

NBER WORKING PAPER SERIES

TECHNOLOGY INNOVATION AND DIFFUSION AS SOURCES OF OUTPUT AND
ASSET PRICE FLUCTUATIONS

Diego A. Comin
Mark Gertler
Ana Maria Santacreu

Working Paper 15029
<http://www.nber.org/papers/w15029>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2009

We appreciate the helpful comments of Marianne Baxter, Paul Beaudry, John Campbell, Larry Christiano, Jordi Gali, Bob King, John Leahy, Martin Lettau, Monika Piazzesi, Martin Schneider, and seminar participants at the Boston Fed, Brown, Boston University, NBER Summer Institute, University of Valencia, University of Houston, Harvard Business School, Swiss National Bank, CEPR-CREI and NBER-EFG meeting. We appreciate the excellent research assistance of Albert Queralto. Financial assistance from the C.V. Starr Center and the NSF is greatly appreciated. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Diego A. Comin, Mark Gertler, and Ana Maria Santacreu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Technology Innovation and Diffusion as Sources of Output and Asset Price Fluctuations
Diego A. Comin, Mark Gertler, and Ana Maria Santacreu
NBER Working Paper No. 15029
June 2009
JEL No. E2,E3

ABSTRACT

We develop a model in which innovations in an economy's growth potential are an important driving force of the business cycle. The framework shares the emphasis of the recent "new shock" literature on revisions of beliefs about the future as a source of fluctuations, but differs by tying these beliefs to fundamentals of the evolution of the technology frontier. An important feature of the model is that the process of moving to the frontier involves costly technology adoption. In this way, news of improved growth potential has a positive effect on current hours. As we show, the model also has reasonable implications for stock prices. We estimate our model for data post-1984 and show that the innovations shock accounts for nearly a third of the variation in output at business cycle frequencies. The estimated model also accounts reasonably well for the large gyrations in stock prices over this period. Finally, the endogenous adoption mechanism plays a significant role in amplifying other shocks.

Diego A. Comin
Harvard Business School
Soldiers Field
Boston, MA 02163
and NBER
dcomin@hbs.edu

Ana Maria Santacreu
Department of Economics
New York University
New York, NY 10012
ams458@nyu.edu

Mark Gertler
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10003
and NBER
mark.gertler@nyu.edu

1 Motivation

A central challenge to modern business cycle analysis is that there are few if any significant primitive driving forces that are readily observable. Oil shocks are perhaps the main example. But even here there is controversy. Not all recessions are preceded by major oil price spikes and there is certainly little evidence that major expansions are fueled by oil price declines. Further, given its low cost share of production, there is debate over whether in fact oil shocks alone could be a source of major output swings. Credit conditions have been a key factor in some of the postwar recessions, including the current one, but not in all.

Motivated by the absence of significant observable shocks, an important paper by Beaudry and Portier (2004) proposes that news about the future might be an important source of business cycle fluctuations. Indeed, the basic idea has its roots in a much earlier literature due to Beveridge (1909), Pigou (1927), and Clark (1934). These authors appealed to revisions in investor's beliefs about future growth prospects to account for business cycle expansions and contractions.

As originally emphasized by Cochrane (1994), however, introducing news shocks within a conventional business cycle framework is a non-trivial undertaking. For example, within the real business cycle framework the natural way to introduce news shocks is to have individual's beliefs about the future path of technology fluctuate. Unfortunately, news about the future path of technology introduces a wealth effect on labor supply that leads to hours moving in the opposite direction of beliefs: Expectation of higher productivity growth leads to a rise in current consumption which in turn reduces labor supply.

Much of the focus of the "news shock" literature to date has been on introducing new propagation mechanisms that deliver the correct cyclical response of hours. Beaudry and Portier (2004) introduce a two sector model with immobile labor between the sectors. Jaimovich and Rebelo (2006) introduce preferences which dampen the wealth effect on labor supply. However, as Chrisiano, Cosmin, Motto, and Rostagno (2007) note, these approaches have difficulty accounting for the high persistence of output fluctuations, as well as the volatility and cyclical behavior of stock prices. These authors instead propose a model based on overly accommodative monetary policy.

In this paper we follow the “news shock” literature in developing a framework that emphasizes revisions in beliefs about future growth prospects as key factor in business fluctuations. The framework differs, however, in that news is tied directly to the evolution of fundamentals that govern these prospects. In particular, growth prospects depend on an exogenously evolving technology frontier. The technologies in the frontier eventually will be used in production. A shock to the growth rate of potential technologies, accordingly, provides news about the future path of the technology frontier.

Unlike in the standard model, however, news about future technology is not simply news of manna from heaven. As in Comin and Gertler (2006), the new technologies have to be adopted prior to being used in production. The firms’ investments in adopting new technologies leads to a shift in labor demand when the news shock hits the economy. For reasonable parametrizations, this substitution effect offsets the wealth effect generating a boom in output, investment consumption and hours worked. This endogenous and procyclical movement of adoption is consistent with the cyclical patterns of diffusion found in Comin (2009). Further, because diffusion of new technologies takes time, the cyclical response to our news shock is highly persistent.

In addition to affecting the propagation of the innovation shock, the endogenous diffusion mechanism also works to amplify and propagate other conventional disturbances to the economy, such as exogenous movements in total factor productivity or shocks to the cost of capital investments. Thus the mechanism we develop is potentially also relevant to business fluctuations driven primarily by factors other than news about future technological prospects.

Finally, our framework also broadly captures the cyclical pattern of stock price movements. Conventional models have problems generating large procyclical movements in stock prices. In these models the value of the firm is the value of installed capital.¹ One immediate problem is that, in the data, the relative price of capital tends to move countercyclically. Of course, by introducing some form of adjustment costs, it is possible to generate procyclical movements in the market price of installed capital. However, absent counterfactually high adjustment costs, it is very difficult

¹One important deviation from this is Hall (2000) that argues that much of the run up in the second half of the 90s does not correspond to the value of installed capital.

to generate empirically reasonable movements in market prices of capital.

Unlike with standard macro models, in our framework firms have the right to the profit flow of current and future adopted technologies, in addition to the value of installed capital. Revisions in beliefs about this added component of expected earnings allow us to capture both the high volatility of the stock market and its lead over output. Further, because the stock market in our model is anticipating the earnings from projects that are productive only when they are adopted in the future, the price-earnings ratio is mean reverting, as is consistent with the evidence.²

Of course, it would be problematic that our model generated the high volatility we observe in price-earnings ratios by inducing an overly volatile or persistent earnings growth process. This is not the case. The first order auto-correlation and standard deviation of earnings growth in the model are approximately in line with the data. This then begs for the question of how, a model such as ours without cyclical variation in risk premia, can produce highly volatile price-earning ratios without overly volatile or persistent processes for dividend growth. The answer is simple: Our process of endogenous slow adoption of technologies induces a process for earnings growth that has a small but highly persistent component. This component generates low frequency fluctuations in the capital share and in earnings growth. Standard models with calibrated processes for earnings growth miss this component and hence have problems inducing large fluctuations in price-earnings ratios. Our macro model, instead, endogenously generates this component due to the endogenous technology adoption process. Reassuringly, when looking at the US macro data we also observe similar low frequency volatility in the capital share and in earnings growth.

Before proceeding we should mention a few closely related papers in the literature. Beaudry, Collard, and Portier (2007) emphasize the expansionary effect of unproductive expenditures in purchasing the rights to new technologies. In our model, instead, the expenditures in technology adoption affect the speed of diffusion of technologies. More generally, there are important differences in the details of the technology and adoption process, as well as the empirical implementation. In addition, we emphasize the implications for stock prices, as well as output and investment dynamics. Iraola and Santos (2007) and Pastor and Veronesi (2008, forthcoming) also study the implications of the arrival of new technologies for the stock market. We differ from

²See for example, Campbell and Shiller (1989).

their analysis in the details of the technology and adoption process, as well as in the empirical implementation.

In section 2 we present a simple expository model to introduce the endogenous technology adoption mechanism and our innovation shock as a prelude to an estimated model that we present in section 4. The model adds to a relatively standard real business model an expanding variety of intermediate goods which determines the level of productivity. Though intermediate goods arrive at an exogenous rate, how many can be used in production depends on the agents' adoption decisions. In section 3 we calibrate the model and analyze the impact of a shock to the evolution of new technologies. As we noted, assuming rational expectations, this shock reveals news about the economy's future growth potential.

In section 4, we move to an estimated model. We combine our model of endogenous technology adoption with a variant of the standard quantitative macroeconomic model due to Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). We differ mainly by having technological change endogenous whereas in the standard model it is exogenous. Section 5 reports the estimates for a sample period covering 1984:1 to 2008:2. Overall, we show that the main findings from the calibrated model are robust to an estimated model that provides a reasonable fit of the data. In addition, our "news/innovation" shock is an important driver of business fluctuations. In particular, it explains 27 percent of output growth (32 percent of HP filtered output).

In section 6 we analyze the implications for the stock market. We show that, broadly speaking, the model captures the overall volatility of stock prices, as well as the co-movement with output. We also show that the model is consistent with a number of findings from the empirical finance literature. Somewhat surprisingly, it can account for the run-up of stock prices in the mid 1990s and also some of the decline preceding the most recent recession. Concluding remarks are presented in section 7.

2 Baseline Model

Our baseline framework is a variation of the Greenwood, Hercowitz, and Krusell (2000) (GHK) business cycle model that features shocks to embodied technological

change, as well as a variable utilization rate of capital. We formulate the process of technological change more explicitly and also allow for endogenous technology adoption.

2.1 Resource Constraints

Let Y_t be gross final output, C_t consumption, I_t investment, G_t government consumption, H_t technology adoption expenses, and O_t firm overhead operating expenses. Then output is divided as follows:

$$Y_t = C_t + I_t + G_t + H_t + O_t \quad (1)$$

In turn, let J_t be newly produced capital and δ_t be the depreciation rate of capital. Then capital, K_t , evolves as follows:

$$K_{t+1} = (1 - \delta_t)K_t + J_t \quad (2)$$

Next, let P_t^k be the price of this capital in units of final output which is our numeraire. Then,

$$J_t = (P_t^k)^{-1} \bar{\mu}^k I_t$$

where $\bar{\mu}^k$ is a weighted markup in the capital goods sector to be characterized below. A distinguishing feature of our framework is that P_t^k evolves endogenously. One key source of variation is the pace of technology adoption, which depends on the stock of available new technologies, as well as overall macroeconomic conditions, as we eventually describe.

2.2 Production

There are two production sectors: one for new capital, J_t , and one for output, Y_t . Within each in sector there are several stages of production.

New capital

A continuum of N_t^k monopolistically competitive firms produce differentiated final capital goods. The aggregate J_t is a CES composite of a continuum of these

differentiated goods as follows:

$$J_t = \left(\int_0^{N_t^k} J_t(r)^{\frac{1}{\mu^k}} dr \right)^{\mu^k}, \text{ with } \mu^k > 1, \quad (3)$$

where $J_t(r)$ is the output produced by the r^{th} final capital goods producer. Free entry determines N_t^k , as we describe below. The parameter μ^k is inversely related to the price elasticity of substitution across new capital goods.

To produce a differentiated capital good, r , a producer combines new structures ($J_t^s(r)$) and new equipment ($J_t^e(r)$) as follows:

$$J_t(r) = \bar{\gamma} (J_t^s(r))^\gamma (J_t^e(r))^{1-\gamma}, \text{ with } \gamma \in (0, 1) \text{ and } \bar{\gamma} = [\gamma^\gamma (1-\gamma)^{1-\gamma}]^{-1} \quad (4)$$

We distinguish between equipment investment and other forms of investment, which we generically label “structures”, for two reasons. First, as emphasized in GHK, embodied technology change influences mainly equipment investment, making it important to disentangle the different forms of capital. Second, over our sample there have been significant fluctuations in both commercial and residential structures that are more likely due to factors such as credit conditions and taxes changes than technological change. By introducing an independent disturbance to structures we can capture these factors, at least in a reduced form way.

