Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 54 (2007) 2486–2511
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Optimal learning and new technology bubbles ☆

Timothy C. Johnson*

London Business School, Regent's Park, London NW1 4SA, UK

Received 17 July 2005; received in revised form 5 March 2007; accepted 7 March 2007
Available online 6 April 2007

Abstract

It is widely believed that there is a fundamental linkage between major technological innovations, speculative fever, and wasteful overinvestment. This paper presents an equilibrium model of investment in a new industry, whose return-to-scale is not known in advance. Overinvestment relative to the full-information case is then optimal as the most efficient way to learn about the new technology. Moreover, the initial overinvestment is accompanied by apparently inflated stock prices and apparently negative expected excess returns in the new industry, which are also fully rational. This suggests a new interpretation of what seems to be stock market driven real bubbles.

JEL classification: G12; D8

Keywords: Adaptive control; Investment; New technology; Bubbles; Learning

> *Every previous technological revolution has created a speculative bubble, and there is no reason why IT should be different.*
>
> Economist, 2000

1. Introduction

It has recently become widely believed that technological breakthroughs inevitably entail economic excess. The pattern of the recent IT-driven boom and bust has led many

☆I am indebted to Viral Acharya, Suleyman Basak, Wouter den Haan, James Dow, Francisco Gomes, Denis Gromb, Robert King, Anthony Neuberger, Anna Pavlova, and seminar participants at LBS and LSE for thoughtful comments and advice.
*Fax: +44 207 724 3317.
E-mail address: tjohnson@london.edu.

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2007.03.004

commentators and historians to note parallels with earlier technological revolutions—railroads, canals, electric power—which ignited a burst of apparent overbuilding by, and apparent overvaluation of, innovating firms.[1] These rapid expansions were all followed by longer adjustment phases during which the initial excesses were damped toward long-run equilibrium values. The description does not fit all waves of innovation or all financial bubbles. But the recurrence of the pattern does raise the question of what it is about the technological innovations that induced these dynamics.

From a business cycle perspective, such episodes are unusual for a number of reasons. In general, overinvestment in response to a technology shock, large or small, is not a feature of standard models. Likewise, the subsequent disinvestment—without any technological regress—is hard to explain as an optimizing policy. Moreover, these responses reverse the normal asymmetry in which sharp recessions are followed by gradual expansions. Similarly reversed is the usual pattern of investment seeming to respond too strongly to cash-flow and not strongly enough to Tobin's q. From an asset pricing perspective, any pattern of apparently predictable negative returns is also very difficult to explain.

This paper suggests one mechanism that can account for these facts. It studies the short-run equilibrium dynamics following the introduction of a new production technology in a standard equilibrium setting. In this context, I show that, when the return-to-scale of the new technology are not known *a priori*, optimal policies can feature both initial overshooting of real investment and predictable deflation in the price of claims to the new sector. This behavior is driven by the incentive to efficiently learn the curvature of the production function—and hence the optimal long-run scale of the new industry—about which agents are uncertain. Indeed, this particular type of uncertainty could be said to be the distinguishing feature of a truly revolutionary technology: there is no historical experience of how it will scale up. No one knows how the interplay of competition, regulation, and costs will work out at vastly greater levels of production than have ever been seen before.

In general, adaptive learning models can induce either caution or experimentation. In my formulation, agents have an incentive to push investment beyond the level that would seem optimal with full information in order to efficiently learn the shape of the production function. As experience grows, this incentive diminishes and investment declines. Market prices of installed capital mirror the gains to be had from learning. Tobin's q for the new industry starts high and then predictably subsides. The model is not intended as a general theory of booms and busts. Nor does it attempt to model either the evolution of the new technology or the process of its adoption.[2] Instead, the goal is to focus on the apparent overshooting of both real and financial quantities that seems to have characterized several important historical periods.

Given the enormous literature on financial bubbles and the still larger one on technology-driven business cycles, it is not surprising that alternative explanations for real bubbles already exist.

[1]See the surveys of Perez (2003) and Bordo (2003).

[2]Models of endogenous growth and "learning-by-doing" also incorporate learning into the optimizing decisions of agents (see Jovanovic, 1997). In contrast to the adaptive control approach, however, there learning is modeled as the accumulated output of a production function for "knowledge", not as the Bayesian evolution of a probability distribution.

