), is trivial. At the organizational level, the effect of lock-in can potentially be huge.

Both forces coalesce to endow producers of IT goods with market power; consequently, I model the IT sector as monopolistically competitive. There is a fixed continuum of IT firms of unit measure that comprise the IT sector. Each IT firm is indexed by j . Information technology goods produced by the whole sector are on a continuum of measure N_t , and are indexed by i . Each firm monopolistically prices its good(s).¹¹ Information technology is

⁸There are several other important forces that affect the IT sector, see Shapiro and Varian (1999) for an excellent overview.

⁹Bakos and Brynjolfsson (1999) study a strategy of bundling a large number of information goods and selling them for a fixed price. 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.

¹⁰Goolsbee 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.

¹¹Thus, the measure of IT goods, N_t , reflects the entire sector's product line. Because any firm produces a zero-measure set of IT 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 IT sector as a whole, I abstract from intra-industry

notorious for depreciating quickly, so I make a further technically simplifying assumption: it depreciates fully every period.

Consider the quantity demanded, $X_t(i)$, which will be later explicitly modeled, for some IT good $i \in [0, N_t]$. The monopolist of the IT good takes $X_t(i)$ as given and sets prices to maximize its profit, subject to a linear production function that is common to all monopolists.¹² My assumptions imply that every monopolist sets the same price, $P_t(i)$, in every period (See Appendix A.1 for the derivation):

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

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

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

Profitability here is kept simple: each IT good producer simply charges a markup over marginal cost, and earns the difference multiplied by the quantity demanded. A shortcoming of the model is that changes in profitability are solely determined by changes in demand, rather than by changes in markups. As an industry evolves in reality, both price and quantity reductions lower incumbent firms' profits (see Jovanovic and MacDonald (1994)). All that matters here, however, is the aggregate amount of profits, and their being procyclical and increasing in the total size of the industry.

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

strategies to substantially simplify the analysis. You can think about this market structure as having the IT sector provide many differentiated products to the industrial sector; for example, the goods could be smart phones, robots, consulting services, and even applications (apps)—any product that enhances productivity.

¹²In detail, the intraperiod dynamic for the IT firm follows. It observes the demand curve and sets its price to maximize profits. From $X_t(i)$ units of IT capital it produces $X_t(i)$ units of the IT good, which it sells to the industrial firm. Note the existing IT capital fully depreciates after production. Therefore, the IT firm uses part of the industrial sector's payment to re-invest $X_t(i)$ units of IT capital for the following period.

obsolete, is no longer demanded by the industrial sector, and thus has zero value. The value for any IT good, V_t , then, can be written recursively as

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

where M_{t+1} is the stochastic discount factor, and $\mathbb{E}[\cdot | \mathcal{F}_t]$ 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 IT goods are expected to have the same value. Thus, information technology firms will conduct relatively more research to create new IT goods when the value of them is high.

Research division

The information technology sector as a whole spends a lot on research and development.¹³ The division for research is contained within the IT sector and is characterized by two conditions. First, any zero-measure IT firm in the IT sector can conduct research subject to a common, decreasing returns to scale technology, parameterized by $\eta_s \in [0, 1]$.¹⁴ 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.

These assumptions jointly lead to a condition where the marginal benefit of research is equated with its marginal cost for every individual firm, each indexed by j :

$$\theta_t \times \eta_s S_t(j)^{\eta_s - 1} \times \mathbb{E}[M_{t+1} V_{t+1} | \mathcal{F}_t] = 1, \forall j \in [0, 1]. \quad (1.2)$$

The left side is the marginal benefit of research, which is interacted with a time-varying

¹³Computing 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). 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. These latter estimates are from my own calculations on Compustat data.

¹⁴Startups in the IT sector have historically been a large beneficiary of venture capital. Metrick and Yasuda (2011, p.18, Exhibit 1-7) document that the four largest recipients of venture capital investment were communications, software, semiconductors and electronics, and hardware, which received 54 percent of total venture capital investment in the post-tech-bubble period.

externality θ_t that is taken as given by an individual firm. The right side is the marginal cost of research. I interpret the externality as a measure of research productivity and implicitly give it the following form:

$$S_t \theta_t = \chi N_t^{1-\eta_s-\eta_k} K_t^{\eta_k}.$$

I choose the scale parameter χ to match the evidence on balanced growth and have S_t denote aggregate research expenditure. I specify the elasticity of new IT good development with respect to research to be η_s , which corresponds with its research production technology. And I index the strength of a capital reallocation friction with the parameter 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 three features:¹⁵

$$\frac{\partial \theta_t}{\partial S_t} < 0, \quad \frac{\partial \theta_t}{\partial N_t} > 0, \quad \frac{\partial \theta_t}{\partial (K_t/N_t)} < 0.$$

The left-most derivative captures decreasing returns to aggregate research expenditure, because some new products, although researched independently, will possibly overlap in function and use. The middle 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-most derivative captures a capital reallocation friction that is plausible for two reasons: one, it is more difficult to integrate IT into more capital, because differences in each type of capital could require a specific approach; two, the accumulation of knowledge about a particular type of capital would reduce the incentive to learn about a new type of capital, as in Atkeson and Kehoe (2007).

Each IT firm chooses to spend $S_t(j)$ units of the final good on research. And the measure ϕN_t remains in the next period. The law of motion for IT goods, then, takes the following

¹⁵A 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.

form:

$$N_{t+1} = \phi N_t + \int_0^1 \theta_t S_t(j)^{\eta_s} dj \quad (1.3)$$

$$= \phi N_t + \theta_t S_t^{\eta_s}, \text{ with } N_0 > 0. \quad (1.4)$$

Together, (1.2) and (1.4) imply an aggregate condition:

$$S_t = (N_{t+1} - \phi N_t) \eta_s \mathbb{E}[M_{t+1} V_{t+1} | \mathcal{F}_t]. \quad (1.5)$$

The left side denotes aggregate research expenditure and the right side summarizes the aggregate benefit of conducting research today: the increment of novel IT goods ($N_{t+1} - \phi N_t$) multiplied by each good's discounted expected value $\mathbb{E}_t[M_{t+1} V_{t+1}]$ multiplied by the share of research revenue expensed during development, η_s .

Plugging (1.5) into (1.4) (and temporarily holding $\eta_k = 0$ for clarity) highlights the model's crucial feature—a tight link between IT-sector growth and the valuation of its goods:

$$1 + g_{N,t+1} \equiv \frac{N_{t+1}}{N_t} = \phi + \chi^{\frac{1}{1-\eta_s}} (\mathbb{E}_t[M_{t+1} V_{t+1}])^{\frac{\eta_s}{1-\eta_s}}, \quad (1.6)$$

where $g_{N,t+1}$ is defined as the IT sector's net entry rate. This equation clearly 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. Thus, my model intimately links innovation and growth to entry, an result consistent with empirical evidence presented by Jovanovic and MacDonald (1994).

1.2.2. Industrial sector

Production function

The industrial sector comprises competitive firms that are identical. Because all firms are

identical, the economy admits a representative firm. The representative firm produces a final good Y_t by combining capital K_t , a composite IT good G_t , and labor L_t , which is subject to a productivity shock A_t :

$$Y_t = \left(K_t^\alpha (A_t L_t)^{1-\alpha} \right)^{1-m} G_t^m, \quad (1.7)$$

where m denotes the share of IT goods in factor income, and α the capital share of non-IT good factor income.¹⁶ I normalized the price of the final good to one. The production function specifies capital, IT goods, and labor as having positive cross-partial derivatives. Consequently, by renting more IT goods the marginal product of the two traditional inputs of production, capital and labor, are enhanced.¹⁷

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 of substitution aggregator

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

The parameter μ measures the degree of variety that each IT good possesses. As μ goes to one, all IT goods are perfect substitutes, and, furthermore, the incentive of the IT sector to conduct research goes to nil; consequently, growth is not sustained. Thus, I maintain the restriction that $\mu > 1$. A result of this restriction is that the industrial firm is more productive if, for example, it uses an equal amount of two IT goods versus if it uses twice as much of one IT good. Thus, when new IT goods are created, it is in the final-good producer's interest to diversify existing demand and include the new spectrum of goods, and to reduce the quantity demanded of each specific IT good. The variable G_t can be

¹⁶The production function is later rewritten as only a parametric function of α , with that share going to capital, and the balance going to labor; see (1.14). This is the more customary interpretation taken by the literature.

¹⁷Plant-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 productivity. Black and Lynch (2001) additionally discover that greater computer usage by nonmanagerial employees raises plant productivity.

thought of as measuring the technological complexity of the final-good producers.

Capital accumulation

I subject the accumulation of capital to Penrose-Uzawa adjustment costs, along the lines of Jermann (1998), with current capital depreciating at rate δ

$$K_{t+1} = (1 - \delta)K_t + \Lambda \left(\frac{I_t}{K_t} \right) K_t, \text{ where } \Lambda \left(\frac{I_t}{K_t} \right) = c_0 + \frac{c_1}{1 - \frac{1}{\zeta}} \left(\frac{I_t}{K_t} \right)^{1 - \frac{1}{\zeta}}. \quad (1.8)$$

The free parameters (c_0 and c_1) of the adjustment cost function are chosen to eliminate adjustment costs in the deterministic steady state (following Kaltenbrunner and Lochstoer (2010)).¹⁸ 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 capital adjustment costs are high; as $\zeta \rightarrow \infty$, marginal capital adjustment costs go to zero.