Formally, the r^{th} capital producer can obtain a unit of structures from P_t^{st} units of final output, where $p_t^{st} (\equiv \log(P_t^{st}))$ evolves exogenously according to:

$$p_t^{st} = \rho^{st} p_{t-1}^{st} + \varepsilon_t^{st}$$

where ε_{st} is a stationary first order disturbance. Generally speaking, p_t^{st} , reflects any factors that could affect the cost of producing structures.

To produce equipment, the r^{th} capital producer uses the A_t^k intermediate capital goods that have been adopted up to time t . In particular, let $I_t^r(s)$ the amount of intermediate capital from supplier s that final capital producer r demands. Then, equipment $J_t^e(r)$ is the following CES composite:

$$J_t^e(r) = \left(\int_0^{A_t^k} I_t^r(s)^{\frac{1}{\theta}} ds \right)^\theta, \text{ with } \theta > 1. \quad (5)$$

where the parameter θ is inversely related to the price elasticity of substitution across intermediate capital goods. The evolution of A_t^k depends on the endogenous technology adoption process that we describe shortly. Observe that there are efficiency gains in producing new equipment from increasing A_t^k . These efficiency gains are ultimately what creates the incentive to adopt new technologies, as we discuss below.

Intermediate capital goods, in turn, use final output as input. To produce one unit of an existing type of intermediate capital goods, a supplier uses one unit of final output, which fixes the marginal cost at unity. Because the supplier has a bit of market power it can charge the final capital goods producer a fixed markup which, given the CES structure, equals θ .

Output

The composite Y_t is a CES aggregate of the output of N_t^y differentiated final goods producers. Let $Y_t(j)$ is the output of producer j . Then:

$$Y_t = \left(\int_0^{N_t^y} Y_t(j)^{\frac{1}{\mu}} dj \right)^{\mu}, \text{ with } \mu > 1, \quad (6)$$

where μ is inversely related to the price elasticity of substitution across goods. As in the capital goods sector, entry and exit determines the number of firms operating.

As do final capital goods firms, final output goods firms use differentiated intermediate inputs. Let $Y_t^j(s)$ the amount of an intermediate good that final goods firm j employs from supplier s and let A_t^y denote the total number of intermediate inputs. Then

$$Y_t(j) = \left(\int_0^{A_t^y} Y_t^j(s)^{\frac{1}{\vartheta}} ds \right)^{\vartheta} \quad (7)$$

Just as with capital goods, an expanding variety of intermediate output goods increases the efficiency of producing final output goods. As we show, this efficiency gain will be reflected in total factor productivity. Similarly, just as with A_t^k , the evolution of A_t^y will depend on endogenous technology adoption.

Intermediate goods used in the output sector are produced using the following Cobb-Douglas technology:

$$Y_t(s) \equiv \int_0^{N_t^y} Y_t^j(s) dj = X_t (U_t(s)K_t(s))^{\alpha} (L_t(s))^{1-\alpha}$$

where X_t is the level of disembodied productivity, U_t denotes the intensity of utilization of capital, and $K_t(s)$ and $L_t(s)$ are the amount of capital and labor rented (hired) to produce the s^{th} intermediate good.

We assume that $x_t(\equiv \log(X_t))$ evolves as follows

$$x_t = x_{t-1} + \varsigma_t \quad (8)$$

where ς_t is first order serially correlated innovation. Given that total factor productivity will depend on X_t , N_t^y and A_t^y , the model allows for both exogenous and endogenous movements in total factor productivity. By estimating the model in section 5, we let the data tell the relative importance of each.

Finally, following Greenwood, Hercowitz, and Huffman (1988), we further assume that a higher rate of capital utilization comes at the cost of a faster depreciation rate, δ . The markets where firms rent the factors of production (i.e. labor and capital) are perfectly competitive.

Free entry

We now characterize the free entry decision that determines the number of producers in the final capital and output goods sectors, N_t^k and N_t^y , respectively: We assume that the per period operating cost of a final goods producer in sector s , o_t^s is

$$o_t^s = b^s \bar{P}_t^k K_t, \text{ for } s = \{y, k\} \quad (9)$$

where b^s is a constant and \bar{P}_t^k is the wholesale price of capital. That is, in order to have balanced growth, the operating costs grow with the replacement value of the capital stock, a measure of the technological sophistication of the economy. In any period, the producer profits must cover this operating cost. Everything else equal, firm profits are decreasing in the total number of firms. Accordingly, free entry pins down both N_t^k and N_t^y .

2.3 Technology

The efficiency of production depends on the exogenous productivity variables (X_t , and P_{st}^k) and on the number of "adopted" intermediate goods in the production of capital, A_t^k , and final output, A_t^y . We characterize next the process that governs the evolution of these variables.

New intermediate goods

Prototypes of new intermediate goods arrive exogenously to the economy. Upon arrival, they are not yet usable for production. In order to be usable, a new prototype must be successfully adopted. This, however, involves a costly investment that we describe below.

Let Z_t^s denote the total number of intermediate goods in sector s (for $s = \{k, y\}$) at time t . Note that Z_t^s includes both previously adopted goods and “not yet adopted” prototypes. The law of motion for Z_t^s is as follows:

$$Z_{t+1}^s = (\bar{\chi}_s \chi_t^{\xi_s} + \phi) Z_t^s \quad (10)$$

where ϕ is the fraction of intermediate goods that do not become obsolete, and χ_t determines the stochastic growth rate of the number of prototypes and is governed by the following AR(1) process

$$\log \chi_t = \rho \log \chi_{t-1} + \varepsilon_t$$

where ε_t is a white noise disturbance.

Note that the shock to the growth rate of intermediate goods is the same across sectors. However, the effect of the shock on the stock of technologies within a sector, measured by the slope coefficient $\bar{\chi}_s$ and the elasticity ξ_s , differs across sectors. Here we wish to capture the idea of spillovers in the innovation process: Innovations that lead to new equipment often make possible new disembodied innovations. For example, the IT revolution made possible e-commerce. It also accelerated the offshoring process and improved the efficiency of inventories management, and so on.

Evidence of this spillover appears in the data: At medium frequencies, movements in relative equipment prices are correlated with movements in TFP. As we show shortly, given that a component of TFP in our model is exogenous, we can calibrate ξ_s to capture this correlation.

We emphasize that in this framework, news about future growth prospects, captured by innovations in χ_t , govern the growth of potential new intermediate goods. Realizing the benefits of these new technologies, however, requires a costly adoption process that we turn to next.

Adoption (Conversion of Z to A)

At each point in time a continuum of unexploited technologies is available to be adopted. Through a competitive process, firms that specialize in adoption try to make these technologies usable. These firms, which are owned by households, spend resources attempting to adopt the new goods, which they can then sell on the open market. They succeed with an endogenously determined probability λ_t^s , for $s = \{k, y\}$. Once a technology is usable, any producer can use it in production immediately.

Note that under this setup there is slow diffusion of new technologies on average (as they are slow on average to become usable) but aggregation is simple as once a technology is in use, all firms have it. Consistent with the evidence (e.g. Comin, 2009), we obtain a pro-cyclical adoption behavior by endogenizing the probability λ_t^s that a new technology becomes usable, and making it increasing in the amount of resources devoted to adoption at the firm level.

Specifically, the adoption process works as follows. To try to make one prototype usable at time $t + 1$, an adopting firm spends h_t^s units of final output at time t . Its success probability λ_t^s is given by

$$\lambda_t^s = \bar{\lambda}^s (\Gamma_t^s h_t^s)^{\rho_\lambda}$$

with $\bar{\lambda}^s > 0$, $0 < \rho_\lambda < 1$, and where Γ_t^s is a factor that is exogenous to the firm, given by

$$\Gamma_t^s = A_t^s / o_t^s$$

We presume that past experience with adoption, measured by the total number of projects adopted A_t^s , makes the process more efficient. In addition to having some plausibility, this assumption ensures that the fraction of output devoted to adoption is constant along the balanced growth path.

The value to the adopter of successfully bringing a new technology into use v_t^s , is given by the present value of profits from operating the technology. Profits π_t^s arise from the monopolistic power of the producer of the new good. Accordingly, given that $\beta\Lambda_{t,t+1}$ is the adopter's stochastic discount factor for returns between $t + 1$ and t , we can express v_t^s as

$$v_t^s = \pi_t^s + \phi E_t [\beta\Lambda_{t,t+1} v_{t+1}^s]. \quad (11)$$

If an adopter is unsuccessful in the current period, he may try again in the subsequent periods to make the technology usable. Let j_t^s be the value of acquiring an

innovation that has not been adopted yet. j_t^s is given by

$$j_t^s = \max_{h_t^s} -h_t^s + E_t\{\beta\Lambda_{t,t+1}\phi[\lambda_t^s v_{t+1}^s + (1 - \lambda_t^s)j_{t+1}^s]\} \quad (12)$$

Optimal investment in adopting a new technology is given by:

$$1 = E_t \left[\beta\Lambda_{t,t+1}\phi\rho\lambda\bar{\lambda}^s (\Gamma_t^s)^{\rho\lambda} (h_t^s)^{\rho\lambda-1} (v_{t+1}^s - j_{t+1}^s) \right] \quad (13)$$

It is easy to see that h_t^s is increasing in $v_{t+1}^s - j_{t+1}^s$, implying that adoption expenditures, and thus the speed of adoption, are likely to be procyclical. Note also that the choice of h_t^s does not depend on any firm specific characteristics. Thus in equilibrium, the success probability is the same for all firms attempting adoption.

2.4 Households

The household sector is reasonably standard. In particular, there is a representative household that consumes, supplies labor and saves. It may save by either accumulating capital or lending to innovators and adopters. The household also has equity claims in all monopolistically competitive firms. It makes one period loans to adopters and also rents capital that it has accumulated directly to firms.

Let C_t be consumption. Then the household maximizes the present discounted utility as given by the following expression:

$$E_t \sum_{i=0}^{\infty} \beta^i \left[\ln C_{t+i} - \mu^w \frac{(L_{t+i})^{1+\zeta}}{1+\zeta} \right] \quad (14)$$

with $\zeta > 0$. The budget constraint is as follows:

$$C_t = W_t L_t + \Pi_t + [D_t + P_t^k]K_t - P_t^k K_{t+1} + R_t B_t - B_{t+1} - T_t \quad (15)$$

where Π_t reflects the profits of monopolistic competitors paid out fully as dividends to households, B_t is total loans the households makes at $t-1$ that are payable at t , and T_t reflects lump sum taxes which are used to pay for government expenditures. The household's decision problem is simply to choose consumption, labor supply, capital and bonds to maximize equation (14) subject to (15).

For the calibrated model we keep the preference parameters β and μ^w fixed. Once we turn to estimation in section 5 we allow these parameters to follow stationary stochastic processes in order to achieve identification.

2.5 Symmetric equilibrium

We defer to the Appendix the formal definition of equilibrium, as a complete characterization of all the relationships. Here we just present the main equations to highlight some key differences with the basic real business cycle (RBC) model, as well as the variation proposed by GHK.