The simplest explanation is just error. Growth rate expectations drive levels of prices and investment, and expectations can be wrong. For important innovations, growth rates themselves are large numbers. Hence small errors can have big consequences. Here two distinct perspectives can be taken. On one hand, the episodes that stand out over the course of history may simply be the most visible instances of what is essentially idiosyncratic error. In this view, overestimated growth rates lead to spectacular bubbles, while underestimated ones, equally often, lead to unremarkable gradual adjustments. On the other hand, the errors could be systematic. This stance has been forcefully argued by Shiller (2000), who links historical "new economy" sentiment to persistent cognitive biases generating irrational exuberance. The behavioral finance literature has built an impressive body of evidence documenting biases in expectations and systematic negative returns to high-growth, new, and low book-to-market stocks. Some recent empirical work (Polk and Sapienza, 2002; Gilchrist et al., 2002) supports the notion of a behavioral link to real investment.

It is worth noting that the behavioral argument does not *per se* require a role for the stock market. Irrationally exuberant agents would presumably overinvest however allocations were implemented. That said, undoubtedly the most widely held view of the recent boom/bust does involve a melding of irrationality with an accelerator-type role for financial markets. The following quote, from the IMF's 2001 World Economic Outlook and written by distinguished economists, certainly reflects a broad public perception.

> As in past technological revolutions, the initial phase of the IT revolution appears to have been characterized by excessive optimism about the potential earnings of innovating firms. This over-optimism led for several years to soaring stock prices of IT firms, which made equity finance cheaper and more readily available, which in turn boosted investment by IT firms.

Equilibrium models in which financial markets amplify fundamental shocks represent another class of potential explanation for real bubbles, not necessarily related to irrationality. Models such as Carlstrom and Fuerst (1997), and Bernanke et al. (1999) instead embed contracting frictions which lead to endogenous variation in the cost of external finance. A large related literature explores the implications of frictions in credit markets (Holmstron and Tirole, 1997; Allen and Gale, 2000). These models can explain excessive real disinvestment in recessions, when financial constraints bind.

Contracting problems can also directly influence investment without reliance on a financial mechanism. In Philippon (2003) real business cycles are amplified by endogenous loosening of corporate governance in expansions. Models of social learning can produce herding by managers when private information is noncontractible. Scharfstein and Stein (1990) and Caplin and Leahy (1994) have used this approach to model overinvestment.

In a similar vein, what might be called local learning models study the properties of economies populated by agents who adapt their behavior in sensible, but not formally optimal, fashion given their experience. This too can lead to overinvestment cascades.[3] These models bridge the gap between rational and behavioral perspectives.

As this brief summary makes clear, there are already a variety of devices that can be used to link technological advances to overinvestment. Relative to these, the model presented here is notably different in not relying on any form of market failure or frictions. The point

[3] Evans and Honkapohja (2001) provide a comprehensive treatment of this line of research.

of this distinction is not to be doctrinaire: undoubtedly failures and frictions play important roles in the type of episodes under consideration (and the model here does not preclude these). Instead, the goal is to point out that the conventional understanding of new technology bubbles may be incomplete. Moreover, if irrational exuberance and financial amplification are not the whole story, then the common interpretation of these episodes as wastefully misallocating resources may be misguided. In the context of this paper's model, in fact, boom followed by bust is the quickest (and most efficient) adjustment path to the long-run optimum.

In modeling learning as an active process to be optimally managed, this paper follows a line of literature begun by Prescott (1972), who first considered the problem of stochastic control when the control affects the information set.[4] No separation principle applies in these settings: that is, the problem cannot be decomposed into separate estimation and optimization stages. This makes analytical solutions impossible except in highly simplified settings. Moreover, even in one- and two-period problems the effects of learning incentives can be complex and ambiguous. In some cases, the intuition that learning can motivate active experimentation (e.g. via increased output or investment) is validated. In others, the opposite intuition holds: the need to learn can induce caution and waiting while knowledge accumulates.[5] A two-period growth model similar to mine is analyzed by Bertocchi and Spagat (1998) who observe regions of both underinvestment and overinvestment relative to the full-information case. The scope for drawing policy conclusions for actual economic problems has thus been quite limited.[6]

More recently, increases in computing power have enabled analysis of the effects of optimal learning in more realistic settings. Wieland (2000b), for example, is able to deduce important implications for a monetary authority learning about an unknown money demand function while also controlling inflation. My aims are similar in scope: to be able to address the quantitative consequences of adaptive control in a dynamic equilibrium, incorporating standard utility functions, non-trivial constraints, and multiple periods. The model is still stylized and incomplete, yet it is able to offer some significant insights into how and when learning can induce bubble-like dynamics in the real economy.