Stationary productivity

In addition, the exogenous source of total factor productivity (TFP) of the firm A_t follows a stationary Markov process¹⁹

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

where ϵ_{t+1} is an independently and identically distributed normal random variable with mean zero and constant variance σ_ϵ^2 . I set the autoregressive coefficient, ρ , near one, making exogenous productivity persistent. This is a common assumption in the Real Business Cycle (RBC) literature, one used to generate business cycles. Because this process is stationary, long-run growth only occurs endogenously through the IT sector's expansion.

¹⁸Explicitly, $c_0 = \frac{1}{1-\zeta}(g_N^* + \delta)$ 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.

¹⁹The Hicks-Neutral measurement of TFP would actually be $A_t^{(1-\alpha)(1-m)}$.

Maximization

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

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

where $M_{t|t+s} = M_{t+1} \cdot M_{t+2} \cdots M_{t+s}$ is the product of future stochastic discount factors from time $t+1$ to $t+s$. The firm's optimality conditions are in Appendix A.1. I can simplify the first-order condition with respect to $X_t(i)$ because of the IT sector's market structure:

$$X_t = \left(\frac{m}{\mu} \right)^{\frac{1}{1-m}} K_t^\alpha (A_t L_t)^{1-\alpha} N_t^{\frac{\mu m - 1}{1-m}}. \quad (1.9)$$

Because A_t is procyclical and persistent, the demand for IT goods, the valuations of these goods, the IT sector's aggregate expenditure on research, and the entry rate of new IT firms are, too.

Output and balanced growth

Using equilibrium conditions, I can rewrite (1.7) as

$$Y_t = \left(\frac{m}{\mu} \right)^{\frac{m}{1-m}} K_t^\alpha (A_t L_t)^{1-\alpha} N_t^{(\mu-1)\frac{m}{1-m}}. \quad (1.10)$$

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

$$\alpha + (\mu - 1) \frac{m}{1-m} = 1. \quad (1.11)$$

From of this restriction, (1.9) implies that X_t is decreasing in N_t , which would be consistent with a competition effect driving down the profits earned by each firm in the IT sector as

its size expands.

1.2.3. Resource constraint, households, and the steady state

Resource constraint

The final good is used for consumption, and for investment in capital, IT, and research:

$$Y_t = C_t + I_t + N_t X_t + S_t. \quad (1.12)$$

Thus, total investment in this economy is the sum of capital investment, I_t , the total investment of the IT sector, $N_t X_t$, including its aggregate expenditure on research, S_t .

Households

The economy is populated by a competitive representative household that derives utility from the consumption flow of the single consumption good C_t . It supplies labor perfectly inelastically, so $L_t = 1$ for all t ; I focus on analyzing the sectors' capital quantities and valuations, not movements in labor supply. The representative household maximizes the discounted value of future utility flows with Epstein and Zin (1989) and Weil (1989) recursive preferences:

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

where γ is the coefficient of relative risk aversion, ψ is the elasticity of intertemporal substitution, and $\vartheta = \frac{1-\gamma}{1-1/\psi}$ is defined for convenience. 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.

This setup implies that the stochastic discount factor in the economy is given by

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \left(\frac{\mathbb{E}_t[U_{t+1}^{1-\gamma}]}{U_{t+1}^{1-\gamma}} \right)^{\frac{\gamma-1/\psi}{1-\gamma}}, \quad (1.13)$$

where the first term is the 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 stochastic dividend $\mathcal{D}_t$. The budget constraint of the household is

$$C_t + \mathcal{S}_{t+1}Q_t + B_{t+1} = W_t L_t + (Q_t + \mathcal{D}_t)\mathcal{S}_t + R_t B_t,$$

where $S_t Q_t$ is the aggregate market capitalization and R_t is the gross real rate of interest.

Steady state

I cannot solve the system in closed form, but I can find three variables, $k^* \equiv \left(\frac{K}{N}\right)^*$, $s^* \equiv \left(\frac{S}{N}\right)^*$, and g_N^* , that solve a system of three nonlinear equations:²⁰

$$\begin{cases} 1 &= M^* \left\{ \alpha(1-m) \left(\frac{Y}{K}\right)^* + (1-\delta) \right\} \\ 1 + g_N^* &= (s^*)^{\eta_s} (k^*)^{\eta_k} + \phi \\ s^* &= M^* V^* (1 + g_N^* - \phi) \end{cases}.$$

1.2.4. The valuation-productivity link

Productivity

²⁰For completeness, $V^* = \Pi^* / (1 - \phi M^*)$, $\Pi^* = (\mu - 1) \left(\frac{m}{\mu}\right)^{\frac{1}{1-m}} k^*$, $M^* = \beta(1 + g_N^*)^{-\frac{1}{\psi}}$, and $\left(\frac{Y}{K}\right)^* = \left(\frac{m}{\mu}\right)^{\frac{m}{1-m}} (k^*)^{\alpha-1}$.

Given the restriction in (1.11), we can rewrite (1.10) as

$$Y_t = \left(\frac{m}{\mu}\right)^{\frac{m}{1-m}} K_t^\alpha (A_t N_t L_t)^{1-\alpha} \equiv K_t^\alpha (Z_t L_t)^{1-\alpha}. \quad (1.14)$$

Thus, our usual measurement of productivity is the product of two components: the exogenous, stationary component, and the endogenous, increasing mass of IT goods; call this product $Z_t \equiv \left(\frac{m}{\mu}\right)^{\frac{m}{1-m}} A_t N_t$. 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.

Focusing on this product, a straightforward derivation for an arbitrary horizon h gives

$$\mathbb{E}_t \left[\log \left(\frac{Z_{t+h}}{Z_t} \right) \right] = (\rho^h - 1) \log(A_t) + \mathbb{E}_t \left[\log \left(\frac{N_{t+h}}{N_t} \right) \right]. \quad (1.15)$$

Because A_t is persistent, the first term on the right-hand-side of (1.15) is near zero (and slightly negative). What does this last equation say? It says that the conditional expectation of the IT sector's growth rate is key for forecasting future productivity.

Stock-market valuation

The value of the stock market $\mathcal{S}_t Q_t$ includes both the IT sector and industrial sector. I normalize the aggregate supply of stock $\mathcal{S}_t$ to one. 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:

$$\mathcal{D}_t = \underbrace{D_t}_{\text{Industrial div.}} + \underbrace{N_t \Pi_t - S_t}_{\text{IT div.}}.$$

This leads to the following observation, whose proof is in Appendix A.1, and which is central

to the paper.

Proposition 1 (Stock market valuation). *The aggregate value of the stock market in this economy is*

$$Q_t = \underbrace{q_t K_{t+1}}_{\text{Industrial sector}} + \underbrace{\frac{N_t(V_t - \Pi_t)}{}}_{\text{Assets-in-place}} + \underbrace{O_t}_{\text{Growth options}}, \quad (1.16)$$

IT sector

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

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

The value of the stock market, therefore, incorporates three elements: the replacement cost of installed capital, evaluated at the ex-dividend value of capital $q_t K_{t+1}$, and the value of IT firms, comprised of part assets-in-place and part growth options. Important is the addition of the state variable $N_{t+s} V_{t+s}$ to the third term. The value of the stock market contains information about the future level of $N_{t+s} X_{t+s}$, the future demand of IT, and correspondingly the future stock of IT capital. Indeed, if the variety of IT goods is expected to expand, then a dear current market valuation can be justified, even if current IT-capital ratio is low. In the absence of a growing measure of IT goods, the value of the IT sector would be simply $N_t(V_t - \Pi_t)$, the sector's current ex-dividend value.

Price-dividend ratio

From (1.16), I define the IT sector's price-dividend ratio:

$$PD_t^{IT} \equiv \frac{N_t(V_t - \Pi_t) + O_t}{N_t \Pi_t - S_t} \quad (1.17)$$

$$= \frac{\mathbb{E}_t \left[\sum_{s=1}^{\infty} M_{t|t+s} \{\Pi_{t+s} N_{t+s} - S_{t+s}\} \right]}{N_t \Pi_t - S_t} \quad (1.18)$$

The last equation shows the information contained in the price-dividend ratio about economy-

wide future growth. As the measure of IT goods is expected to grow, and thus contribute to the productivity of the industrial sector, the price-dividend ratio of the IT sector encodes this information and consequently will be high, all else equal.

Predictive regression

The tight relationship between the decomposition of productivity in (1.15) and the IT sector's price-dividend ratio in (1.18) is a strong prediction of the model. Does it hold in the data? I test this prediction with the following regression:

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

where $TFP_{t \rightarrow t+h} = TFP_{t+1} + \dots + TFP_{t+h}$ 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 A.2). 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 1 in two panels. Panel A documents economic significance. 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. But the real economic significance is ascertained from the last column, which reports the expected change in TFP for a one standard-deviation change in the IT sector's price-dividend ratio. 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. Even more 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.

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 t -statistics, generated by bootstrapping 10,000 samples of the long horizon regression under the null of no predictability.²¹

Table 1: TFP-forecasting regressions I

The regression equation is $TFP_{t \rightarrow t+h} = a + b \times PD_t^{IT} + \epsilon_{t \rightarrow t+h}$. The dependent variable $TFP_{t \rightarrow t+h}$ 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, from 1971Q1–2012Q4. Panel A's standard errors 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 Hodrick (t_H), Newey-West (t_{NW}), and a Monte Carlo bootstrap method (t_{MC}), developed by Kilian (1999) and used in Goyal and Welch (2008). Data sources and definitions for the IT sector are detailed in Appendix A.2.