In the canonical RBC model, capital is the only endogenous state. Here there are two additional endogenous states, the stocks of adopted technologies in the output and capital production sectors, A_t^y and A_t^k , respectively. The relevant equations of motion are thus given by:

$$K_{t+1} = (1 - \delta(U_t))K_t + (P_t^k)^{-1} \bar{\mu}^k I_t, \quad (16)$$

$$A_{t+1}^s = \lambda_t^s [Z_t^s - A_t^s] + \phi A_t^s, \text{ for } s = \{k, y\}. \quad (17)$$

with

$$P_t^k = \mu^k (N_t^k)^{-(\mu^k - 1)} (P_t^{st})^\gamma \left(\theta (A_t^k)^{-(\theta - 1)} \right)^{1 - \gamma} \quad (18)$$

$$\lambda_t^s = \bar{\lambda}^s \left(\frac{A_t^s h_t^s}{o_t^s} \right)^{\rho_\lambda} \quad (19)$$

and where the evolution of the stock of new technologies in each sector, Z_t^s , is given by equation (36). Note that also in contrast to both the RBC model and GHK, the relative price of capital depends positively on the stock of adopted technologies in the capital goods sector, as measured by A_t^k , as well as the degree entry, measured by N_t^k . In addition, the fraction of unadopted technologies that come on line, λ_t^s , depends on endogenously determined adoption expenditures, h_t^s , and is likely to vary procyclically, as equation (13) suggests.

In turn, aggregate production, consumption/saving, and factor market equilibria are given by

$$Y_t = X_t (A_t^y)^{\theta - 1} (N_t^y)^{\mu - 1} (U_t K_t)^\alpha L_t^{1 - \alpha} \quad (20)$$

$$E_t \left\{ \beta C_t / C_{t+1} \cdot \left[\alpha \frac{Y_{t+1}}{\mu K_{t+1}} + (1 - \delta(U_{t+1}) P_{t+1}^k) / P_t^k \right] / P_t^k \right\} = 1 \quad (21)$$

$$(1 - \alpha) \frac{Y_t}{L_t} = \mu \mu^w L_t^\zeta / (1 / C_t) \quad (22)$$

$$\alpha \frac{Y_t}{U_t} = \mu \delta'(U_t) P_t^k K_t \quad (23)$$

and where equation (1) gives the economy-wide resource constraint. In contrast to the standard formulation (with variable utilization of capital), total factor productivity is endogenous and depends on both the stock of adopted technologies in the output sector, A_t^y , as well as the degree of entry, N_t^y . Thus not only does embodied technological change depend on adoption and cyclical entry, the same is true for disembodied technological change.

We can now get a sense of how “news” about technology plays out in this model. In our model, the news is about future technological prospects (i.e., future values of Z_t), as opposed to future technology per se. For those prospects to be realized, resources need to be invested in adopting these potential technologies. As a result, the news shock sparks a contemporaneous rise in aggregate demand driven by the desire to speed up adoption (32). Output increases to meet the rise in demand via three channels: a rise in the utilization rate, increased entry, and a rise in hours worked. Two factors work to offset the standard wealth effect, which produces a decline in hours in the standard model. First, the response of utilization and entry to increased demand (stemming from increased adoption expenditures) raises the marginal product of labor, everything else equal, enhancing the rise in labor demand. Second, given that adoption expenditures are effectively a form of saving, consumption increases by less than it might otherwise in the standard model, as households substitute some current consumption for increased investment in technology. This moderating of the consumption rise, dampens the negative wealth effect on labor supply.

As we illustrate in the next section, in contrast to the standard model, news about improved technological prospects increases current output, hours and consumption. Endogenous technology adoption plays a critical role, along with endogenous utilization of capital. Endogenous entry improves the quantitative performance, but is not needed for the main qualitative arguments regarding the cyclical responses to news about future technological prospects. Finally, given that there are rents associated with both adopted and prospective technologies, there are as well implications for the cyclical behavior of asset prices. However, we defer a discussion of asset prices until section 6.

3 Model Simulations of “Innovation” Shocks

In this section we present simulations of the impact of a shock to the growth rate of prospective new technologies. As we have been noting, one can interpret this shock as capturing news about the economy’s growth potential. Our goal here is to elucidate the basic mechanisms. Thus, we work with a calibrated version of our simple baseline model. In the subsequent section we enrich the model to enable it to capture short run cyclical dynamics and also to estimate most the model parameters..

3.1 Calibration

The calibration we present here is meant as a reasonable benchmark. The model’s behavior is robust to modest variations around this benchmark.

To the extent possible, we use the restrictions of balanced growth to pin down parameter values. Otherwise, we look for evidence elsewhere in the literature. There are a total of eighteen parameters. Ten appear routinely in other studies. The other eight relate to the adoption processes and also to the entry/exit mechanism. Table 1 reports the value for these parameters. We defer the discussion of the calibration of the standard parameters and of the more trivial non-standard parameters to the Appendix.

There two key sets of parameters that are specific to our model. The first is the sectoral elasticity parameter in equation (36) that governs sensitivity of the growth rate of potential new technologies to movements in the exogenous disturbance χ_t . We normalize the elasticity for the creation of new capital goods technologies, ξ_k , to unity. The elasticity for the creation of new output goods technologies, ξ_y , affects the correlation between TFP growth and the growth rate of the relative price of equipment, particularly at medium and low frequencies where cyclical factors are less important. We can accordingly use information about this co-movement at medium and low frequencies (i.e. cycles with periods between 8 and 50 years, following Comin and Gertler (2006)).to pin down ξ_y . In particular, our model implies that the covariance between medium term growth in TFP, and the relative price of equipment, and their variances depend on the variance of χ_t , the variance of x_t (the exogenous component of TFP) and ξ_y . Hence, we can use these three moments in the data to identify ξ_y . This yields an estimate for ξ_y of approximately 0.6. Our results are quite robust to

variation in ξ_y between 0.5 and 0.8.

The second set of parameters are the two that govern the technology adoption process in equation lambda (19), $\bar{\lambda}^s$ and ρ_λ . $\bar{\lambda}^s$ governs the average adoption lag and ρ_λ governs the elasticity of adoption with respect to adoption investments. We set $\bar{\lambda}^s$ so that the average adoption lag is approximately 5 years which is a reasonable benchmark within the productivity literature (e.g. Mansfield, 1989). We set ρ_λ to 0.9 to match a time series regression of the rate of decline in the relative price of capital on US adoption expenditures measured by development costs by the NSF.³

3.2 Model Simulations

We now analyze the effect of a positive shock to the growth rate of new technologies. To compare with the literature, we first consider a variation of the model that eliminates the key features we have introduced.. In particular we suppose that technology diffusion is instantaneous and exogenous and that firm entry and exit is shut off. In this case, our experiment closely mimics the "news" shock scenario analyzed in the literature: The expected increase in the arrival of new technologies leads to an expected increase in the growth rate of productivity that is independent of any actions that individual firms or households make.⁴ As Figure 1 shows, the increase in the expected new technology arrival rate initially reduces labor supply and output. At work is the wealth effect, noted by Cochrane (1994) and many others.

We next return to our baseline model by adding back the relevant features. In this instance, as Figure 2 shows, the increase in the expected technology arrival rate produces an initial increase in both output and hours. Now the increase in expected productivity growth is not simply manna from heaven. Rather, it may be realized only if resources are devoted to technology adoption. Further, the more resources are devoted, the faster the technology will be adopted. The initial increase in labor demand in part reflects an intertemporal substitution effect: Because more labor and capital is needed for adoption in the future, it is optimal to build up the capital stock

³This estimate is consistent with the very high pro-cyclicality of the speed of adoption estimated by Comin (2009).

⁴The arrival of new technologies simultaneously affects both future disembodied and embodied technological change, as in our baseline model. The results are qualitatively the same if the shock just affects one type of technology change or the other.

today, before the technologies come in line. The associated rise in capital utilization and entry increases the marginal product of labor, everything else equal, contributing to the increase in labor demand. This in turn leads to an increase in real wages and labor supply.

What is key to producing a positive co-movement between output and expected technology growth is the combination of slow diffusion and costly adoption. We illustrate this point in Figure 3 by examining the response of output and hours for different variations of the model. The top panel is our baseline. In the second panel we keep endogenous adoption but remove entry and exit. As the figure shows, the output and hours responses is weaker than in the baseline case, but qualitatively the same. One other difference, is that consumption declines initially. By contrast, the agglomeration effect from entry in our baseline boosts output sufficiently to introduce an increase in consumption. In the bottom panel we also remove endogenous adoption. New technologies diffuse exogenously at the same rate as in the steady state of our baseline. As the panel shows, output and hours decline at the onset of the shock, as in the conventional literature. Thus it appears that within our framework endogenous technology adoption is key to getting the right co-movement.

Though we do not report the results here, endogenous entry alone does not generate the right quantitative co-movements in response to innovation shocks.⁵ Entry interacts with endogenous adoption to magnify the overall response of real activity. Intuitively, the agglomeration effects from entry expand output and investment, which in turn raises profitability and enhances the incentives to adopt.

Finally, it is the case, as in Comin and Gertler (2006), that the endogenous technology feature of our model introduces a significant propagation mechanism that operates over the medium term. The acceleration in the speed of adoption after a news shock improves the overall efficiency of production of capital and output as reflected, respectively, by the medium and long term fluctuations of the relative price of capital and TFP.

As we show in section 5, the mechanism we have just outlined propagates not only the innovation shock but also other shocks that may disturb the economy.

⁵The impulse response functions for this case are reported in the extended estimated model below.

4 An Extended Model for Estimation

In this section we generalize our model and then estimate it. We add some key features that have proven to be helpful in permitting the conventional macroeconomic models (e.g. Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007)) to capture the data. Our purpose here is twofold. First we wish to assess whether the effects of our news shock that we identified in our baseline model are robust in a framework that provides an empirically reasonable description of the data. Second, by proceeding this way, we can formally assess the contribution of our innovation shock as we have formulated them to overall business cycle volatility.

4.1 The Extended Model

The features we add include: habit formation in consumption, flow investment adjustment costs, nominal price stickiness in the form of staggered price setting, and a monetary policy rule.

To introduce habit formation, we modify household preferences to allow utility to depend on lagged consumption as well as current consumption in the following simple way:

$$E_t \sum_{i=0}^{\infty} \beta^i b_{t+i} \left[\ln(C_{t+i} - vC_{t+i-1}) - \mu_{t+i}^w \frac{(L_{t+i})^{1+\zeta}}{1+\zeta} \right] \quad (24)$$

where the parameter v , which we estimate, measures the degree of habit formation. In addition, the formulation allows for two exogenous disturbances: b_t is a shock to household's subjective discount factor and μ_t^w is a shock to the relative weight on leisure. The former introduces a disturbance to consumption demand and the latter to labor supply. Adding flow adjustment costs leads to the following formulation for the evolution of capital:

$$K_{t+1} = (1 - \delta_t)K_t + J_t \left(1 - \eta \left(\frac{J_t}{(1 + g_K)J_{t-1}} - 1 \right)^2 \right) \quad (25)$$

with $J_t = (P_t^k)^{-1} \bar{\mu}^k I_t$. η , another parameter we estimate, measures the degree of adjustment costs and g_K is the steady state growth rate of capital. We note that the adjustment costs are external and not at the firm level. Capital is perfectly

mobile between firms. In the standard formulation (e.g. Justiniano, Primiceri, and Tambalotti (2008)), the relative price of capital is an exogenous disturbance. In our model it is endogenous. As equation (52) suggests, P_t^k depends inversely on the volume of adopted technologies A_t^k and the cyclical intensity of production of new capital goods, as measured by N_t^k .