The paper is organized as follows. The next section gives the details of the economic setting. The information structure and the optimization problem are described, and the weaknesses and driving assumptions of the model are discussed. Section 3 presents solutions which establish the occurrence of the overinvestment effect and demonstrate that asset prices as well as investment become inflated. Section 4 examines patterns of returns that can arise under the model, and relates these to the empirical finance literature. The effects produced in the model can be large enough to account for the anomalous returns observed in new or high growth stocks. On this basis, the model appears consistent with available measurement of price bubbles. The final section briefly summarizes the paper's contribution.

[4] These are known as adaptive control problems in the systems literature. Åström (1987) provides a brief survey.

[5] This was first observed in MacRae (1975). Some explicit results in the two-period problems are derived by Mirman et al. (1993). Datta et al. (2002) establish conditions under which experimentation can reduce information.

[6] Seminal contributions to the analysis of adaptive problems also include Chow (1975), Rothschild (1974), and Grossman et al. (1977). See Wieland (2000a) for further references.

2. The model

This section presents a description of the model and reviews its key assumptions. While the formal optimization problem is complex, the underlying intuition is straightforward. Investment in a new sector produces an externality in the form of faster learning about returns-to-scale.

2.1. Description

Consider a single-agent production economy with one good. Time is discrete and is indexed by t. For simplicity, take the time intervals to be of unit size (i.e. 1 year). At the start of each period, the agent chooses to either consume his supply, W, of the good or invest it in one of two production technologies. The first is a linear storage technology, with certain return R, which is known to the agent. The second is a risky technology, new at $t = 0$, whose returns-to-scale are not precisely known. Denote the quantity invested in the riskless and risky technologies as $K^{(0)}$ and $K^{(1)}$, respectively. Let C be consumption and $K = K^{(0)} + K^{(1)} = W - C$ be the total capital stock. The agent's problem is to choose policies $(C, K^{(0)}, K^{(1)})$ to maximize lifetime expected utility, taking into account the effects of current actions, not just on future wealth, but also on the probability measure determining *future* expectations.

This set-up is meant to capture the situation facing investors when confronted with a totally new opportunity, whose ideal scale is hard to predict.[7] Here, the riskless technology summarizes the production opportunities of the "old" economy. As such, its important feature is not that it is riskless (all the main features of the model are preserved if it is stochastic), but that its production function is known. As information about the new technology changes, the economy can adjust along two dimensions by either changing the overall level of investment or shifting resources between old and new sectors.

Wealth evolves according to the law of motion:

$$W_{t+1} = Y^{(1)}_{t+1} + Y^{(0)}_{t+1} + (1-\delta)K^{(1)}_t + K^{(0)}_t \tag{1}$$

$$= Y^{(1)}_{t+1} + (1-\delta)K^{(1)}_t + (1+R)K^{(0)}_t, \tag{2}$$

where δ is the (known, constant) depreciation rate of capital in the new sector, and the output of that sector, $Y^{(1)}_{t+1}$, is given by

$$Y^{(1)}_{t+1} = Af(K^{(1)}_t)\exp\{\varepsilon_{t+1}\}. \tag{3}$$

Here A is a known constant and ε_{t+1} is a mean-zero normal random variable with known variance σ^2.

The risky sector's production function, $f()$, is assumed to belong to a generalization of the constant elasticity class $f(K) = K^{\alpha}$. Specifically (suppressing the superscript), the

[7] While formally a real opportunity, the new technology could also represent an innovative financial opportunity, such as hedge funds or emerging markets, whose capacity to absorb investment is hard to assess.