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	$t_H(b)$	$t_{NW}(b)$	$t_{MC}(b)$
1 year	0.01	1.0	2.5	1.7
2 year	0.03	2.1	4.1	4.2
3 year	0.04	3.4	3.4	7.1
4 year	0.05	4.7	6.7	9.4
5 year	0.06	4.9	6.7	10.8

I further analyze this prediction in Table 2. Consistent with a research expenditure affecting the economy with a lag, the effects of the price-dividend ratio are stronger at longer horizons.

²¹This 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.

Horse races are also run between the IT sector's and the industrial sector's price-dividend ratios. The IT sector's price-dividend ratio drives out the industrial sector's when comparing the two measures over all horizons.

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

1.3. Deterministic model analysis

Before proceeding to the calibrated quantitative analysis, I consider a simplified version of the model to show that its dynamics are consistent with broad movements in the data. This version cannot accurately match the data, so I calibrate the full model that can in Section 1.4. All insight that follows carries over to the full model.

The simplifications follow:

- The economy is non-stochastic
- The representative household is risk neutral
- Capital readjustment costs are nonexistent
- The industrial firm replaces depreciated capital, adjusting for growth

The first and second simplifications allow for a closed-form solution of the model. A deterministic economy sets $A_t = 1$ for all t . Risk neutrality sets the stochastic discount factor to a constant $M_{t+1} = \beta = \frac{1}{1+r}$ for all t , where r can be interpreted as the real interest rate. The third simplification sets $\eta_k = 0$. Finally, the last simplification sets $I_t = (\delta + g_N^*)K_t$ for all t , making the dynamic between k_t and k_{t+1} simple and putting the focus on the IT

Table 2: 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, from 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 A.2.

Horizon h	Statistic	TFP			Adjusted TFP		
		(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

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

sector's growth rate:

$$k_{t+1} \frac{N_{t+1}}{N_t} = k_t(1 + g_N^*). \quad (1.19)$$

Steady state

The simplifications and (1.9) imply

$$\frac{N_t X_t}{K_t} = \left(\frac{m}{\mu} \right)^{\frac{1}{1-m}} \left(\frac{K_t}{N_t} \right)^{\alpha-1}, \quad (1.20)$$

and so the analysis can simply focus on the ratio $k_t \equiv \frac{K_t}{N_t}$, which is an inverse mapping of the IT-capital ratio.

From (A.2) the steady-state optimality condition for the industrial firm's (normalized) capital choice can be rearranged to give

$$\left(\frac{K}{N} \right)^* \equiv k^* = \left(\frac{\alpha(1-m) \left(\frac{m}{\mu} \right)^{\frac{m}{1-m}}}{r + \delta} \right)^{\frac{1}{1-\alpha}}. \quad (1.21)$$

Putting (1.20) and (1.21) together gives a simple equation for the steady-state IT-capital ratio:

$$\left(\frac{NX}{K} \right)^* = \frac{(r + \delta) \left(\frac{m}{1-m} \right)}{\alpha \mu}.$$

Increases in the user cost of capital ($r + \delta$) make capital more expensive to hold and hence increase the steady-state IT-capital ratio. As m increases, IT goods make up a larger share of production, and thus increases the ratio. For the opposite reason, increasing the importance of capital in production, α , decreases the ratio. Increasing μ effectively increases

the price paid—the user cost—for IT goods, and will thus decrease the ratio.²²

Transition analysis

In this paper, I relate stock prices to the future growth rates of the economy. I structure the analysis by starting the economy at an IT-capital ratio $\frac{N_0 X_0}{K_0}$ lower than its steady-state value $\left(\frac{NX}{K}\right)^*$ and then by running and observing the system's dynamics as it converges to its steady state:

$$\frac{N_0 X_0}{K_0} < \left(\frac{NX}{K}\right)^* \Leftrightarrow k_0 > k^*.$$

At this point, define the first time T when k_t is in the ε -neighborhood of k^* and where k_t , from time T on, will be treated as approximately equal to k^* the following period:²³

$$T = \inf \{t : |k_t - k^*| \leq \varepsilon\} \text{ and } k_t \approx k^*, \forall t > T.$$

The time T refers to how long the transition takes to get within an epsilon of the steady state. It also introduces an element to the analysis that would otherwise be absent because k_t would asymptotically approach (and never reach in finite time) k^* . In addition, it buys a decomposition of the ex-dividend value of an IT good at time t :

$$V_t - \Pi_t \approx (\mu - 1) \left(\frac{m}{\mu}\right)^{\frac{1}{1-m}} \left(\underbrace{\sum_{s=1}^T \beta^{s+1} k_{t+s}^\alpha}_{\text{Transition path}} + \underbrace{\beta^{T+1} \frac{(k^*)^\alpha}{1-\beta}}_{\text{Steady state}} \right). \quad (1.23)$$

²²Solving the present value of an IT good in the steady state gives

$$V^* = \frac{\Pi^*}{1-\beta} = \frac{(\mu-1)X^*}{1-\beta} = \frac{(\mu-1) \left(\frac{m}{\mu}\right)^{\frac{1}{1-m}} (k^*)^\alpha}{1-\beta}. \quad (1.22)$$

Plugging the above equation into (1.6) gives the economy's steady-state growth rate,

$$g_N^* = \phi + \chi^{\frac{1}{1-\eta_s}} (\beta V^*)^{\frac{\eta_s}{1-\eta_s}} - 1.$$

²³The definition of T is actually a function of epsilon— $T(\varepsilon)$ —but for brevity this dependence is ignored.

Hence, the ex-dividend value of an IT good reflects information about the duration of the transition path $(T - t)$ and the distance to the steady state $(|k_t - k^*|)$.

From these simplifications, I present a proposition that summarizes the model's salient properties (see Appendix A.1 for details):

Proposition 2 (Deterministic transition dynamics). *Consider starting the economy at $k_0 > 0$ and assume there exists a sequential bound on growth, $\{\bar{g}_{N,t+1}\}_{t=0}^{\infty}$, then*

- *The system $\mathbf{x}_t \equiv \{k_t, g_{N,t+1}, V_t\}$ converges monotonically to its steady state $\mathbf{x}^*$*
- *The value of an IT good is an increasing function of both T and k_t for all t :*

$$V_t(\tilde{T}) - V_t(T) > 0, \text{ for } \tilde{T} > T \qquad \frac{\partial V_t}{\partial k_t} > 0, \forall t$$

The sequential bound on growth simplifies the proof and is consistent with the dynamics of the full model. Its definition is in Appendix A.1. The intuition follows for the case $k_0 > k^*$. Because an IT good's value is increasing in all future discounted k_t 's, its initial value is greater than its steady-state value. This incentivizes the IT sector's research division to develop relatively more new IT goods. Because aggregate research expenditure has decreasing returns, k_t does not immediately reach its steady state: decreasing returns act like a variable adjustment cost, inducing a multi-period transition. As the measure of IT goods expands tomorrow, next period's k_t decreases (and therefore next period's $\frac{N_t X_t}{K_t}$ increases), reducing an IT good's value. This process repeats until $k_t \leq k^* + \varepsilon$, at which point the economy reaches its steady state in the following period and remains there.

The proposition's second bullet point implies that increasing either T or k_0 increases the value of an IT good. The difference is subtle, and clarifies the model's use in disentangling the two effects. Figure 1 plots the price-dividend ratio of the IT sector while varying either the distance to the steady state $(k_t - k^*)$ or the duration of the transition to the steady state (T) . The top panel shows the price-dividend ratio as a function of T . The ratio falls

asymptotically until it nears the end of the transition, at which point the ratio starkly falls. The bottom panel varies the distance between the initial IT-capital ratio and its steady state. The transition is relatively more gradual for varying distances. The point to take from this exercise is that a sharply falling price-dividend ratio signals that the end of the transition period is near.

The transition paths of this deterministic model are qualitatively consistent with the data. But the full model presented in the next section will also be quantitatively consistent.

1.4. Calibration and quantitative analysis

I present the model analysis in two parts. In the first part, I calibrate the model to match historical data over the period 1974–2012. I do this in two steps:

1. Fix an initial IT-capital ratio $\frac{N_0 X_0}{K_0}$ near the 1974 data point
2. Simulate an entire shock sequence $\{A_t\}_{t=1,2,\dots}$ many times for a given set of parameters
 - For each simulation, compute model quantities and prices
 - Average the model's output across simulations and match it to the data

In step one, I pick the initial IT-capital ratio to also match the data's price-dividend ratios of both sectors. I calibrate the model in step two to agree with informative asset pricing data: price-dividend ratios, growth rates, and discount rates. This method puts structure on financial market data that is consistent with the underlying macroeconomic quantities. I can then use the model's structure to observe the remaining, and currently unobserved, dynamics.

That said, in the second part I analyze the model's entire transition path, which includes the years 1974 and runs until the IT-capital ratio hits its long-run share at time T , which will now be defined. Consider a *convergence time* T adapted to our filtration $\mathcal{F}_t$ that is defined at

Figure 1: Deterministic model: Price-dividend ratio plots

Both panels plot the price-dividend ratio of the IT sector as a function of two of its arguments: the length of the transition path (T) and the distance of the model's input ratio to its steady state ($k_t - k^*$). The top panel varies T while the rest of the model is held constant. The bottom panel varies D such that $k_0 = k^* \times D$ while the rest of the model is held constant.

the first moment when the economy's IT-capital ratio has crossed its unconditional expected value from below:

$$\begin{aligned}
T &\equiv \inf \left\{ t : \frac{N_t X_t}{K_t} \geq \mathbb{E} \left[\frac{N_t X_t}{K_t} \right] \right\} \\
&= \inf \left\{ t : \frac{N_t X_t}{K_t} \geq \left(\frac{m}{\mu} \right)^{\frac{1}{1-m}} \left[\left(\frac{K}{N} \right)^* \right]^{\alpha-1} \exp \left\{ \frac{1}{2} (1-\alpha)^2 \frac{\sigma_\epsilon^2}{1-\rho^2} \right\} \right\}. \quad (1.24)
\end{aligned}$$

1.4.1. Calibration

Table 3 summarizes the choice of parameters. The model is calibrated at a quarterly frequency. The equilibrium is computed numerically using a high-order perturbation method (Schmitt-Grohe and Uribe (2004)) that takes into account the high volatility of stock market prices. Note that the calibration here is unorthodox: we do not observe the entire time series with 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.