We model nominal price rigidities by assuming the final output goods producing firms (6) set nominal prices on a staggered basis. For convenience, we now restrict entry in this sector and instead fix the number of these firms at the steady state value N . Following Smets and Wouters (2007) and Justiniano, Primiceri, and Tambalotti (2008), we used a formulation of staggered price setting due to Calvo (1983), modified to allow for partial indexing. In particular, every period a fraction $1 - \xi$ are free to optimally reset their respective price. A fraction ξ instead adjust price according to a simple indexing rule based on lagged inflation. Let $P_t(j)$ be the nominal price of firm j 's output, P_t the price index and $\Pi_{t-1} = P_t/P_{t-1}$ the inflation rate. Then, the indexing rule is given by:

$$P_{t+1}(j) = P_t(j) (\Pi_t)^{\iota_p} (\Pi)^{1-\iota_p} \quad (26)$$

where Π and ι_p are parameters that we estimate: the former is the steady state rate of inflation and the latter is the degree of partial indexation. The fraction of firms that are free to adjust, choose the optimal reset price P_t^* to maximize expected discounted profits given by.

$$E_t \sum_{s=0}^{\infty} \xi^s \beta^s \Lambda_{t,s} \left\{ \left[\frac{P_t^*}{P_{t+s}} \left(\prod_{j=0}^s (\Pi_{t+j})^{\iota_p} (\Pi)^{1-\iota_p} \right) \right] Y_{t+s}(j) - W_{t+s} N_{t+s}(j) - D_{t+s} K_{t+s}(j) \right\} \quad (27)$$

given the demand function for firm j 's product (obtained from cost minimization by final goods firms):

$$Y_t(j) = \left(\frac{P_t(j)}{P_t} \right)^{\frac{-\mu}{\mu-1}} Y_t \quad (28)$$

Given the law of large numbers and given the price index, the price level evolves according to

$$P_t = [(1 - \xi)(P_t^*)^{\frac{\mu-1}{\mu}} + \xi(P_{t-1})^{\frac{\mu-1}{\mu}}]^{\frac{\mu}{\mu-1}} \quad (29)$$

Finally, define R_t^n as the nominal rate of interest, defined by the Fisher relation $R_{t+1} = R_t^n E_t \Pi_{t+1}$. The central bank sets the nominal interest rate R_t^n according to a

simple Taylor rule with interest rate smoothing, as follows:

$$\frac{R_t^n}{R^n} = \left(\frac{R_{t-1}^n}{R^n}\right)^{\rho_r} \left(\left(\frac{\Pi_t}{\Pi}\right)^{\phi_p} \left(\frac{Y_t}{Y_t^0}\right)^{\phi_y}\right)^{1-\rho_r} \exp(\mu_{mp,t}) \quad (30)$$

where R^n is the steady state of the gross nominal interest rate and Y_t^0 is trend output, and $\mu_{mp,t}$ is an exogenous shock to the policy rule.

Including habit formation and flow investment adjustment costs give the model more flexibility to capture output, investment, and consumption dynamics. We include nominal rigidities and a Taylor rule for two reasons. First, doing so allows us to use the model to identify the real interest rate which enters the first order conditions for both consumption and investment. The nominal interest rate is observable but expected inflation is not. However, from the model we identify expected inflation. Second, with nominal rigidities, the market real interest rate need not equal the flexible price equilibrium real rate of interest (i.e. the "natural rate of interest".) This will permit the model to simultaneously account for the relatively smooth behavior of observed market real interest rates and relatively volatile behavior of asset prices, as section 6 makes clear.⁶

We emphasize that the critical difference in our framework is the endogenous component of both embodied and disembodied productivity. The standard model treats the evolution of both of these phenomena as exogenous disturbances. In our model the key primitive is the innovation process. Shocks to this process influence the pace of new technological opportunities which are realized only by a costly adoption process.

5 Estimation

5.1 Data and Estimation Strategy

We estimate the model using quarterly data from 1984:I to 2008:II on seven key variables in the US economy: output, consumption, equipment investment, non-equipment investment, inflation, nominal interest rates and hours. The Appendix describes the sources and transformations of the data used in the estimation.

⁶One widely employed friction that we do not add is nominal wage rigidity. While adding this feature would help improve the ability of the model in certain dimensions, we felt that at least for this initial pass at the data, the cost of added complexity outweighed the marginal gain in fit.

The model contains seven structural shocks. Five appear in the standard models: the household’s subjective discount factor, the household’s preference for leisure, government consumption; the monetary policy rule, and the growth rate of TFP. The key new shock in our model is the disturbance to the growth rate of potential new intermediate capital goods, which we refer to as an ”innovation” shock. As we have been noting, since this shock signals opportunities for future growth, it is also similar in spirit to a ”news shock”. Finally, we allow for an exogenous shock to the cost of producing non-equipment investment, but are agnostic about the deep underlying source of this shock.

We continue to calibrate the parameters of the embodied technology process. However, as in the standard quantitative macroeconomic framework we estimate the rest of the parameters of the model, using Bayesian techniques, as in An and Schorfheide (2007)

5.2 Priors and Posterior Estimates and Model Fit

Table 2 presents the prior distributions for the structural parameters along with the posterior estimates. Tables 3 presents the same information for the serial correlation and standard deviation of the stochastic processes. To maintain comparability with the literature, for the most part we employ the same priors as in Justiniano, Primiceri, and Tambalotti (2008). Overall, the parameter estimates are very close to what has been obtained elsewhere in the literature (e.g. Smets and Wouters (2007), Primiceri, Schaumburg, and Tambalotti (2006), and Justiniano, Primiceri, and Tambalotti (2008)).

To get a sense of how well our model captures the data, Table 4 presents the standard deviations of several selected variables in our model and in two reasonable competing alternatives. These are as follows: The first is our model but with endogenous adoption shut off. The second is a version of the conventional DSGE model. In particular, we make diffusion instantaneous, shut off entry, and also eliminate the distinction between equipment and structures. In effect, this alternative model is identical to Primiceri, Schaumburg, and Tambalotti (2006) and very similar to Christiano, Eichenbaum, and Evans (2005), and Smets and Wouters (2007), though without wage rigidity (in order to be comparable to our baseline model). Overall, our

baseline model is in line with the data while the two alternatives do a poorer job in matching the volatility of output and investment and (for the model with exogenous adoption) hours worked. More formally, Table 5 shows that of the three alternatives, the marginal likelihood for our baseline model is highest. Intuitively, the endogenous adoption structure allows for more flexible lag dynamics, which improves the ability of the model to fit the data.

To assess how important the innovation shock is as a business cycle driving force, Tables 6 and 7 report the contribution of each shock to the unconditional variance of output, consumption and equipment and structures investment and hours worked. We explore the variance decomposition both for the growth rate (Table 6) and the HP filtered level (Table 7).

The innovation shock accounts for 27 percent of output growth fluctuations and 32 in HP filtered output. It is of nearly equal importance to the neutral technology shock, which accounts for 43 percent of fluctuations in output growth and 34 percent in HP filtered output. Investment shocks combined, however, account for more the half the high frequency variation in output, in keeping with the findings of Justiniano, Primiceri, and Tambalotti (2008). The difference in our model is that we disentangle shocks to equipment versus non-equipment investment and also endogenize the pace of technological change. The shock to non-equipment investment is the third most important in explaining approximately 11 percent of output growth fluctuations, and 25 percent of HP filtered output. The other 4 shocks seem much less important in explaining output fluctuations, representing a combined 20 percent of output growth fluctuations and less than 9 percent of HP filtered output.

5.3 Estimated Impulse Response Functions

Next we analyze the impulse responses to our innovation/news shock using the estimated model. Figure 2 presents the results for our model (solid line) and for the version with exogenous adoption (dashed line).⁷ As Figure 4 shows, the qualitative patterns are very similar to what we obtained from the calibrated model. The economy with exogenous adoption experiences a recession in response to a positive news shock. In contrast, in our model, there is a positive and prolonged response of output,

⁷Just to be clear, the version with exogenous adoption has also endogenous entry, as our model.

investment, consumption and hours worked.

In contrast to the simple calibrated model we analyzed earlier, the responses of output and investment in the estimated model are humped-shaped, reflecting the various real frictions such as investment adjustment costs that are now present. The response of hours relative to output, however, is somewhat weaker. The introduction of the various frictions has likely dampened the overall hours response. This is somewhat mitigated in conventional models by incorporating wage rigidity.

The speed of technology adoption (first panel in the third row) strongly reacts to the arrival of news about future technology. This is the case because of the sharp increase in the value of new adopted technologies in response to the news shock (second panel in the third row). As we shall see below, this mechanism plays a key role in inducing fluctuations in the stock market.

The estimated model not only delivers a plausible response to the innovation shock, but does so to the other shocks as well. Figures 5 and 6 report the impulse response functions of our baseline model to the structures shock and to the neutral technology shock (solid lines). (To save space we only report results for the major shocks, but the responses to the other shocks are reasonable as well.) As with a positive news shock, a positive shock to TFP or to structures leads to an increase in output, hours, investment and adoption expenses. In response to a TFP shock, consumption, initially, experiences a very small decline due to the large substitution effect introduced by technology adoption and entry. After, that, consumption increases. For the shock to structures, instead, consumption is pro-cyclical. It is also worth noting that, because these shocks induce pro-cyclical fluctuations in the value of adopted technologies, they also generate large, pro-cyclical fluctuations in the speed of adoption of new technologies.

In Figures 5 and 6 we also report (in dashed lines) the impulse responses to the structures TFP shocks of the version of our model with exogenous adoption (i.e. constant λ^s , for $s = \{k, y\}$). One striking observation from this figures is that the response of the models to these shocks is significantly more muted when adoption is exogenous than when it is endogenous. Accordingly, the endogenous adoption mechanism greatly amplifies the model's response not only to the news shock but also to the other shocks considered here. Thus, even in instances where our innovation shock is not the key driving force, the endogenous technology mechanism we have

characterized may be relevant.

5.4 Historical Decompositions

To get a better feel for the role of our innovation shock and the two other major shocks, structure and TFP, in output fluctuations, we present a historical decomposition of the data. Figure 7 present three panels. Each plots the contribution to output growth the model implies for one of the three major shocks. The top panel reports results for the innovation shock, the middle for the structures shock, and the bottom for the TFP shock.

As the top panel indicates, the innovation shock contributes significantly to cyclical output growth. In particular, the shock seems to play a prominent role in recessions and early stages of the expansions. As one might expect, it also appears to play a role in the late 1990s period of high output and productivity growth.

The structures shock is very important in the recession of the early 1990s and also the period of slow growth at the end of our sample, which just precedes the most recent recession. These results are consistent with the role that the contraction in commercial structures played in the 1990s recessions and the collapse of housing investment in the very recent period. In each instance, of course, credit conditions likely influenced the slowdown in structures. In this respect, our structures shock may capture in a reduced form way the influence of credit conditions. A more explicit modeling of this phenomenon would be of interest, though.

6 The Stock Market

6.1 Theory

In standard macro models, the market value of corporations is equal to the value of installed capital. This creates a serious challenge for these models. Since capital is a stock, the short run evolution of the value of installed capital is driven by the dynamics of the price of installed capital, which for reasonable adjustment costs is not very different from the price of new capital. In the data, the price of new capital is countercyclical and moves approximately as much as output. The stock market,

however, is strongly pro-cyclical and moves about ten times more than output. A theory that equalizes the two variables will have to be inconsistent with the empirical behavior of at least one of the two.

Unlike standard macro models, in our framework firms have the rights to the profit flows from selling current and future adopted technologies. Thus, the market value of companies is given by the present discounted value of these profits in addition to the value of installed capital. Formally, the value of the stock market Q_t is composed of four terms as shown in (31).

$$\begin{aligned}
Q_t = & \overbrace{P_t^{insk} K_t}^{\text{Value of installed capital}} + \sum_{s=\{k,y\}} \overbrace{A_t^s (v_t^s - \pi_t^s)}^{\text{Value of adopted technologies}} \\
& + \sum_{s=\{k,y\}} \overbrace{(j_t^s + h_t^s)(Z_t^s - A_t^s)}^{\text{Value of existing not adopted technologies}} + E_t \left[\sum_{s=\{k,y\}} \overbrace{\sum_{i=0}^{\infty} \beta \Lambda_{t,t+1+i} j_{t+1+i}^s (Z_{t+1+i}^s - \phi Z_{\tau+i}^s)}^{\text{Value of future non-adopted technologies}} \right]
\end{aligned} \tag{31}$$

where P_t^{insk} is the value of a unit of installed capital in the firm (i.e. the shadow value of a unit of capital to the firm). We note that Iraola and Santos (2007) have previously derived a similar expression for stock market value, also based on a framework in the spirit of Comin and Gertler (2006).