posited form is

$$
\begin{aligned}
f(K) &= \exp\left\{\alpha\left(\frac{(K/\underline{K})^b - 1}{b}\right)\right\} \cdot 1_{\{K \geqslant \underline{K}\}} \\
&= \exp\left\{\alpha\left(\log\left(\frac{K}{\underline{K}}\right) + \frac{b}{2}\log\left(\frac{K}{\underline{K}}\right)^2 + \frac{b^2}{6}\log\left(\frac{K}{\underline{K}}\right)^3 \cdots\right)\right\} \cdot 1_{\{K \geqslant \underline{K}\}} \\
&\equiv \left(\frac{K}{\underline{K}}\right)^{\alpha(1+o(b))} \quad \text{for } K \geqslant \underline{K}.
\end{aligned} \tag{4}
$$

The form of this function is discussed further below. For now, it suffices to observe that, (i) from the last line, it reduces to the constant elasticity case for small values of b, and (ii) from the second line, it is also an extension of the translog class (Christensen et al., 1973) which would have log output quadratic in log capital.[8] The important point is that $b<0$ is a known structural constant, but α is unobservable and must be estimated. Thus, the agent knows the production function up to a parameter which determines its curvature, and hence returns-to-scale. The estimation problem is non-trivial because output of the new sector is also subject to unobservable random perturbations, for example due to exogenous cost shocks. These are taken to be independent over time. Since the stochastic shocks are unobservable, agents must choose their time t policies without knowing either of the second two terms of the output equation (3).

The information structure of the problem is as follows. At time 0, agents have a normal prior over α with mean m_0 and variance v_0 or precision $\pi_0 \equiv 1/v_0$. After making their allocation decisions, they observe $Y_1^{(1)}$ or equivalently $y_1^{(1)} = \log Y_1^{(1)}$, where

$$
y_{t+1}^{(1)} = a + \alpha\left(\frac{(K_t^{(1)}/\underline{K})^b - 1}{b}\right) + \varepsilon_{t+1}
$$

and $a = \log A$. Although they know the values of a and b they cannot separate the noise term from the term due to investment. But, conditional on the observation, they can and do update their beliefs according to Bayes' law. Because of the normality of ε and of the prior distribution, the posterior distribution stays normal at all times, with mean and variance updated according to the recursion

$$
\pi_{t+1} = \pi_t + \left(\frac{x_t}{\sigma}\right)^2, \tag{5}
$$

$$
\pi_{t+1} m_{t+1} = \pi_t m_t + \left(\frac{y_{t+1}^{(1)} - a}{x_t}\right)\left(\frac{x_t}{\sigma}\right)^2, \tag{6}
$$

where x_t is defined to be $x(K_t^{(1)}) = ((K_t^{(1)}/\underline{K})^b - 1)/b$. Hence, as the first equation shows, investing more in the risky technology directly buys more knowledge about it.

A second crucial consequence of the normal conjugate set-up is that conditional beliefs about the process $y_t^{(1)}$ stay normally distributed at all times. Specifically, after integrating

[8] A minimum investment threshold $\underline{K}$ is incorporated just to maintain the consistency of the interpretation that higher values of α are better in that they correspond to higher output for all K. This plays no important role in the results, and hereafter $\underline{K}$ will be normalized to unity.

out uncertainty about α, $y^{(1)}_{t+1}$ at time t is subjectively distributed as

$$\mathcal{N}(a + m_t x_t, \sigma^2 + v_t x_t^2). \tag{7}$$

This expression immediately demonstrates one of the key motivations for acquiring information. Although current investment increases next period's return uncertainty (x is increasing in K) and so has an ambiguous impact on the attractiveness of next period's payoff, it also always lowers v_{t+1} according to (5). This unambiguously decreases the variance of all *future* period's returns. Hence K buys an improvement in the future investment opportunity set, through higher Sharpe ratios.

To formally define the agent's problem, assume a standard time-separable utility function, $u(C)$, with subjective discount factor β. The agent is to choose policies $(C, K^{(0)}, K^{(1)})$ to maximize infinite-horizon expected utility:

$$\mathrm{E}^{\mathscr{F}_0}\left[\sum_{t=0}^{\infty} \beta^t u(C_t) \mid W_0, m_0, \pi_0\right].$$

The current state is characterized by current wealth, W_t, and the two sufficient statistics describing agent's beliefs. The somewhat redundant notation is meant to emphasize that those parameters also describe the current information set $\mathscr{F}_t$ and hence the probability measure with respect to which the expectation is taken.[9]

To complete the specification of the problem, I constrain the set of feasible policies by imposing $C \geqslant 0, K^{(0)} \geqslant 0, K^{(1)} \geqslant 0$, on the grounds that physical capital, like consumption, cannot be negative. Together with the budget constraint, $W = C + K$, the policy space may be described by the two variables $\imath \equiv K/W$ and $\omega \equiv K^{(1)}/K$.