Table 3: Calibration

Parameter	Value	Description	Source/Target
Household			
β	0.9915	Subjective discount factor	Match levels of PD ratios
ψ	0.9	Elasticity of intertemporal substitution	Match β risk exposures and growth rates
γ	10	Coefficient of relative risk aversion	Match risk premia
Industrial			
α	0.3	Non-IT capital share of factor income	Literature
δ	0.02	Capital depreciation rate	Literature
ζ	1.01	Adjustment cost parameter	Literature
ρ	0.9945	Exogenous TFP persistence	Match macro autocorrelation
σ_ϵ	0.0175	Productivity volatility	Match macro volatility
IT			
χ	1.61	Scale parameter	Match balanced growth evidence
$\mu - 1$	0.144	IT net markup	Compustat, Goeree (2008)
ϕ	0.95	IT good/firm survival rate	Compustat (default), BEA (depreciation), Li (2012)
η_s	0.855	Elasticity of new IT goods wrt R&D	CRSP/Compustat (PD), SF Fed (TFP), Griliches (1990)
η_k	0.24	Capital reallocation friction	Match transition path, PD ratio transitions, growth rates

The calibrated parameters imply that the steady-state IT-capital ratio, $\mathbb{E} \left[\frac{N_t X_t}{K_t} \right]$ is 0.44, so for every 100 units of industrial capital, there 44 units of IT capital. The model's median convergence time is 2033, and puts the revolution's duration at 60 years.²⁴

Information technology sector

The six parameters here to be calibrated are χ , ϕ , η_s , η_k , μ , and m . They are discussed in turn. I set the scale parameter χ to 1.61 to match balanced growth evidence and generate an annual consumption growth rate of two percent.

The rate of obsolescence of an IT good $1 - \phi$ in the model should capture two features: a high rate of economic obsolescence and default, as weaker firms without competitive advantages 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 A.2. Both methods produce results near 3 percent. Because ϕ is interpreted as a measure, I assume economic obsolescence and delisting are independent and add the two measures together to get $1 - \phi^{\text{Annual}} = 16.5 + 3 = 19.5$ percent, or nearly $\phi = 0.95$ at a quarterly frequency.

To estimate η_s , I approximate (1.6) to get

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

and then substitute this equation into (1.14) to yield

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

²⁴The 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.

Table 4: Estimates of η_s

This table estimates the parameter η_s from the data by running the following regression: $TFP_{t \rightarrow t+h} = a + b \log PD_t^{IT} + \epsilon_{t+h}$. The dependent variable $TFP_{t \rightarrow t+h}$ 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 (log) price-dividend ratio for the IT sector, which is adjusted for repurchases. Data are quarterly, from 1971Q1–2012Q4. The model counterpart is $\log(Z_{t+1}/Z_t) = (\rho - 1) \log A_t + \frac{\eta_s}{1-\eta_s} \log \mathbb{E}_t[M_{t+1}V_{t+1}] + \epsilon_{t+1}$. The parameter $\hat{\eta}_s$ is retrieved from the estimate of $\hat{b}$ by the equation: $\eta_s(b) = \frac{b}{1+b}$. Ninety-five percent confidence intervals are constructed using the delta method: $se(\hat{\eta}) = \eta'_s(\hat{b}) \left[se(\hat{b}) \right]^2 \eta'_s(\hat{b})$, where $se(\hat{b})$ is computed with the Newey-West (1987) adjustment with three lags. Data are described in Appendix A.2.

	Lower 95%	$\hat{\eta}_s$	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

This resembles a linear regression equation. It can be taken directly to the data to estimate η_s . I provide estimates in Table 4. Because the 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. I pick 0.855.

Estimating the parameter that governs the cost of capital readjustment η_k is difficult. Jovanovic and Rousseau (2002) provide estimates of learning laws, a friction of capital reallocation, for general purpose technologies, like IT, within a range of 0.2 to 0.62. I choose 0.24. I discipline this choice by having this single parameter 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.

The parameter $(\mu - 1)$ governs the average markup charged on IT goods. It is tough to

Table 5: 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 A.2.

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

measure accurately, especially given the IT sector’s heterogeneity of products.²⁵ 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 5) 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. I use 14.3 percent.

The IT share of factor income parameter m determines the importance of IT in the production. This is unknown by construction, because the steady state has not yet been observed. The choice is disciplined, however, by the balanced growth condition in (1.11) which specifies m given μ and α , two parameters that are plausibly easier to measure.

Industrial sector

The parameters here are α , δ , ζ , ρ , and σ_ϵ . 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. The adjustment cost parameter ζ is set to 1.01, which falls in line

²⁵For 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.

with estimation evidence (see, for instance, Jermann (1998), Kaltenbrunner and Lochstoer (2010), and Croce (2012)). I choose the persistence parameter ρ to be an annualized value of 0.978 to match the first-order autocorrelation of consumption growth. Remember that measured productivity is a composition of the exogenous and endogenous components in the model. Finally, I pick the volatility of the exogenous TFP process σ_ϵ to be 0.0175 to generate plausible macroeconomic volatilities.

Households

Households are characterized by recursive preferences, which are governed by three parameters γ , ψ and β . Substantial empirical work has been done on these parameters, see Bansal, Kiku and Yaron (2012), and this is followed here by setting $\gamma = 10$, $\psi = 0.9$, and $\beta = 0.9915$ to produce reasonable levels for price-dividend ratios. The elasticity of substitution parameter is usually assumed to be greater than one in much of the long-run risks literature (Croce (2012), Bansal and Yaron (2004)). The model, however, requires a value less than one to match the observed relationship of risk exposures (betas), as described in the next section.

1.4.2. Transition calibration

I match five transition paths: the IT-capital ratio, both sectors' price-dividend ratios, and the IT sector's average sales and net entry growth rates. The first path is the variable of interest. The latter four ensure that the model's asset pricing variables are consistent with financial market data. I discuss discount rates in the next section.

IT-capital ratio

Figure 2 plots the IT-capital ratio of the data versus that generated by the model. I calibrate the model to 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 the data on price-dividend ratios as well. The IT-capital ratio data are only available up until 2006. Appendix A.2 discusses in detail

Figure 2: Model calibration I: Input ratio

This figure plots the IT-capital ratio, the ratio of the IT sector's quantity of capital services to the industrial sector's. 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 this figure is in Appendix A.2. I fix an initial IT-capital ratio $\frac{N_0 X_0}{K_0} < \mathbb{E} \left[\frac{NX}{K} \right]$ and calibrate the model to match the length and curve of the data.

its construction. An important disciplining device is the model's other transitions.

Price-dividend ratios

The price-dividend ratio of the IT sector is defined in (1.18). The price-dividend ratio of the industrial sector is $q_t K_{t+1}/D_t$. Table 6 lists the values of the start- and end-points of the model that is consistent with the data available for the IT-capital ratio. In Figure 3 I plot the system's transitions. The model is able to match the industrial sector's price-dividend ratio to the time series data. Within its confidence 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”. What is important is that the data reverted back to the model’s implied value after the boom.

Table 6: Price-dividend ratios

This table reports price-dividend statistics generated by the model and compares them to the data. The model takes its calendar-time counterparts and estimates the average (across simulations) value of the price-dividend ratios of the IT and industrial sector. The data comes from an average of the previous sixteen quarters from data at year-end 1974 and 2006. The value ∞ represents the model’s steady-state value. I discuss the construction of price-dividend ratios in Appendix A.2.

Model	IT	IND	Ratio
1974	123	39	3.2
2006	44	28	1.6
∞	40	27	1.5
Data			
1974	121	33	3.7
2006	41	29	1.4

IT sector average sales growth rates

I plot the transition of the IT sector’s average sales growth rate in Figure 4. The model matches the fast, initial increase displayed by the data and then its drawn out path to convergence. This fast increase is consistent with a competition driving down the sales generated per firm. In the model, initially few firms dominate the marketplace. Over time, as more firms enter the marketplace, the industrial firm reduces the quantity demanded of each IT firm’s good. Consequently, sales and profit earned per firm falls.

IT sector net entry rates

I depict in Figure 5 the transition of the IT sector’s net entry rate. The model matches the sharp decline displayed initially by the data and then its drawn out path to convergence. The model is unable to get its mean to be negative to match the data after the dot-com boom. Indeed, this is impossible in the current setup because the entire economy’s growth

Figure 3: Model calibration II: Price-dividend ratios

This figure plots time series paths 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 errors. Standard errors are estimated from the standard deviation of point estimates across simulations. Ten-thousand simulations are run. The solid 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 A.2. Data are quarterly and are smoothed with a Hodrick-Prescott filter with a smoothing parameter equal to 1600.