The first term in (31) captures the fact that the market values the capital stock installed in firms. The second term reflects the market value of adopted intermediate goods that are currently used to produce new capital and output. The third term corresponds to the market value of existing intermediate goods which have not yet been adopted. The final term captures the market value of the intermediate goods that will arrive in the future. The rents associated with the arrival of these prototypes also have a value which is priced in by the market.

Of course, only the first term appears in conventional models. It is the last three terms, however, that account for the enhanced volatility of asset prices within our framework. Unlike the first term, the last three are highly pro-cyclical since both current and future profits as well as the flow of current and future technologies increase sharply in booms and decline (relative to trend) in recessions. While the shadow value of a unit of installed capital is procyclical, the replacement cost is countercyclical.

Indeed, the estimates of our model will suggest that overall the value of installed capital is countercyclical on average. Thus it is the terms that reflects the value of current and expected future technologies that ultimately account for the strong procyclical volatility of asset prices within our framework.^{8,9}

6.2 Impulse responses of Stock Market Variables

Figure 8 plots the responses of the stock market and its components to the news shock. The stock market jumps as soon as the news about the future technology hits the economy. In particular, following the same positive news shock that led output to increase initially by about 5% (Figure 4), stock prices increase by about 10 times more. This boom in the stock market occurs despite the fact that the value of installed capital (third panel in first row, Figure 9) declines driven by the decline in the relative price of capital (second panel in first row) which, as in the data moves roughly as much as output (Comin and Gertler, 2006). What drives the stock market boom is the expectation of higher profits from selling intermediate goods in both the near term and over the long run.

The output and investment booms drive up the demand for intermediate goods. The persistence of the output and investment responses to the shock induces higher profits per adopted intermediate good not only upon impact but also in the future. Furthermore, the growth rate of the number of adopted intermediate goods also increases. This is the case for two reasons. First, adoption intensity jumps in response to the increase in the market value of an adopted intermediate good. As a result, unadopted intermediate goods become usable in production more quickly. Second, with the innovation shock, the rate at which unadopted intermediate goods arrive in the economy increases. Hence, the number of intermediate goods that can potentially

⁸Quantitatively, the most important terms to explain the evolution of the stock market are the value of adopted technologies.

⁹Greenwood and Jovanovic (1999) and Hobbijn and Jovanovic (2001) argue that the decline in stock market value during the 70s was driven by the arrival of new technologies that led to a decline in the market value of incumbent companies that were going to become uncompetitive in the new technological era. Unlike the 70s, the innovations that arrived in the 90s and 2000s made incumbent companies more productive. Indeed, many of the applications of the new technologies were developed by incumbent companies (e.g. Internet Explorer, the iphone or the ipod).

be adopted also increases. Though, the arrival of these new technologies does not affect output immediately, it is immediately reflected in the stock market, Q_t . Figure 8 illustrates this phenomenon: There are sharp immediate increases in the value of: adopted technologies (first panel in second row); existing technologies that have not been adopted (second window in second row); and the technologies that have not arrived in the economy yet (third panel in second row).

There are other interesting observations from Figure 9. First, the response of the stock market to the shock is persistent. This is the case because of the persistence in the responses of output, investment and in the number of current and future intermediate goods.¹⁰ Second, the stock market leads output. Intuitively, this is the case because the stock market value at t incorporates the value of future profits which strongly co-move with future output. The response of output, instead is hump-shaped as a result of the frictions that impede a full adjustment in response to the shock. As we show below, the lead of the stock market over GDP is a salient feature of the data.

Our model also has implications for the evolution of the price-dividend ratio. The natural definition of dividends from (31) is capital rental income plus profits from the sale of intermediate goods minus adoption expenses.¹¹ We find that the price-dividend ratio is mean reverting (Figure 9, first panel third row). Intuitively, this is the case because the market's response to the shock declines after the initial impact. In contrast, the slower response of output leads to a more persistent evolution of the profits of intermediate goods producers which are a key component of dividends. As a result, the price-dividend ratio is mean reverting, which is consistent with the evidence in the literature.¹²

So far we have focused on the responses of the stock market to a positive news shock. However, the market responds very similarly to all the other shocks we have considered in the estimation. Consequently, all the findings uncovered for the innovation shock also hold for these other shocks. To save space, we just report the

¹⁰Of course, the persistence of the shock also contributes towards the persistence of Q_t . However, a significant share of the persistence in Q_t is endogenous to the model as will be more clear from the impulse responses to the price of capital and TFP shocks which have significantly less persistence than the news shock.

¹¹Note that the profits for final output and capital producers are equal to the entry costs.

¹²See for example, Campbell and Shiller (1989).

responses to the shocks that were most important in the variance decomposition: the shock to the price of structures and the TFP shock. The market responses to these shocks are reproduced in Figures 9.

The last row in Tables 6 and 7 report the importance of each shock for the evolution of the stock market value both in first differences (Table 6) and HP filtered (Table 7). The most significant shock when using first differences is the TFP shock which explains 84% of the variance with the innovation shock in a distant second (with 15%). For the HP-filtered stock market, the importance of the TFP shock declines to 42% and the contribution of the innovation shock increases to 52%.

Note that in our model stock prices lead movements in TFP. This is also true for movements in stocks prices that are orthogonal to TFP, which is consistent with the evidence in Beaudry and Portier (2004). In particular, within our model the innovation shock does not affect current measured TFP nearly as much as it affects it in the future. Stock prices, instead, rise immediately. (Compare Figures 4 and 8.) It is also the case that other shocks generate this pattern. For example, a shock to structures also influences expected future productivity due to the endogenous diffusion mechanism. Again, stock prices increase immediately, consistent with the BP finding. (Compare Figures 5 and first row of Figure 9).

6.3 Unconditional moments: Model vs. Data

How well does the model fare in matching the stock market in the data? To answer this question, we first compare some basic unconditional moments in the model and in the data (Table 8). Specifically, we simulate 1000 runs of the estimated model each 98 quarters long and compute the volatility and first order autocorrelation of the first differences and HP filtered levels of the stock market and dividends. Then we compare these moments with various data counterparts. For the stock market, we use both the market value of all stocks traded in the US markets and the S&P500 both deflated by the GDP deflator.¹³ It is harder to find the right data counterpart to the dividends in our model. We report two different variables. The dividends

¹³When computing the market value of publicly-traded companies we do not consider the market value of corporate debt due to lack of data on this component of the value of companies.

distributed by publicly traded companies¹⁴ and the compensation to capital from the NIPA tables, both seasonally adjusted.¹⁵

The first finding is that, the volatility of the stock market in the model is approximately two thirds of the volatility in the data. That is true both when comparing the model with the market value and with the S&P500. For example, the average standard deviation of stock market growth in the model is 5.2% while in the data it is 7.7% both for the growth of the market value and of the S&P500.

This gap in the volatility between the model and the data is almost reassuring since our model abstracts from countercyclical risk premia which many authors have stressed is an important component of high frequency fluctuations in the stock market. In particular, Campbell and Shiller (1989) show that revisions in expectations about future dividend growth from simple VAR models cannot account for the observed variation in price-dividend ratios. On the other hand, our model suggests that the contribution of cyclical movements in profits to overall stock market volatility is surely greater than what much of the literature has suggested.

Interestingly, our model is consistent with the Campbell-Shiller tests. Specifically, when conducting a Campbell-Shiller test on data simulated from our model we also find that revisions in expected future dividend growth, when expected future dividends are computed using the simple VARs in CS, only account for a small fraction of the fluctuations in price-dividend ratios of the simulated series.¹⁶ Since in our model none of the fluctuations in the price-dividend ratio are driven by fluctuations in risk premia, this shows that the CS test surely underpredicts the contribution of expected dividend growth to asset price fluctuations. In other words, in our model, and surely in the world too, the dynamics of dividends are rather complex. The simple VARs used by CS cannot properly capture this complexity and, as result, the expected dividend growth series from the VAR forecast are much less volatile than if a more sophisticated model of the economy was used. (Below we comment on what features of the dividend growth process are not captured by the VARs.)

¹⁴Specifically, we follow Campbell and Shiller (1989) and compute the dividends from the value weighted returns including and excluding distributions from COMPUSTAT.

¹⁵That is income minus compensation to employees.

¹⁶Specifically, using the simple one lag 2-variable VAR in Campbell and Shiller (1988) in 1000 (98 quarters-long) simulations, the ratio of predicted over actual standard deviation of the (log) price-dividend ratio is 0.24 with a 95 confidence interval of (0.11, 0.46).

Our model does not perform as well in reproducing the volatility of dividends. In the model, the average standard deviation of dividend growth is 1.27% while the data counterparts are much more volatile (8.7% for the dividends of publicly traded companies). This difference is in part due to the gap between the model and data definition of dividends. In particular, the model measure includes rental income to capital while the data does not. The NIPA measure of dividends includes rental payments to capital and its volatility (2.1%) is closer to the model.

Table 8 also reports the first order autocorrelation of the stock market variables. We find that the average persistence of both the stock market and dividends in the model simulations are very similar to the data both in growth rates and HP-filtered.

These findings raise the question of how a model such as ours, which does not incorporate time-varying risk, can generate large fluctuations in the price-earnings ratio without overly volatile or persistent earnings growth. The answer to this question is that our process of endogenous slow adoption of technologies induces a process for earnings growth that has a small but highly persistent component. This component generates low frequency fluctuations in the capital share and in earnings growth. We illustrate this in Table 8 where we report the volatility of medium term fluctuations in earnings growth. These fluctuations correspond to cycles with periods of length between 8 and 50 years. The main observation is that both in the data and in the model there are significant medium term fluctuations in earnings/dividends growth. In the data the standard deviation of these fluctuations is slightly higher. For NIPA, the standard deviation is 0.0032. In our model, the mean standard deviation is 0.0015 with a 95% confidence interval of (0.0006, 0.0027). Using COMPUSTAT, the volatility is higher (around, 0.01). Further, in the historical series generated from our model, earnings growth is quite highly correlated with the actual data over the medium term. The correlation with NIPA is 0.72 and with COMPUSTAT is 0.61.

One important driver of these low frequency fluctuations in earnings growth in our model is the low frequency variation in the capital share. Our model is also able to generate variation in the capital share consistent with the data (see Table 8). Specifically, the standard deviation of the medium term fluctuations in the log capital share in the US is 0.018. In our simulations the average standard deviation is 0.025 with a 95% confidence interval of (0.0096, 0.03).

Fluctuations in the capital share are important to match the inability of current

price dividend ratios to forecast future dividend growth. In particular, Beeler and Campbell (2008) show that a drawback of the long-run predictability models of Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2007) is that in these models the price-dividend ratio overpredicts consumption growth fluctuations. The endogenous fluctuations in the capital share reflect the wedge that exists in our model between consumption and income. While, in our model, output may increase significantly in response to expansionary shocks, consumption does not increase as much since agents find that the shocks also increase the return to other activities such as adopting new technologies, increasing the capital stock or entering in the production of final capital goods. As a result, current price-dividend ratios do not predict so well consumption growth over the short and medium term.