2.2. Discussion of assumptions

The economy described above is mostly standard apart from the informational problem which is the paper's focus. Several features merit comment, however, in order to clarify which are essential for the main results.

First, the technologies of the "old" and "new" sectors look quite different in several ways. As already mentioned, taking the old sector's return to be constant is not essential. The main dynamics are all preserved if it is random, provided it is not too correlated with the new sector's returns. Likewise, having the old sector be linear is also just for convenience; the results are not affected if it too has decreasing returns-to-scale.

One can say more: the solutions below are literally isomorphic to those of an economy in which both sectors are homogeneous in the level of wealth. Let $z = K^{(1)}/W = \imath\omega$ and suppose Eq. (3) is replaced by

$$Y^{(1)}(K_t^{(1)}, W_t) = Y^{(1)}(z_t, W_t) = Af(z_t)\,\mathrm{e}^{\varepsilon_{t+1}} W_t$$

with the function $f()$ still given by (4). Then the law of motion becomes

$$W_{t+1} = [B(\imath - z_t) + \widetilde{A}(z_t)]W_t,$$

[9]Since all expectations hereafter are taken with respect to the agent's information, the superscript will be dropped.

where $B \equiv (1+R)$ and

$$\widetilde{A}(z_t) \equiv A\,\mathrm{e}^{\alpha x(z_t)+\varepsilon_{t+1}} + (1-\delta)z_t$$

with the function $x()$ as defined above. This homogeneous model would see both sectors grow at the same long-run rate, and so might be more appropriate for a long horizon analysis. In fact, the first-order conditions for this second model are exactly the same as for the original specification when $K_t^{(1)}$ is replaced by $(W_0/W_t)K_t^{(1)}$, where $W_0 = \underline{K}/\underline{z}$ is the ratio of the scaling factors in the output function. Hence the results presented for risky investment in the next section are equally valid for the homogeneous model, with $K^{(1)}$ interpreted as being scaled by the current level of wealth.

Next, the seemingly complicated form of the risky production function requires explanation. Its only purpose is to bound output while retaining the basic features of the constant elasticity class. This has important implications for tractability. The agent's beliefs, being normally distributed, will not rule out $\alpha > 1$, i.e. an explosive production function. This leads to the possibility of an infinite value function, and the problem is then not well defined.[10] An equally workable assumption would be to simply cap output, as in e.g. $f(K) = |\min(K, \bar{K})|^{\alpha}$.

It is important to point out that bounding the production functions (by taking $b < 0$ in Eq. (4)) does not induce any of the model's effects. In fact, it works against them: allowing for more explosive behavior would lead to even more overinvestment as uncertainty increases. Fig. 1 plots the production functions used here for various values of α, with $b = -0.25$ (which is the value used in the computations) on the left and $b = 0$ on the right. As the graph makes clear, the two models look quite similar locally. It is only as $K^{(1)}$ and α both get large that the difference becomes apparent.

In any model with unobservable parameters, one is justified in asking why agents are assumed to know some structural constants and not others. In reality, investors will know even less about new opportunities than the curvature parameter α. The simplified set-up here is meant to focus on just one aspect of the full problem. In thinking about the forecasting task, however, α is fundamentally different from the other parameters of the production function a (or A) and σ, in that the latter two can be learned locally. That is, even small scale experience with the technology (e.g. in the laboratory) can pin them down with high accuracy. So treating these as known seems consistent with the objective of capturing the uncertainty that could not possibly have been previously learned for a fundamentally new business.

One parameter that definitely does affect experimentation is risk aversion. The computations below will focus on the log utility case, and the main qualitative features survive under relative risk aversion levels in the single digits. With higher risk aversion, overinvestment is damped when the new sector is a large fraction of the economy, as experimentation adds significantly to consumption risk. (Robustness results are discussed further in Section 3.) In the cases explored below, however, the new industry will typically remain small with respect to total wealth. In the numerical solutions, the long-run optimal level of risky investment is usually in the vicinity of 5–10% of wealth. The impact of risk aversion only begins to be felt when there becomes a significant probability of very large returns-to-scale.