Figure 4: Model calibration III: IT sector's average sales growth rate

This figure plots the average sales growth rate per firm of the IT sector. The dashed line is the model's average simulation path. The dotted lines are two times the model's standard errors. Standard errors are estimated from the standard deviation of point estimates across simulations. Ten-thousand simulations are run. The solid line is data. The data use Compustat data for the IT sector to calculate 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 smoothed with a Hodrick-Prescott filter with a smoothing parameter equal to 1600. In the model, the variable is $\log\left(\frac{\Pi_{t+1}}{\Pi_t}\right)$. IT firms are identified by NAICS codes in Appendix A.2.

Figure 5: Model calibration IV: IT sector’s firm-count growth rate

This figure plots the net entry rate of the IT sector. The dashed line is the model’s average simulation path. The dotted lines are two times the model’s standard errors. Standard errors are estimated from the standard deviation of point estimates across simulations. Ten-thousand simulations are run. The solid line is data. The data use the two-year compound annual growth rate of the growth rate of public IT firms. Data are quarterly and are smoothed with a Hodrick-Prescott filter with a smoothing parameter equal to 1600. IT firms are identified by NAICS codes described in Appendix A.2.

is driven by the expansion of the IT sector.²⁶ 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.

1.4.3. Quantitative analysis

Risk exposures

Until now, there has been little discussion of risk and discount rates. The stochastic discount factor specified in (1.13) implies two sources of risk: the first source relates to innovations

²⁶The model can be extended to incorporate another sector of growth in the economy. The current setup’s advantage is that it lucidly links the IT sector’s market valuation to future economy-wide growth.

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 premia and long-run risk, 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. 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.²⁷ The factors are standardized (to mean zero and variance one) before running the following regression for each sector i :

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

Figure 6 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 across the transition path. Why does it change? Consider a standard log-linearization of the return of the sector

$$\log(1 + r_{IT,t+1}) = \kappa_0 + \log\left(\frac{D_{t+1}^{IT}}{D_t^{IT}}\right) + \kappa_1 \log PD_{t+1}^{IT} - \log PD_t^{IT},$$

and compare this to (1.17), which I rewrite to highlight a distinction

$$PD_t^{IT} = \underbrace{\frac{\mathbb{E}_t \left[\sum_{s=1}^T M_{t|t+s} \{ \Pi_{t+s} N_{t+s} - S_{t+s} \} \right]}{N_t \Pi_t - S_t}}_{\text{Transition path}} + \underbrace{\frac{\mathbb{E}_t \left[\sum_{s=T+1}^{\infty} M_{t|t+s} \{ \Pi_{t+s} N_{t+s} - S_{t+s} \} \right]}{N_t \Pi_t - S_t}}_{\text{Steady state}}.$$

We can see that as calendar time t approaches the convergence time T , the information

²⁷The return on wealth is defined recursively by the equation $W_t = C_t + \mathbb{E}_t[M_{t+1}W_{t+1}]$ or equivalently by $r_{C,t+1} \equiv \frac{W_{t+1}}{W_t - C_t}$. Data estimates of the long-run risk factor are described in Appendix A.2.

Figure 6: Model: Rolling risk exposures II

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. Ten-thousand simulations are run. Estimation details are provided in Section 1.4.

contained in the price-dividend ratio changes, which directly influences the return of the IT sector as shown above. Indeed, the IT sector becomes more exposed to short-run risk as the transition elapses. Because the initial size of IT sector is small, it absorbs a small amount of the risk to consumption. As the transition proceeds, IT 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,

it becomes a hedge. An adverse shock to the expected growth rate benefits the IT 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 long-run growth. Information technology, in the steady state, acts as a hedge against shocks to long-run growth.

Moments

I report in Table 7 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. It is important to get these 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.

Table 7: Model and data macroeconomic statistics

This table reports macroeconomic growth and volatility statistics generated by the model and compares them to the data. The model is simulated on a quarterly basis and then time-aggregated to an annual one. Data statistics are calculated from Bureau of Economic Analysis's data over the annual period 1974–2012. In the model, 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 A.2.

Statistic	Data	Model
<i>Consumption growth</i>		
mean(Δc)	2.00	1.95
std(Δc)	2.27	2.52
AC1(Δc)	0.39	0.43
<i>Business cycle</i>		
$\sigma_{\Delta c}/\sigma_{\Delta y}$	0.61	0.96
$\sigma_{\Delta INV}/\sigma_{\Delta c}$	4.38	3.14

The model is able to generate a risk-free rate that has both a low volatility and a high

persistence (first-order autocorrelation is 0.99 at a quarterly rate) as shown in Table 8. It is higher than commonly estimated because it is calibrated to match the levels of the sectors' price-dividend ratios.

Table 8: 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 $\mathbb{E}[r_i^{\text{LEV}} - r_f] \equiv \mathbb{E}[r_i - r_f](1 + \overline{D/E})$, 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 the model's estimates after including stochastic volatility. Panel C reports my estimates of the data's 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: Benchmark model			
	Mean	Stdev	Sharpe ratio
$\mathbb{E}[r_f]$	3.86	0.8	0
$\mathbb{E}[r_{IND}]$	8.25	5.75	0.80
$\mathbb{E}[r_{IT}]$	4.39	3.46	0.21
$\mathbb{E}[r_{MKT}]$	6.46	5.40	0.51

Panel B: Model with stochastic volatility

	Mean	Stdev	Sharpe ratio
$\mathbb{E}[r_f]$	2.08	3.36	0
$\mathbb{E}[r_{IND}]$	9.82	15.2	0.50
$\mathbb{E}[r_{IT}]$	4.24	11.0	0.20
$\mathbb{E}[r_{MKT}]$	7.61	15.6	0.35

Panel C: 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

In the same table, the equity returns 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, moreover, does not specify volatility as being stochastic as in Bansal and Yaron (2004), which could aid in generating risk premia and much greater equity volatility. Calibrating an additional equation for stochastic volatility (with parameters close to Bansal et al.'s (2012) choices, but with a higher shock volatility) such as the following

$$\begin{aligned}\log(A_{t+1}) &= \rho \log(A_t) + \sigma_t \epsilon_{t+1} \\ \sigma_{t+1} &= \bar{\sigma} + \rho_\sigma(\sigma_t - \bar{\sigma}) + \eta_{t+1},\end{aligned}$$

where $\epsilon \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, $\bar{\sigma} = 0.0175$, $\rho_\sigma = 0.985$, and $\eta \stackrel{iid}{\sim} \mathcal{N}(0, 0.0135^2)$, aids in generating more realistic return processes.

The model also generates a value premium if the industrial sector is taken to proxy for value stocks and the IT sector proxies for growth stocks. This is consistent with Zhang's (2005) work, who generates the value premium by appealing to an asymmetric adjustment cost. Taken together, these estimates of risk exposures and these asset pricing moments corroborate the model's ability to match the relevant financial market discount rates.

Full dynamics

I provide intuition on four variables's average transition dynamics that are displayed in Figure 7. 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 how I specify the research externality. Initially, the rate of expansion of the IT sector is slow, because the capital reallocation friction, whose strength is indexed by the parameter η_k , drags heavily on research productivity. As time passes, the strength of this friction wanes; consequently, innovation begins feeding on itself, as the IT 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 reducing the rate of innovating and the expansion of IT. Because of the restriction on balanced growth, eventually IT, capital, and the rest of the economy grow, on average, at the same rate.

Figure 7: Model: Full transition paths

The top-left panel plots the IT-capital ratio. The top-right panel plots the quarterly 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. Ten-thousand simulations are run and the figures below show the average across simulations.

The bottom-left panel displays the price-dividend ratios of the two sectors. Because the value of an IT good is initially very high, the price-dividend ratio of the IT sector is also high. The final-good firm's production function specifies positive cross-partial derivatives for capital and IT, so its value, as well, is 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.

The top-right panel plots quarterly productivity growth. 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.

The bottom-right panel features a marked drop in per firm profitability resulting from a decline in the quantity demanded for each IT good X_t that takes place as the IT sector grows. This is consistent with competition increasing for every firm as the transition runs its course. 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 8 plots the density of convergence times. The distribution is skewed right. This is 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 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.

Distributions of productivity

Finally, Figure 9 plots the distributions of historical TFP growth and model-implied future TFP growth. There are two things to notice. The first is that the historical distribution is

Figure 8: Model: Density of convergence times

This figure plots the density of convergence times, as described in Section 1.4. The convergence time is defined as $T \equiv \inf \left\{ t : \frac{N_t X_t}{K_t} \geq \mathbb{E} \left[\frac{N_t X_t}{K_t} \right] \right\}$. Ten-thousand simulations are run.

much more symmetric than the future distribution. This is because both distributions are normalized per year, and 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. 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 continued to fall, lowering the incentive to conduct research and to produce new IT goods. As a result, future TFP growth of the economy is expected to be lower than before.²⁸

²⁸A reservation with this statement is that long-run TFP growth is solely generated by the IT sector and, from the balanced growth condition, is set to two percent per year in the long-run, which is higher than the historical rate of 0.9 percent. To adjust the current model’s implication, I take the ratio of means to adjust the forecasted TFP growth rate per year. Therefore, expected TFP growth per year is $0.87 \times 2.7/4.5 = 0.52$, a reduction of 35 basis points. Note this is similar to Robert J. Gordon’s “educational plateau adjustment” of 27 basis points he calculated in *The Economist’s* online debate.

Figure 9: Model: Densities of TFP growth per year

This figure plots the density of productivity growth per year. The top figure plots the historical amount. The actual amount observed in the data was 0.87 percent TFP growth per year. The bottom figure generates the future TFP growth from 2012 until the stopping time T is reached. Both cumulative growth rates are divided by the number of years. Ten-thousand simulations are run. The figures plot the entire density of all the simulations. Further details are provided in Section 1.4.