This is illustrated in Table 9. The first column reports the effect of the price-dividend ratio on cumulative consumption growth in the data over three horizons (4, 12 and 20 quarters). The second column reports the same coefficient using the historical evolution of the price-dividend ratio as predicted by our model. The difference between the estimates in these two columns are not only statistically insignificant but very close to zero. In the third column, we report the same coefficients when using the 1000 simulations from our model both for consumption growth and for the price-dividend ratio. Now the average point estimates are quite far from the point estimates in the first two columns but these are well inside the 95% confidence interval. Based on this we conclude that our model does a fair job in reproducing the long-run predictability tests.

Another difficulty encountered by many asset pricing models is the difficulty of explaining simultaneously the volatility of stock prices and the risk-free rate. When they explain the former they tend to generate a risk-free rate that is too volatile. This is not the case in our model. The standard deviation of the real interest rate in the quarterly data is 0.0056 while the average standard deviation in our 1000 simulations is 0.0061 with a 95% confidence interval of (0.0053, 0.007). Intuitively, our estimated model has nominal price rigidities and short term nominal rates set by a Taylor rule. Thus, the observed market real rate is not overly volatile. What is volatile is the unobserved “natural rate of interest”, i.e. the real rate that would arise if prices were perfectly flexible.

Table 8 also reports the moments for the stock market series generated from a

model with a conventional real and monetary sector similar to Justiniano, Primiceri, and Tambalotti (2008). Overall, this model fails to account for the volatility of stock prices. In particular, while the average volatility of stock market growth in our model is 5.2% and of the HP-filtered stock market value is 6.3%, the equivalent statistics from this alternative model are both 2%. Hence, the more conventional model is unable to generate the observed large fluctuation in asset prices.

In addition to the variance and autocorrelation, another important feature of the stock market in the data is that it leads output, unconditionally. This is illustrated in Figure 10 which plots the cross-correlogram of HP-filtered output and the stock market value in the data. Overall, the model captures the lead in the stock market. Specifically, it plots the average cross-correlogram of output and the stock market in the 1000 runs of our model together with the 95% confidence interval. As in the data, the stock market in the model strongly co-moves contemporaneously with output. Further, there is a lead of about one quarter of the stock market over output which is also consistent with the data.

The pattern of co-movement of the stock market and output is another dimension where our model differs from the conventional framework. Figure 10 also plots the average cross-correlogram between output and the stock market for this model. Two observations are worth making. First, the contemporaneous co-movement between output and the stock market is negative rather than positive. This is driven by the shocks to the relative price of capital which, as in Justiniano, Primiceri, and Tambalotti (2008) are an important source of fluctuations when this model is estimated. A shock that reduces the price of capital, causes an output expansion but, despite the presence of adjustment costs, a reduction in the price of installed capital. Since capital is fixed in the short run, this shock causes a decline in the value of the capital stock which is the stock market in this model. Second, the co-movement pattern between output and the stock market in this model does not capture the observed lead of the stock market over output.

6.4 Historical evolution of the stock market

How closely does the stock market value predicted by the model given the estimated shocks track the actual evolution of the US stock market? Figure 11 plots the evo-

lution of the predicted and actual (real) value of the stock market together with the S&P500 deflated using the GDP deflator. The stock market value in the data is the value of all publicly-traded companies deflated also by the GDP deflator.

The finding is that, to a first order, the predicted stock market value tracks fairly closely the actual series. In particular, the model captures the relatively slow growth between 1984 and 1994, the acceleration starting in 1994-95.¹⁷ The peak takes place in 2001 rather than in 2000. Then there is a small decline though not nearly as pronounced as in the 2001 crash. The model also captures the recovery until the end of 2007. Finally, it captures the decline in the stock market in 2008.

Beyond the qualitative patterns, the model does a surprisingly good job in capturing the magnitude of the run up during the second half of the 90s. While the (real) US stock market went from a value of \$2.59 trillion in 1984:I¹⁸ to \$18.28 trillion in 2000:I, our model predicts an increase from \$2.59 trillion to \$16.68 trillion in 2001:I. The similarity of these increases is somewhat surprising, given that we have not used any information from the stock market to estimate the model.

The predictions of the model for the evolution of the stock market in 2008 are also worth noting. In particular, the model predicts a decline in the stock market value of 18% which is approximately half of the decline that experienced the S&P500. It is important to stress, though, that our model abstracts from financial factors that appear to be relevant in the sharp decline in stock prices since October 2008. Further, the data used in the estimation of the model and identification of the shocks runs only until the second quarter of 2008. It is interesting though that the macroeconomic conditions identified in the estimation were sufficient to generate such a significant drop in asset prices in the context of our model.

7 Conclusions

We have modified a conventional business cycle model to allow for changes in the rate of growth of new technologies and endogenous technology diffusion. An "innovation" shock has the flavor of a news shocks because it influences expectations of

¹⁷The most important component in (31) to account for the upward trend during the second half of the 90s is the value of adopted innovations.

¹⁸All these figures are in 2000 US dollars.

future growth without affecting current productivity. As we, show, with endogenous diffusion, news about future growth prospects produces movements in current output and hours that is positively correlated with the news. In this way the paper addresses a conundrum in the literature, originally identified by Cochrane (1994). We also find that in an estimated version of the model, the innovation shock accounts for nearly a third of the variation of output fluctuations, and even more at the business cycle frequencies. The model also accounts surprisingly well for asset price movements, at least relative to most other business cycle models.

Our endogenous technology diffusion mechanism is also relevant to other disturbances besides innovation shocks. For example, the mechanism amplifies and propagates the impact of a shock to structures on the movement of both output and asset prices. As we noted, our structures shock, which affects both residential and non-residential investment may in a reduced form sense partly capture movements in credit frictions. Indeed, our historical decomposition suggests that this structures shock was important in both the 1990-91 recession and the period leading up to the current recession, episodes where disruptions in credit markets appear to have affected structures investment. Even though the initiating disturbance does not involve technology, the endogenous diffusion mechanism works to propagate the effects of the shock on output and the stock market. Explicitly modeling the interactions between credit marker frictions and our endogenous diffusion mechanism, we think, is an important next step to take.

A Appendix

A.1 Symmetric Equilibrium

This section describes the complete set of equations that determine the symmetric equilibrium.¹⁹

A symmetric equilibrium in this economy is defined as an exogenous stochastic sequence, $\{X_t, G_t, P_{st}^k, \xi_t\}_{t=0}^{\infty}$, an initial vector $\{A_{n0}, Z_{n0}, K_{n0}\}$, a sequence of parameters, a sequence of prices $\{P_t^k, \bar{P}_t^k, P_{et}^k\}_{t=0}^{\infty}$, endogenous variables, $\{Y_t, C_t, \Lambda_{rt+1}, I_t, I_t^e, I_t^s, J_t, U_t, L_t, h_t^s, N_t^y, N_t^k, v_t^s, j_t^s, \pi_t^k, \pi_t^y, \lambda_t^s\}$ for $s = \{k, y\}$, and laws of motion $\{A_{t+1}^s, Z_{t+1}^s, K_{t+1}^s\}_{t=0}^{\infty}$ such that,

- The state variables $\{A_{t+1}^s, Z_{t+1}^s, K_{t+1}^s\}_{t=0}^{\infty}$ satisfy the laws of motions in equations (34) to (36)
- The endogenous variables solve the producers and consumers problems in equations (37) to (54)
- Feasibility is satisfied in equations (32) and (33)
- Prices are such that the markets clear

The equilibrium relations of this economy are:

Resource Constraint:

$$Y_t = C_t + G_t + \frac{P_t^k J_t}{\bar{\mu}^k} + \overbrace{\frac{\mu - 1}{\mu} Y_t + \frac{\mu^k - 1}{\mu^k} I_t}^{\text{Entry Costs}} + \sum_{s=\{k,y\}} \overbrace{(Z_t^s - A_t^s) h_t^s}^{\text{Adoption Costs}} \quad (32)$$

Aggregate production:

$$Y_t = X_t (A_t^y)^{\vartheta-1} (N_t^y)^{\mu-1} (U_t K_t)^{\alpha} L_t^{1-\alpha} \quad (33)$$

where total factor productivity, $X_t (A_t^y)^{\vartheta-1} (N_t^y)^{\mu-1}$, depends on the stock of adopted intermediate output goods A_t^y .

Evolution of endogenous states, K_t and A_t^y and A_t^k :

¹⁹Note that there are 23 equations and 23 variables.

$$K_{t+1} = (1 - \delta(U_t))K_t + (P_t^K)^{-1}\bar{\mu}^k I_t, \quad (34)$$

where $\bar{\mu}^k \equiv \frac{\mu^k \theta}{\theta\gamma + (1-\gamma)}$ is the average markup in the production of new capital.

$$\bar{\mu}^k \equiv \frac{\mu^k \theta}{\theta\gamma + (1-\gamma)}$$

$$A_{t+1}^s = \lambda_t^s [Z_t^s - A_t^s] + \phi A_t^s, \text{ for } s = \{k, y\}. \quad (35)$$

and where the evolution of the stock of new technologies in each sector, is

$$Z_{t+1}^s = (\bar{\chi}_s \chi_t^{\xi_s} + \phi) Z_t^s \quad (36)$$

Factor market equilibria for L_t , and U_t :

$$(1 - \alpha) \frac{Y_t}{L_t} = \mu \mu^w L_t^\zeta / (1/C_t) \quad (37)$$

$$\alpha \frac{Y_t}{U_t} = \mu \delta'(U_t) P_t^K K_t \quad (38)$$

New Capital:

Let I_t^e denote the amount output devoted to producing equipment and I_t^s denote the amount devoted to structures. Then the optimal pricing of equipment, and structures capital goods and final capital goods implies that

$$\frac{P_t^k J_t}{\mu^k} = \theta I_t^e + I_t^s \quad (39)$$

where from cost minimization:

$$\frac{\theta I_t^e}{I_t^s} = \frac{1 - \gamma}{\gamma} \quad (40)$$

and

$$J_t = (P_t^k)^{-1} \bar{\mu}^k I_t \quad (41)$$

Consumption/Saving:

$$E_t \left\{ \beta \Lambda_{t,t+1} \cdot \left[\alpha \frac{Y_{t+1}}{\mu K_{t+1}} + (1 - \delta(U_{t+1})) P_{t+1}^K / P_t^k \right] \right\} = 1 \quad (42)$$

where

$$\Lambda_{rt+1} = C_t/C_{t+1} \quad (43)$$

Optimal adoption of innovations in sector $s = \{k, y\}$:

$$1 = \phi\beta E_t \left[\Lambda_{t+1} \frac{A_t^s}{o_t^s} \lambda' \left(\frac{A_t^s}{o_t^s} h_t^s \right) (v_{t+1}^s - j_{t+1}^s) \right] \quad (44)$$

with

$$v_t^s = \pi_t^s + \phi\beta E_t [\Lambda_{t+1} v_{t+1}^s] \quad (45)$$

and

$$\pi_t^k = \left(1 - \frac{1}{\theta}\right) (1 - \gamma) \frac{I_t}{A_t^k \mu_k} \quad (46)$$

$$\pi_t^y = \left(1 - \frac{1}{\vartheta}\right) \frac{Y_t}{A_t^y \mu} \quad (47)$$

$$j_t^s = -h_t^s + \phi\beta E_t [\Lambda_{t+1} [\lambda_t^s v_{t+1}^s + (1 - \lambda_t^s) j_{t+1}^s]] \quad (48)$$

where

$$\lambda_t^s = \bar{\lambda}^s \left(\frac{A_t^s h_t^s}{o_t^s} \right)^{\rho_\lambda} \quad (49)$$

Free entry into production of final goods and final capital goods:

$$\frac{\mu - 1}{\mu} \frac{Y_t}{N_t^y} = o_t^y \quad (50)$$

$$\frac{\mu^k - 1}{\mu^k} \frac{I_t}{N_t^k} = o_t^k \quad (51)$$

Relative price of retail and wholesale capital

$$P_t^K = \mu^k (N_t^k)^{-(\mu^k - 1)} (P_{st}^K)^\gamma (P_{et}^K)^{1-\gamma} \quad (52)$$

where P_{et}^K is equal to

$$P_{et}^K = \theta (A_t^k)^{-(\theta - 1)} \quad (53)$$

and the wholesale price of capital is

$$\bar{P}_t^K = \theta^{(1-\gamma)} (A_t^k)^{-(1-\gamma)(\theta-1)} (P_{st}^K)^\gamma \quad (54)$$

The exogenous variables, $\{X_t, G_t, P_{st}^k, \xi_t\}_{t=0}^\infty$, follow an AR(1) process.

A.2 Calibration

We begin with the standard parameters. A period in our model corresponds to a quarter. We set the discount factor β equal to 0.98, to match the steady state share of investment to output. Based on steady state evidence we also choose the following numbers: (the capital share) $\alpha = 0.35$; (the equipment share) $(1 - \gamma) = 0.17/0.35$; (government consumption to output) $G/Y = 0.2$; (the depreciation rate) $\delta = 0.015$; and (the steady state utilization rate) $U = 0.8$.²⁰ We set the inverse of the Frisch elasticity of labor supply ζ to unity, which represents an intermediate value for the range of estimates across the micro and macro literature. Similarly, we set the elasticity of the change in the depreciation rate with respect to the utilization rate, $(\delta''/\delta')U$ at 0.15 following Rebelo and Jaimovich (2006). Finally, based on evidence in Basu and Fernald (1997), we fix the steady state gross valued added markup in the final output, μ , equal to 1.1 and the corresponding markup for the capital goods sector, μ^k , at 1.15.

We next turn to the “non-standard” parameters. To approximately match the operating profits of publicly traded companies, we set the gross markup charged by intermediate capital (θ) and output goods (ϑ) to 1.4 and 1.25, respectively. Following Caballero and Jaffe (1993), we set ϕ to 0.99, which implies an annual obsolescence rate of 4 percent. The steady state growth rate of the relative price of capital, depends on $\bar{\chi}_k$, the markup θ , the obsolescence rate and ξ_k . We normalize ξ_k to 1. To match the average annual growth rate of the Gordon quality adjusted price of equipment relative to the BEA price of consumption goods and services (-0.035), we set $\bar{\chi}_k$ to 3.04 percent.

The growth rate of GDP in steady state depends on the growth rate of capital and on the growth rate of intermediate goods in the output sector. To match the average annual growth rate of GDP per working age person over the postwar period (0.024) we set $\bar{\chi}_y$ to 2.02 percent.

For the time being, we also need to calibrate the autocorrelation of the shock to future technologies. When we estimate the model, this will be one of the parameters we identify. One very crude proxy of the number of prototypes that arrive in the

²⁰We set U equal to 0.8 based on the average capacity utilization level in the postwar period as measured by the Board of Governors.

economy is the number of patent applications. The autocorrelation of the annual growth rate in the stock of patent applications is 0.95. This value is consistent with the estimate we obtain below and is the value we use to calibrate the autocorrelation of χ_t .

We now consider the parameters that govern the adoption process. We use two parameters to parameterize the function $\lambda^s(\cdot)$ as follows:

$$\lambda_t^s = \bar{\lambda}^s \left(\frac{A_t^s h_t^s}{o_t^s} \right)^{\rho_\lambda}$$

These are $\bar{\lambda}^s$ and ρ_λ . To calibrate these parameters we try to assess the average adoption lag and the elasticity of adoption with respect to adoption investments. Estimating this elasticity is difficult because we do not have good measures of adoption expenditures, let alone adoption rates. One partial measure of adoption expenditures we do have is development costs incurred by manufacturing firms trying that make new capital goods usable, which is a subset of the overall measure of R&D that we used earlier. A simple regression of the rate of decline in the relative price of capital (the relevant measure of the adoption rate of new embodied technologies in the context of our model) on this measure of adoption costs and a constant yields an elasticity of 0.9. Admittedly, this estimate is crude, given that we do not control for other determinants of the changes in the relative price of capital. On the other hand, given the very high pro-cyclicality of the speed of adoption estimated by Comin (2009), we think it provides a plausible benchmark value.

Given the discreteness of time in our model, the average time to adoption for any intermediate good is approximately $1/\lambda + 1/4$. Mansfield (1989) examines a sample of embodied technologies and finds a median time to adoption of 8.2 years. However, there are reasons to believe that this estimate is an upper bound for the average diffusion lag. First, the technologies typically used in these studies are relatively major technologies and their diffusion is likely to be slower than for the average technology. Second, most existing studies oversample older technologies which have diffused slower than earlier technologies.²¹ For these reasons, we set $\bar{\lambda}^s$ to match an average adoption lag of 5 years and a quarter.²²

²¹Comin and Hobjin (2007) and Comin, Hobjin, and Rovito (????).

²²It is important to note that, as shown in Comin (2009), a slower diffusion process increases

We next turn to the entry/exit mechanism. We set the overhead cost parameters so that the number of firms that operate in steady state in both the capital goods and final goods sector is equal to unity, and the total overhead costs in the economy are approximately 10 percent of GDP.

A.3 Data

The vector of observable variables is:

$$[\Delta \log Y_t \quad \Delta \log C_t \quad \Delta \log I_t^e \quad \Delta \log I_t^s \quad R_t \quad \Pi_t \quad \log(L_t)]$$

Following Smets and Wouters (2007), and Primiceri, Schaumburg, and Tambalotti (2006) and Justiniano, Primiceri, and Tambalotti (2008), we construct real GDP by dividing the nominal series (GDP) by population and the GDP Deflator. Real series for consumption and investment in equipment and structures are obtained similarly. Consumption corresponds only to personal consumption expenditures of non-durables and services; while non-equipment investment includes durable consumption, structures, change in inventories and residential investment. Labor is the log of hours of all persons divided by population. The quarterly log difference in the GDP deflator is our measure of inflation, while for nominal interest rates we use the effective Federal Funds rate. Because we allow for non-stationary technology growth, we do not demean or detrend any series.

the amplification of the shocks from the endogenous adoption of technologies because increases the stock of technologies waiting to be adopted in steady state. In this sense, by using a higher speed of technology diffusion than the one estimated by Mansfield (1989) and others we are being conservative in showing the power of our mechanism.

References

- [1] Sungbae An and Frank Schorfheide. Bayesian analysis of dsge models. *Econometric Reviews*, 2007.
- [2] Bansal, Kiku, and Yaron. Risks for the long run: Estimation and inference. *mimeo, Duke University*, 2007.
- [3] Ravi Bansal and Amir Yaron. Risks for the long run. *Journal of Finance*, 2004.
- [4] Susanto Basu and John G. Fernald. Returns to scale in u.s. production: Estimates and implications. *Journal of Political Economy*, 1997.
- [5] Paul Beaudry, Fabrice Collard, and Frank Portier. Gold rush fever in business cycles. *mimeo, University of British Columbia*, 2007.
- [6] Paul Beaudry and Frank Portier. An exploration into pigou’s theory of cycles. *Journal of Monetary Economics*, 2004.
- [7] Beeler and Campbell. The long-run risks model and aggregate asset prices: An empirical assessment. *mimeo, Harvard University*, 2008.
- [8] William H. Beveridge. Unemployment: A problem of industry. *Longmans, Green, and Co.*, 1909.
- [9] Ricardo J. Caballero and Adam B. Jaffe. How high are the giants’ shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth. *NBER Macroeconomics Annual, O. J. Blanchard and S. Fischer editors, MIT press*, 1993.
- [10] John Campbell and Robert J. Shiller. The dividend ratio model and small sample bias : A monte carlo study. *Economic Letters*, 1989.
- [11] Lawrence J. Christiano, Illut Cosmin, Roberto Motto, and Massimo Rostagno. Monetary policy and stock market boom-bust cycles. *mimeo, Northwestern University*, 2007.

- [12] Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 2005.
- [13] John M. Clark. Strategic factors in business cycles. *NBER*, 1934.
- [14] John Cochrane. Shocks. *Carnegie-Rochester Conference Series on Public Policy*, 1994.
- [15] Diego Comin. On the integration of growth and business cycles. *Empirica*, 2009.
- [16] Diego Comin and Mark Gertler. Medium-term business cycles. *American Economic Review*, 2006.
- [17] Diego Comin and Bart Hobjin. Lobbies and technology diffusion. *Review of Economics and Statistics*, 2007.
- [18] Diego Comin, Bart Hobjin, and Emilie Rovito. A new approach to measuring technology with an application to the shape of diffusion curves. *Journal of Technology Transfer*, 2008.
- [19] Jeremy Greenwood, Zvi Hercowitz, and Gregory W. Huffman. Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 1988.
- [20] Jeremy Greenwood, Zvi Hercowitz, and Per Krusell. The role of investment-specific technological change in the business cycle. *European Economic Review*, 2000.
- [21] Jeremy Greenwood and Boyan Jovanovic. The it revolution and the stock market. *AEA Papers and Proceedings*, 1999.
- [22] Robert E. Hall. e-capital: The link between the labor market and the stock market in the 1990s. *Brookings Papers on Economic Activity*, 2000.
- [23] Bart Hobjin and Boyan Jovanovic. The it revolution and the stock market: Evidence. *American Economic Review*, 2001.

- [24] Miguel Angel Iraola and Manuel Santos. Productivity and stock returns. *Work in progress*, 2007.
- [25] Nir Jaimovich and Sergio Rebelo. Can news about the future drive the business cycle? *mimeo, Northwestern University*, 2006.
- [26] Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti. Investment shocks and business cycles. *CEPR discussion paper No. 6739*, 2008.
- [27] Edwin Mansfield. The diffusion of industrial robots in japan and the united states. *Research Policy*, 1989.
- [28] Lubos Pastor and Pietro Veronesi. Technological revolutions and stock prices. *American Economic Review*, 2008, forthcoming.
- [29] Arthur C. Pigou. Industrial fluctuations. *Macmillan and Co.*, 1927.
- [30] Giorgio E. Primiceri, Ernst Schaumburg, and Andrea Tambalotti. Intertemporal disturbances. *NBER working paper No. 12243*, 2006.
- [31] Frank Smets and Raf Wouters. Shocks and frictions in u.s. business cycles: A bayesian dsge approach. *American economic Review*, 2007.