1.5. Conclusion

In this paper, I build an asset pricing model that endogenously links economy-wide growth to innovation in the IT sector, whose 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. I then calibrate the model's transition paths to match historical data on factor shares, price-dividend ratios, growth rates, and discount rates.

This new method I develop puts structure on financial market data to forecast future economic outcomes. I apply this methodology to study the IT sector's transition towards its long-run share in the US economy, along with its implications for future growth. Future work could apply this methodology to revolutions of the past, such as the electricity revolution, to assess its predictive ability. It could also be applied in principle to study other phenomena, such as the discovery of a large, exhaustible energy resource.

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 indicates the IT sector's transition ends in 2033, six decades after its 1974 inception. My analysis also suggests that the sector's transition is likely to end within the decade before 2033 than within the decade after.

APPENDIX

A.1. Proofs and derivations

Constant markups

The price, $P_t(i)$, of IT good i is chosen to maximize the IT firm's profit. Take the quantity demanded, $X_i(P_t(i); P_t(i') \neq i), K_t, L_t, A_t, G_t) \equiv \arg \max_{X_t(i)} \Pi_t(i)$, by the representative final-good firm for this particular good as given. For brevity, write $X_t(i)$ instead of $X_i(P_t(i); P_t(i') \neq i), K_t, L_t, A_t, G_t)$. Each monopolist solves the following static profit maximization problem each period:

$$\max_{P_t(i)} \Pi_t(i) \equiv P_t(i)X_t(i) - X_t(i)$$

Differentiating with respect to $P_t(i)$ and plugging in (A.4) gives

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

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

First-order conditions

$$I_t : \quad q_t = \frac{1}{\Lambda' \left(\frac{I_t}{K_t} \right)} \quad (\text{A.1})$$

$$K_{t+1} : \quad q_t = \mathbb{E}_t \left[M_{t+1} \left\{ \alpha(1-m) \frac{Y_{t+1}}{K_{t+1}} + q_{t+1} \left((1-\delta) - \Lambda' \left(\frac{I_{t+1}}{K_{t+1}} \right) \cdot \left(\frac{I_{t+1}}{K_{t+1}} \right) + \Lambda \left(\frac{I_{t+1}}{K_{t+1}} \right) \right) \right\} \right]. \quad (\text{A.2})$$

$$L_t : \quad W_t = (1-\alpha)(1-m) \frac{Y_t}{L_t} \quad (\text{A.3})$$

$$X_t(i) : \quad P_t(i) = (K_t^\alpha (A_t L_t)^{1-\alpha}) \mu m \left[\int_0^{N_t} X_t(i)^{\frac{1}{\mu}} di \right]^{\mu m - 1} \frac{1}{\mu} X_t(i)^{\frac{1}{\mu} - 1}. \quad (\text{A.4})$$

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 $X_t(i)$ and can be simplified, leading to (1.9). 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 final-good firm, consistent with the literature's usual treatment.

Market values (Proposition 1)

Proof. First, multiply (A.2) by K_{t+1} to give

$$\begin{aligned}
q_t K_{t+1} &= \mathbb{E}_t \left[M_{t+1} \left\{ \alpha(1-m)Y_{t+1} + q_{t+1} \left((1-\delta)K_{t+1} - \Lambda' \left(\frac{I_{t+1}}{K_{t+1}} \right) \cdot I_{t+1} + \Lambda \left(\frac{I_{t+1}}{K_{t+1}} \right) K_{t+1} \right) \right\} \right] \\
&= \mathbb{E}_t \left[M_{t+1} \left\{ \alpha(1-m)Y_{t+1} + q_{t+1}K_{t+2} - q_{t+1}\Lambda' \left(\frac{I_{t+1}}{K_{t+1}} \right) \cdot I_{t+1} \right\} \right], \text{ by (1.8)} \\
&= \mathbb{E}_t [M_{t+1} \{ \alpha(1-m)Y_{t+1} - I_{t+1} + q_{t+1}K_{t+2} \}], \text{ by (A.1).}
\end{aligned}$$

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 a transversality condition gives the result. The simpler proof that does not involve adjustment costs can be attained by specifying $q_t = 1$ in the last equation for all t .

Second, what needs to be shown is that

$$N_t(V_t - \Pi_t) + O_t = \mathbb{E}_t \left[\sum_{s=1}^{\infty} M_{t|t+s} \{ \Pi_{t+s} N_{t+s} - S_{t+s} \} \right].$$

A similar claim was made by Iraola and Santos (2009) and Comin, Gertler and Santacreu (2009). By specifying $\phi = 1$ and a constant discount factor $M_{t|t+s} = \beta^s$, one attains the simple version.

Begin with the observation that

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

Iterating this equation forward and assuming the absence of bubbles in equilibrium (see Santos and Woodford (1997)) proves the claim. $\square$

Deterministic transition dynamics (Proposition 2)

I present the result of the proposition in two parts. The first part shows that the system, when starting from any $k_0 > 0$, k_t monotonically converges to k^* . The second part introduces the time $T = \inf\{t : |k_t - k^*| > \varepsilon\}$, and analyzes the effects of its introduction.

For convenience, define here the following system objects:

- $\pi \equiv (\mu - 1) \left(\frac{m}{\mu}\right)^{\frac{\eta_s}{1-\eta_s}}$, so $\Pi_t = \pi k_t^\alpha$
- $k_{t+1}(1 + g_{N,t+1}) = (1 + g_N^*)k_t$
- $k^*(1 + \bar{g}_{N,t+1}) = (1 + g_N^*)k_t$, where $\{\bar{g}_{N,t+1}\}_{t=0}^\infty$ is a sequential bound on growth
- $V_t = \pi k_t^\alpha + \beta V_{t+1} = \pi \sum_{s=0}^\infty \beta^s k_{t+s}^\alpha$
- $1 + g_{N,t+1} = \phi + \chi^{\frac{1}{1-\eta_s}} (\beta V_{t+1})^{\frac{\eta_s}{1-\eta_s}}$

Proof. To show the first part, it's sufficient to show the following two-implication relation:

$$k_0 > k^* \text{ and } \exists \{\bar{g}_{N,t+1}\}_{t=0}^\infty \text{ s.t. } g_{N,t+1} < \bar{g}_{N,t+1}, \forall t \Rightarrow k_t > k^*, \forall t \Rightarrow k_{t+1} < k_t, \forall t.$$

First take the right-most implication and suppose that, for a contradiction, $\exists \tilde{t} : k_{\tilde{t}+1} \geq k_{\tilde{t}}$ and $k_t > k^*, \forall t$. If there is one $\tilde{t}$ such that $k_{\tilde{t}+1} \geq k_{\tilde{t}}$, then $g_{N,\tilde{t}+1} \leq g^*$. Because V_{t+1} is a bijective function of $g_{N,t+1}$, then $V_{\tilde{t}+1} \leq V^*$. Using the definition of V yields

$$\pi \sum_{s=0}^\infty \beta^s k_{\tilde{t}+s+1}^\alpha \leq \frac{\pi (k^*)^\alpha}{1 - \beta},$$

which implies there is at least one k_{t+s} that is less than k^* , a contradiction.

The left-most implication can be proved by induction: I'll show the base case, when $k_0 > k^*$ implies $k_1 > k^*$; and then show the induction step, where $k_t > k^*$ implies $k_{t+1} > k^*$. For the base case, suppose that, for a contradiction, $k_1 \leq k^* < k_0$. There are two cases: (i) $g_{N,1} > g_N^*$, or (ii) $g_{N,1} \leq g_N^*$. If (i) holds, then we can use the bound on growth to get

$k^* = \left(\frac{1+g_N^*}{1+\bar{g}_{N,1}} \right) k_0 < \left(\frac{1+g_N^*}{1+g_{N,1}} \right) k_0 = k_1$, a contradiction. If (ii) holds, then $k_1 = \left(\frac{1+g_N^*}{1+g_{N,1}} \right) k_0 > \left(\frac{1+g_N^*}{1+\bar{g}_{N,1}} \right) k^* \geq \left(\frac{1+g_N^*}{1+g_{N,1}} \right) k_1$, again a contradiction. To show the induction step, the same contradictions to prove the base case can be used. Because $\{k_t\}_{t=0}^\infty$ is a monotone sequence and is bounded below by k^* , then $k_t \rightarrow k^*$.

Now I can show that for all t we have $V_t > V_{t+1}$, $g_{N,t+1} > g_{N,t}$, and $\bar{g}_{N,t+1} > \bar{g}_{N,t}$. For any t , we have

$$\begin{aligned} V_{t+1} &= \pi k_{t+1}^\alpha + \pi \beta k_{t+2}^\alpha + \dots \\ &< \pi k_t^\alpha (1 + \beta + \beta^2 + \dots), \text{ by the two-implication relation} \\ &= \pi \frac{k_t^\alpha}{1 - \beta}, \end{aligned}$$

which can be rearranged to give $V_t = \pi k_t^\alpha + \beta V_{t+1} > V_{t+1}$, which holds for all t . Because $g_{N,t+1}$ is a bijection of V_{t+1} , then $g_{N,t+1} < g_{N,t}$, for every t . Finally, by definition of $\bar{g}_{N,t+1}$:

$$k^* = \left(\frac{1+g_N^*}{1+\bar{g}_{N,t+1}} \right) k_t = \left(\frac{1+g_N^*}{1+\bar{g}_{N,t}} \right) k_{t-1},$$

so by the two-implication relation, $\bar{g}_{N,t+1} > \bar{g}_{N,t}$, for all t . In sum, $\mathbf{x}_t = \{k_t, g_{N,t+1}, \bar{g}_{N,t+1}, V_t\}$ converges monotonically to $\mathbf{x}^*$. Because the IT-capital ratio is an inverse mapping of k_t , it similarly converges. The argument for $0 < k_0 < k^*$ is similar and analogously requires the sequence $\{\bar{g}_{N,t+1}\}_{t=0}^\infty$ s.t. $g_{N,t+1} < \bar{g}_{N,t+1}$, $\forall t$.