Standard Parameters	Value
β	0.98
δ	0.015
G/Y	0.2
α	0.35
α_s	0.17/0.35
ζ	1
θ	0.7
$\bar{\theta}$	0.8
U	0.8
$(\delta''/\delta')U$	0.15
μ	1.1
μ_w	1.2
μ_k	1.15
Non-standard Parameters	Value
$\bar{\chi}_y$	so that growth rate of $y=0.024/4$
$\bar{\chi}_k$	so that growth rate of $p_{et}^K=-0.035/4$
ϕ	0.99
$\bar{\lambda}^y$	so that $\lambda^y=0.2/4$
$\bar{\lambda}^k$	so that $\lambda^k=0.2/4$
ρ_λ	0.9
ξ_y	0.6

Table 1: Calibrated parameters

Parameter	Prior	Posterior			
	Distribution	max	mean	5%	95%
ν	Beta (0.50,0.10)	0.502	0.565	0.104	0.952
ρ_r	Beta (0.65,0.10)	0.642	0.623	0.518	0.800
ξ	Beta (0.5,0.10)	0.565	0.557	0.366	0.758
ι_p	Beta (0.5,0.10)	0.488	0.487	0.280	0.694
η	Normal (1.00,0.50)	1.305	1.185	0.818	1.510
ϕ_p	Gamma (1.70,0.30)	1.707	1.944	1.226	2.746
ϕ_y	Gamma(0.125,0.10)	0.079	0.082	0.062	0.106
ζ	Gamma (1.20,0.10)	1.193	1.344	1.150	1.516
$\frac{\delta''U}{\delta'}$	Gamma (0.10,0.10)	0.025	0.022	0.003	0.043

Table 2: Prior and Posterior Estimates of Structural Coefficients

Coefficient	Prior	Posterior			
	Distribution	max	mean	5%	95%
ρ^b	Beta (0.25,0.05)	0.235	0.230	0.185	0.284
ρ^m	Beta (0.25,0.05)	0.248	0.247	0.186	0.301
ρ^w	Beta (0.35,0.10)	0.346	0.349	0.331	0.364
ρ	Beta (0.95,0.15)	1.000	0.999	0.999	0.999
ρ^g	Beta (0.6,0.15)	0.349	0.894	0.893	0.894
ρ^{st}	Beta (0.95,0.15)	1.000	0.999	0.999	0.999
σ_{rd}	IGamma(0.25, ∞)	0.285	0.292	0.255	0.337
σ_w	IGamma (0.25, ∞)	0.254	0.263	0.254	0.272
σ_g	IGamma (0.25, ∞)	0.252	0.267	0.248	0.287
σ_b	IGamma (0.25, ∞)	0.252	0.261	0.227	0.296
σ_m	IGamma (0.25, ∞)	0.251	0.268	0.191	0.352
σ_x	IGamma (0.25, ∞)	0.253	0.277	0.269	0.287
σ_s	IGamma (0.25, ∞)	0.306	0.206	0.164	0.245

Table 3: Prior and Posterior Estimates of Shock Processes

Observable	Data	Endogenous	Exogenous	Benchmark
ΔY_t	0.50	0.63	0.78	1.18
ΔI_t^e	2.92	2.91	2.24	1.40
ΔI_t^s	2.80	2.77	2.18	2.00
ΔC_t	0.33	0.43	0.31	0.40
ΔL_t	0.66	0.60	0.30	0.66

Table 4: Standard deviations in data and alternative models

Specification	Log Marginal
Benchmark	1906
Exogenous Adoption	2092
Endogenous Adoption	2337

Table 5: Log-Marginal Density Comparison

Observ.	Gov.	Lab.Sup.	Int.Pref.	Innov.	Neutr.Tech.	Inves.	Mon.Pol.
ΔY_t	3.45	0.38	9.94	27.15	42.57	10.62	5.89
ΔI_t^e	0.07	0.08	0.74	49.36	35.15	13.67	0.93
ΔI_t^s	0.08	0.09	0.83	33.53	42.05	22.13	1.29
ΔC_t	0.16	1.70	19.38	18.05	40.03	9.43	11.25
ΔL_t	1.61	32.34	0.99	13.69	49.04	1.64	0.69
ΔQ_t	0.27	0.59	0.01	14.83	84.14	0.16	0.00

Table 6: Variance Decomposition

Observ.	Gov.	Lab.Sup.	Int.Pref.	Innov.	Neutr.Tech.	Inves.	Mon.Pol.
Y_t	1.45	0.21	3.84	32.29	34.24	24.78	3.19
I_t^e	0.07	0.06	0.62	35.52	38.00	24.03	1.71
I_t^s	0.08	0.07	0.72	36.92	39.93	20.64	1.65
C_t	0.31	3.61	16.91	15.93	25.60	24.31	13.33
L_t	2.09	35.87	0.75	20.06	29.16	11.24	0.84
Q_t	4.10	1.85	0.11	51.83	41.68	0.18	0.26

Table 7: Variance Decomposition (HP Filtered)

	Volatility			Autocorrelation		
	Data ^a	Our Model	Conven. Model	Data	Our Model	Conven. Model
Growth rate stock market value	0.077	0.052 (0.045, 0.059)	0.021 (0.018, 0.024)	-0.04 (-0.25, 0.17)	0 (-0.2, 0.19)	-0.18 (-0.35, -0.01)
Growth rate S&P500	0.077			-0.03 (-0.24, 0.17)		
HP-filtered stock market value	0.103	0.063 (0.049, 0.079)	0.02 (0.016,0.023)	0.71 (0.53, 0.88)	0.67 (0.49,0.8)	0.45 (0.27, 0.6)
HP-filtered S&P500	0.107			0.76 (0.61, 0.91)		
Dividend growth (COMPUSTAT), s.a ^b	0.087	0.0127 (0.0107, 0.014)	0.014 (0.012, 0.016)	-0.56 (-0.83, -0.29)	-0.36 (-0.51, -0.2)	-0.25 (-0.43, -0.06)
Profit growth (NIPA)	0.022			-0.24 (-0.67, 0.18)		
HP-filtered dividends (COMPUSTAT), s.a	0.072	0.0106 (0.009, 0.0127)	0.0134 (0.011, 0.016)	0.29 (0.06, 0.52)	0.3 (0.04, 0.49)	0.46 (0.25, 0.64)
HP-filtered profits (NIPA)	0.022			0.53 (0.28, 0.82)		
Medium term ^c dividend growth (COMPUSTAT), s.a	0.011	0.0015 (0.0006, 0.0027)	0.001 (0.0005,0.002)	0.99 (0.97,1)	0.99 (0.99,1)	0.99 (0.99,1)
Medium term profit growth (NIPA), s.a	0.0031			0.99 (0.97,1)		
(Log) capital share	0.025	0.041 (0.019, 0.082)	0.03 (0.027,0.037)	0.83 (0.7,0.96)	0.93 (0.81,0.99)	0.39 (0.18,0.58)
Medium term (log) capital share	0.0186	0.03 (0.0096, 0.063)	0.01 (0.027,0.037)	0.99 (0.97,1)	0.99 (0.99,1)	0.99 (0.99,1)

Table 8: Volatility of Stock Market variables

^aIn the stock market data, the period is 1984:I to 2008:II

^bSeasonally Adjusted

^cMedium term variables are computed by applying Band Pass filter that isolates fluctuations with periods between 8 and 50 years

Horizon (in quarters)	Data ^a	Model Historical series	Model simulated series
4	0.001 (-0.0087, 0.0107)	-0.0025 (-0.008, 0.0029)	-0.0028 (-0.0288, 0.0154)
12	0.0034 (-0.0174, 0.024)	-0.0037 (-0.015, 0.007)	-0.0484 (-0.1352, 0.0352)
20	0.0031 (-0.0194, 0.025)	-0.004 (-0.02, 0.012)	-0.0985 (-0.2094, 0.0351)

^aCoefficient reported is β from the following regression: $\sum_{\tau=1}^T \Delta c_{t+\tau} = \alpha + \beta x_t + \varepsilon_t$, where x_t is the price-dividend ratio and T is the horizon.

Table 9: Long-run predictability of consumption growth




Figure 1: Impulse responses to an innovation shock in conventional model (immediate diffusion)

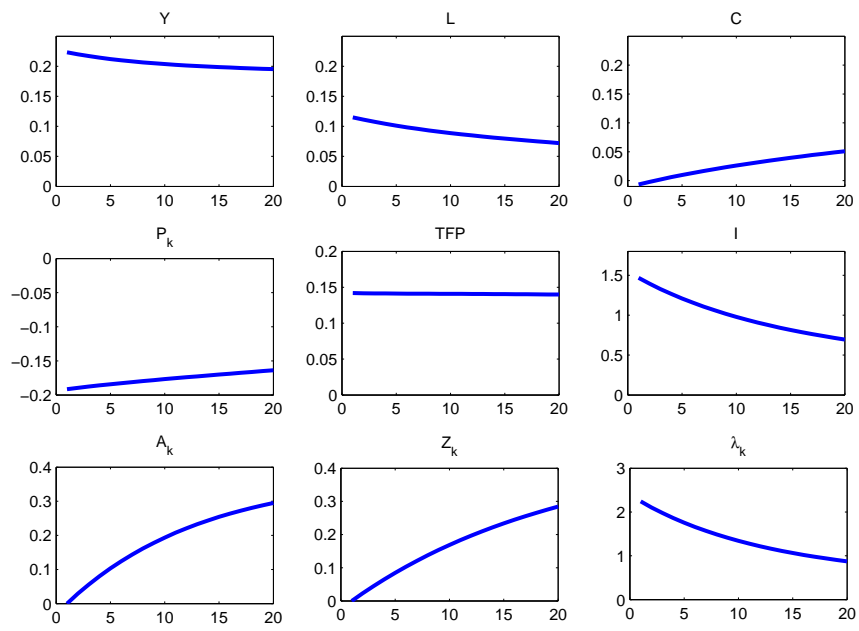


Figure 2: Impulse responses to an innovation shock in baseline model (slow diffusion, endogenous adoption)




Figure 3: Robustness: Impulse responses to innovation shock. Top row: baseline model (slow diffusion, endogenous adoption). Middle row: baseline model without entry. Bottom row: baseline model without endogenous adoption.




Figure 4: Estimated impulse responses to innovation shock, our model (solid) and model with entry and exogenous adoption (dashed).




Figure 5: Estimated impulse responses to structures shock, our model (solid) and model with entry and exogenous adoption (dashed).




Figure 6: Estimated impulse responses to TFP shock, our model (solid) and model with entry and exogenous adoption (dashed).




Figure 7: Historical decomposition of output growth. Data in dotted green and counterfactual in solid blue, for innovation shock (first panel), structures shock (second panel), and TFP shock (third panel).




Figure 8: Impulse responses to innovation shock for stock market value and its components: installed capital (first row, third column), adopted technologies (second row, first column), unadopted technologies (second row, second column), and future unadopted technologies (second row, third column).




Figure 9: Impulse responses of stock market variables to positive shock to structures (first row) and TFP (second row).

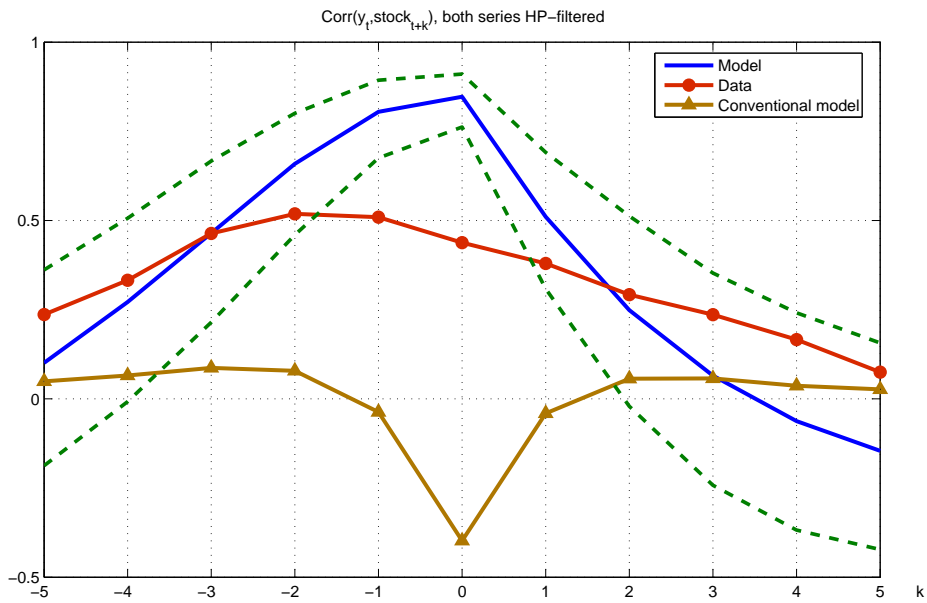


Figure 10: $Corr(y_t, stock_{t+k})$ in the data (first panel), our model (second panel), and conventional model (third panel).




Figure 11: Stock market value in model (solid blue), data (dotted green) and S&P500 (triangled red, right axis).