We're interested in the price-dividend ratio's dynamics, so let's now look at that. The price-dividend ratio in (1.17) can have the numerator and the denominator looked at separately. The numerator is the aggregate ex-dividend market value of the IT sector and is equal to

$$V_t^{IT} = \sum_{s=1}^{\infty} \beta^s N_{t+s} \Pi_{t+s} \Rightarrow \frac{V_t^{IT}}{N_t} = \sum_{s=1}^{\infty} \beta^s \frac{N_{t+s}}{N_t} \Pi_{t+s}.$$

Because k_t and $g_{N,t+1}$ follow decreasing sequences, so does the aggregate ex-dividend market value of the IT sector per firm.

The denominator comprises the aggregate IT sector dividend $N_t \Pi_t$ and the IT sector research expenditure S_t , which per firm is $\Pi_t - \frac{S_t}{N_t}$. While the difference is always positive, its relationship as a function of time is ambiguous. On the one hand, dividends per firm will decrease overtime, but so does the amount of research expenditure per firm. The determination of which effect will dominate depends on the parameters and steady state of the model. When uncertainty is introduced in Section 1.4, the analysis of the denominator is further confounded and becomes a quantitative question. Note that if research expenditure per firm falls faster than dividends per firm, then the price-dividend ratio of the IT sector will decline as a function of time.

For the second part, we introduce the time T . Write $V_t(T)$ to acknowledge V 's dependence on T . Start the economy at $k_0 > k^T = k^* + \varepsilon$, and fix ε and consequently T . For the last bullet, the value of an IT good, $V_t(T)$, can simply be differentiated with respect to k_t to get

$$\frac{\partial V_t(T)}{\partial k_t} = \pi \alpha k_t^{\alpha-1} > 0.$$

For the second part of the fourth bullet, consider decreasing ε such that the convergence time T increases by one period: $\tilde{T} = T + 1$. Then define

$$V_t(\tilde{T}) = \pi \left(k_t^\alpha + \beta k_{t+1}^\alpha + \cdots + \beta^{T-t} k_T^\alpha + \beta^{\tilde{T}-t} k_{\tilde{T}}^\alpha + \beta^{\tilde{T}-t+1} \frac{(k^*)^\alpha}{1-\beta} \right) \text{ for } 0 \leq t \leq \tilde{T}.$$

Simply subtracting $V_t(T)$ from $V_t(\tilde{T})$ shows $V_t(\tilde{T}) - V_t(T) = \pi \beta^{\tilde{T}-t} k_{\tilde{T}}^\alpha > 0$. □

A.2. Data construction

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.

Construction of IT-capital ratio

Consistent with Jorgenson et al.’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 capital services, which has a close analog to the quantities $N_t X_t$ 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.

Price-dividend ratios

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

$$R_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t}, \quad RX_{t+1} \equiv \frac{P_{t+1}}{P_t}.$$

Price-dividend ratios are then constructed as the inverse of

$$\frac{D_{t+1}}{P_{t+1}} = \frac{R_{t+1}}{RX_{t+1}} - 1. \quad (\text{A.5})$$

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). Data are restricted to stocks trading on the NYSE, AMEX, and NASDAQ exchanges, having share codes 10 and 11, and being US-headquartered firms.

Because the incidence of firms which repurchase shares has increased, an alternative measure of payouts to equity shareholders is used.¹ 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 et al.'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 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

¹Fama 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.

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. Moreover, 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.²

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 similar 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 be classified by management as either an expenditure before taxes and earnings or as investment in a capital asset. Thus, this discretion can be used to manipulate earnings.

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.

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

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.

Estimation of default

Following Campbell, Hilscher and Szilagyi (2008) and Boualam, Gomes and Ward (2012), I proxy default with performance-related delisting events on CRSP.³ 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:

1. Kalman filter

³Delisting codes used are 500, 550, 552, 560, 561, 574, 580, and 584. They are defined at <http://www.crsp.com/products/documentation/delisting-codes>.

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 9. The Kalman filter method follows Croce (2012). 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 : $r_{C,t} = \text{proj}[\Delta c_{t+1}|X_t] = \hat{\beta} X_t$.

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 9: 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 A.2.

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

Figure A plots rolling 50-quarter regressions of beta estimates 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 because it specifies the intertemporal elasticity of substitution to be less than one (the calibration uses 0.9).

Figure 10: Model calibration V: Rolling risk exposures I

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 errors. Standard errors are estimated from the standard deviation of point estimates across simulations. Ten-thousand simulations are run. 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. Estimation details are provided in Section A.2.

BIBLIOGRAPHY

- Abel, Andrew B. and Janice C. Eberly**, “Investment, Valuation, and Growth Options,” *The Quarterly Journal of Finance*, March 2012, 2 (1).
- Acemoglu, Daron**, *Introduction to Modern Economic Growth*, Princeton University Press, 2009.
- Aghion, Phillipe and Peter Howitt**, “A Model of Growth Through Creative Destruction,” *Econometrica*, March 1992, 60 (2), 323–351.
- Atkeson, Andrew and Patrick J. Kehoe**, “Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions,” *The American Economic Review*, March 2007, 97 (1), 64–88.
- Bakos, Yannis and Erik Brynjolfsson**, “Bundling Information Goods: Pricing, Profits, and Efficiency,” *Management Science*, December 1999, 45 (12), 1613–1630.
- Bansal, Ravi and Amir Yaron**, “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *The Journal of Finance*, August 2004, 59 (4), 1481–1509.
- , **Dana Kiku, and Amir Yaron**, “An Empirical Evaluation of the Long Run Risks Model for Asset Prices,” *Critical Finance Review*, 2012, 1, 183–221.
- , **Robert F. Dittmar, and Christian T. Lundblad**, “Consumption, Dividends, and the Cross Section of Equity Returns,” *The Journal of Finance*, August 2005, 60 (4), 1639–1672.
- Bartel, Ann, Casey Ichnowski, and Kathryn Shaw**, “How Does Information Technology Affect Productivity? Plant-level Comparisons of Product Innovation, Process Improvement, and Worker Skills,” *Quarterly Journal of Economics*, November 2007, 122 (4), 1721–1758.
- Beaudry, Paul and Franck Portier**, “An Exploration Into Pigou’s Theory of Cycles,” *Journal of Monetary Economics*, September 2004, 51 (6), 1183–1216.
- and —, “Stock Prices, News, and Economic Fluctuations,” *The American Economic Review*, September 2006, 96 (4), 1293–1307.
- Belo, Frederico**, “Production-Based Measures of Risk for Asset Pricing,” *Journal of Monetary Economics*, March 2010, 57 (2), 146–163.
- Berk, Jonathan B., Richard C. Green, and Vasant Naik**, “Optimal Investment, Growth Options, and Security Returns,” *The Journal of Finance*, October 1999, 54 (5), 1553–1607.

- Black, Sandra E. and Lisa M. Lynch**, “How to Compete: The Impact of Workplace Practices and Information Technology on Productivity,” *The Review of Economics and Statistics*, August 2001, 83 (3), 434–445.
- Boualam, Yasser, Joao F. Gomes, and Colin Ward**, “Understanding the Behavior of Distressed Stocks,” December 2012. Working paper.
- Brynjolfsson, Erik and Andrew McAfee**, *Race Against The Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, Digital Frontier Press, 2011.
- Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel**, “Is the Information Technology Revolution Over?,” Technical Report 25, Centre for the Study of Living Standards Spring 2013. International Productivity Monitor.
- Campbell, John Y. and Robert J. Shiller**, “Stock Prices, Earnings, and Expected Dividends,” *The Journal of Finance*, July 1988, 43 (3), 661–676.
- , **Jens Hilscher, and Jan Szilagyi**, “In Search of Distress Risk,” *The Journal of Finance*, December 2008, 63 (6), 2899–2939.
- Carlson, Murray, Adlai Fisher, and Ron Giammarino**, “Corporate Investment and Asset Price Dynamics: Implications for the Cross-section of Returns,” *Journal of Finance*, December 2004, 59 (6), 2577–2603.
- Cochrane, John H.**, “A Cross-Sectional Test of an Investment-Based Asset Pricing Model,” *The Journal of Political Economy*, June 1996, 104 (3), 572–621.
- , “Stocks as Money: Convenience Yield and the Tech-Stock Bubble,” in William C. Hunter, George G. Kaufman, and Michael Pomerleano, eds., *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies*, Cambridge and London: MIT Press, 2003.
- , “Presidential Address: Discount Rates,” *The Journal of Finance*, August 2011, 66 (4), 1047–1108.
- Colacito, Riccardo and Mariano M. Croce**, “Risks for the Long Run and the Real Exchange Rate,” *Journal of Political Economy*, February 2011, 119 (1), 153–181.
- Comin, Diego and Mark Gertler**, “Medium-Term Business Cycles,” *The American Economic Review*, June 2006, 96 (3), 523–551.
- , —, and **Ana Maria Santacreu**, “Technology Innovation and Diffusion as Sources of Output and Asset Price Fluctuations,” June 2009. Harvard Business School Working Paper 09-134.
- Cooper, Ilan**, “Asset Pricing Implications of Nonconvex Adjustment Costs and Irreversibility of Investment,” *The Journal of Finance*, February 2006, 61 (1), 139–170.

- Cowen, Tyler**, *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better Again*, Dutton, 2011.
- Croce, Mariano Massimiliano**, “Long-run Productivity Risk: A New Hope for Production-Based Asset Pricing?,” May 2012. SSRN working paper 1106361.
- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *The Journal of Finance*, July 1989, 57 (4), 937–969.
- Fama, Eugene F. and Kenneth R. French**, “Disappearing Dividends: Changing Firm Characteristics or Lower Propensity to Pay?,” *Journal of Financial Economics*, 2001, 60, 3–43.
- Gala, Vito**, “Irreversible Investment and the Cross-Section of Stock Returns in General Equilibrium,” 2010. Working paper.
- Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas**, “Displacement Risk and Asset Returns,” *Journal of Financial Economics*, September 2012, 105 (3), 491–510.
- , **Stavros Panageas, and Jianfeng Yu**, “Technological Growth and Asset Pricing,” *The Journal of Finance*, August 2012, 67 (4), 1265–1292.
- Goeree, Michelle Sovinsky**, “Limited Information and Advertising in the US Personal Computer Industry,” *Econometrica*, September 2008, 76 (5), 1017–1074.
- Gomes, João F., Leonid Kogan, and Lu Zhang**, “Equilibrium Cross-Section of Returns,” *Journal of Political Economy*, August 2003, 111 (4), 693–732.
- , —, and **Motohiro Yogo**, “Durability of Output and Expected Stock Returns,” *Journal of Political Economy*, October 2009, 117 (5), 941–986.
- Goolsbee, Austan and Peter J. Klenow**, “Evidence on Learning and Network Externalities in the Diffusion of Home Computers,” *Journal of Law and Economics*, October 2002, 45, 317–343.
- Gordon, Robert J.**, “Does the New Economy Measure up to the Great Inventions of the Past?,” *Journal of Economic Perspectives*, Fall 2000, 14 (4), 49–74.
- , “Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds,” August 2012. NBER Working paper No. 18315.
- , “Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds,” Technical Report 25, Centre for the Study of Living Standards Spring 2013. International Productivity Monitor.
- Gourio, François**, “Disaster Risk and Business Cycles,” *American Economic Review*, October 2012, 102 (6), 2734–66.

- Goyal, Amit and Ivo Welch**, “A Comprehensive Look at The Empirical Performance of Equity Premium Prediction,” *The Review of Financial Studies*, 2008, 21 (4), 1455–1508.
- Greenwood, Jeremy and Boyan Jovanovic**, “The Information-Technology Revolution and the Stock Market,” *The American Economic Review*, May 1999, 89 (2), 116–122.
- and **Mehmet Yorukoglu**, “1974,” *Carnegie-Rochester Conference Series on Public Policy*, June 1997, 46, 49–95.
- , **Zvi Hercowitz**, and **Per Krusell**, “Long-Run Implications of Investment-Specific Technological Change,” *The American Economic Review*, June 1997, 87 (3), 342–362.
- Griliches, Zvi**, “Patent Statistics as Economic Indicators: A Survey,” *Journal of Economic Literature*, December 1990, 28 (4), 1661–1707.
- Grossman, Gene M. and Elhanan Helpman**, *Innovation and Growth in the Global Economy*, MIT Press, 1991.
- Grullon, Gustavo and Roni Michaely**, “Dividends, Share Repurchases, and the Substitution Hypothesis,” *The Journal of Finance*, August 2002, 57 (4), 1649–1684.
- Hobijn, Bart and Boyan Jovanovic**, “The Information-Technology Revolution and the Stock Market: Evidence,” *The American Economic Review*, December 2001, 91 (5), 1203–1220.
- Hodrick, Robert J.**, “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, 1992, 5 (3), 357–386.
- Iraola, Miguel A. and Manuel S. Santos**, “Long-term Asset Price Volatility and Macroeconomics Fluctuations,” July 2009. Working paper.
- Jaruzelski, Barry, John Loehr, and Richard Holman**, “The 2012 Global Innovation 1000 Study: Making Ideas Work,” Technical Report, Booz and Company 2012.
- Jermann, Urban J.**, “Asset Pricing in Production Economies,” *Journal of Monetary Economics*, February 1998, 41 (2), 257–275.
- Jorgenson, Dale W. and Kevin J. Stiroh**, “Raising the Speed Limit: US Economic Growth in the Information Age,” *Brookings Papers on Economic Activity*, February 2000, 31 (1), 126–236.
- , **Mun S. Ho**, and **Jon D. Samuels**, “Information Technology and US Productivity Growth: Evidence from a Prototype Industry Production Account,” *Journal of Productivity Analysis*, 2011, 36, 159–175.

- Jovanovic, Boyan and Glenn M. MacDonald**, “The Life Cycle of a Competitive Industry,” *Journal of Political Economy*, April 1994, 102 (2), 322–347.
- and **Peter L. Rousseau**, “Moore’s Law and Learning by Doing,” *Review of Economic Dynamics*, 2002, 5, 346–375.
- and —, “General Purpose Technologies,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 1, Elsevier B.V., 2005, chapter 18, pp. 1182–1221.
- Kaltenbrunner, Georg and Lars A. Lochstoer**, “Long-Run Risk through Consumption Smoothing,” *The Review of Financial Studies*, 2010, 23 (8), 3190–3224.
- Kilian, Lutz**, “Exchange Rates and Monetary Fundamentals: What Do We Learn from Long-Horizon Regressions?,” *Journal of Applied Econometrics*, September/October 1999, 14 (5), 491–510.
- Kogan, Leonid**, “Asset Prices and Real Investment,” *Journal of Financial Economics*, September 2004, 73 (4), 411–431.
- and **Dimitris Papanikolaou**, “Firm Characteristics and Stock Returns: The Role of Investment-Specific Shocks,” *Review of Financial Studies*, 2013. Forthcoming.
- and —, “Growth Opportunities, Technology Shocks and Asset Prices,” *The Journal of Finance*, 2013. Forthcoming.
- , —, **Amit Seru**, and **Noah Stoffman**, “Technological Innovation, Resource Allocation, and Growth,” 2012. Working paper.
- , —, and **Noah Stoffman**, “Technological Innovation: Winners and Losers,” 2013. Working paper.
- Kung, Howard and Lukas Schmid**, “Innovation, Growth, and Asset Prices,” April 2012. Working paper.
- Li, Wendy C. Y.**, “Depreciation of Business R&D Capital,” Technical Report, Bureau of Economic Analysis October 2012. R&D Satellite Account Paper.
- Loualiche, Erik**, “Asset Pricing with Entry and Imperfect Competition,” 2013. Working Paper.
- Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan**, “The Wealth-Consumption Ratio,” *The Review of Asset Pricing Studies*, 2013, 3 (1), 38–94.
- Metrick, Andrew and Ayako Yasuda**, *Venture Capital and the Finance of Innovation*, John Wiley and Sons, Inc., 2011.

- Moore, Gordon E.**, “No Exponential is Forever: But Forever can be Delayed!,” in “Solid-State Circuits Conference,” Vol. 1 of *Digest of Technical Papers* 2003, pp. 20–23.
- Newey, Whitney K. and Kenneth D. West**, “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, May 1987, *55* (3), 703–708.
- Novy-Marx, Robert**, “An Equilibrium Model of Investment Under Uncertainty,” *The Review of Financial Studies*, 2007, *20* (5), 1461–1502.
- Oberfield, Ezra**, “Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture,” January 2013. Working paper.
- Ofek, Eli and Matthew Richardson**, “DotCom Mania: The Rise and Fall of Internet Stock Prices,” *The Journal of Finance*, June 2003, *58* (3), 1113–1138.
- Opp, Christian C.**, “Cycles of Innovation and Financial Propagation,” January 2010. Working paper.
- Papanikolaou, Dimitris**, “Investment Shocks and Asset Prices,” *Journal of Political Economy*, October 2011, *119* (4), 639–685.
- Pástor, Ľuboš and Pietro Veronesi**, “Was There a NASDAQ Bubble in the Late 1990s?,” *Journal of Financial Economics*, July 2006, *81* (1), 61–100.
- and —, “Technological Revolutions and Stock Prices,” *The American Economic Review*, September 2009, *99* (4), 1451–1483.
- Rauh, Joshua D. and Amir Sufi**, “Explaining Corporate Capital Structure: Product Markets, Leases, and Asset Similarity,” *Review of Finance*, 2012, *16* (1), 115–155.
- Rebelo, Sergio**, “Long-Run Policy Analysis and Long-Run Growth,” *The Journal of Political Economy*, 1991, *99* (3), 500–521.
- Romer, Paul M.**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, *98* (5), S71–S102.
- Santos, Manuel S. and Michael Woodford**, “Rational Asset Pricing Bubbles,” *Econometrica*, January 1997, *65* (1), 19–57.
- Schmitt-Grohe, S. and M. Uribe**, “Solving Dynamic General Equilibrium Models Using a Second-Order Approximation to the Policy Function,” *Journal of Economic Dynamics and Control*, 2004, *28*, 755–775.
- Shapiro, Carl and Hal R. Varian**, *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, 1999.

Stambaugh, Robert F., “Predictive Regressions,” *Journal of Financial Economics*, 1999, *54* (1), 375–421.

Weil, Phillipe, “The Equity Premium Puzzle and the Risk-free Rate Puzzle,” *Journal of Monetary Economics*, November 1989, *24* (3), 401–421.

Zhang, Lu, “The Value Premium,” *The Journal of Finance*, February 2005, *60* (1), 67–103.