[10]The constant elasticity case can be solved on a finite horizon. Solutions exhibit the same characteristics as those computed below.

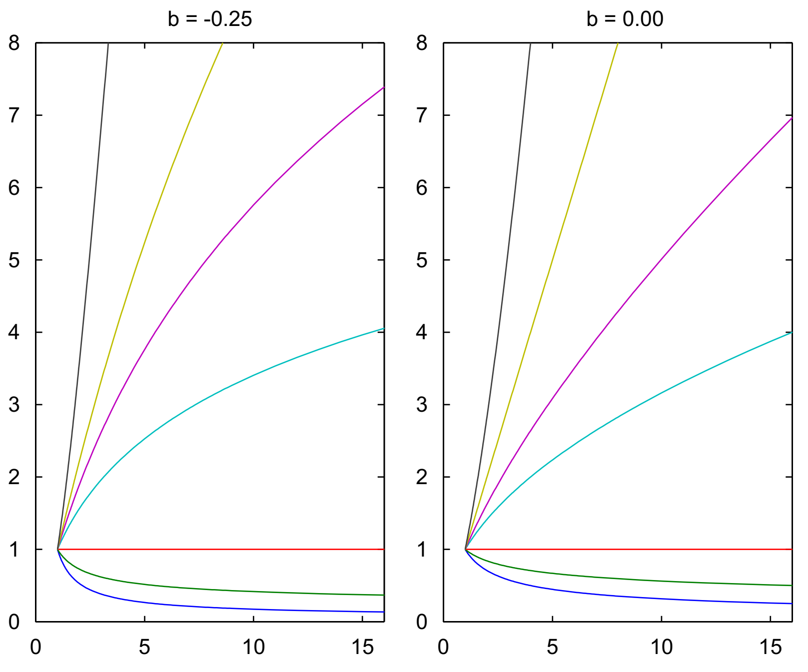

Fig. 1. Production functions. The figure shows the form of the production functions, $f(K)$, used in the paper. The left panel shows curves corresponding to $\alpha = \{-1, -0.5, 0, 0.7, 1.0, 1.25, 2\}$ for the case $b = -1/4$. The right panel is the constant elasticity case $b = 0$, and uses $\alpha = \{-0.5, -0.25, 0, 0.5, 0.7, 1, 1.5\}$.

It is worth emphasizing at this point that, although the structure of the economy is that of a two-sector real business cycle model, the focus is on the short-run dynamics of the new sector, rather than on business cycle effects. This emphasis makes the paper quite distinct from the line of research that has used similar models to tackle, for example, the equity premium puzzle or the comovement of aggregate consumption and investment. In particular, the reader should not think of the risky and riskless investments as ‘‘stocks’’ and ‘‘bonds’’, respectively. More consistent with the paper's interpretation would be to view the riskless technology as ‘‘the market’’ i.e. the aggregate of all other sectors' productive activity, so that R is the certainty-equivalent opportunity cost of funds.

An important assumption of the model is the absence of frictions or irreversibilities in adjusting the level of $K^{(1)}$. This assumption might be defended by imagining the model embedded in an overall growing economy in which actual disinvestment (net of depreciation) would be rare. Symmetric adjustment costs, too, would not diminish the incentive to arrive at the long-run optimum in the fewest possible steps, which drives overinvestment.[11] Nonsymmetric costs, however, clearly hurt the main argument by providing a disincentive to overcommit capital.

Likewise a disincentive to experiment could arise if individual firms are able to free-ride on the information externality created by the risky experimentation of others. As with

[11]Balvers and Cosimano (1990) provide an example of an active learning problem with adjustment costs in which the incentive to learn cumulates over time. Experimentation then occurs in periodic bursts when the shadow price of additional information exceeds a threshold cost.

adjustment costs, though, the effect may be to induce delay and clustering of individual actions, and not necessarily to reduce investment conditional on action being undertaken.[12] The paper will sidestep these issues by maintaining a single-agent assumption. Here the agent should be thought of as a firm. Decentralizing households does not present a problem for the results. What is important is that the firm is able to protect the information benefits of its activities. The most straightforward way to do that is to envision protected ownership rights to the new production technology.[13] This is not the only possibility. First-mover advantages and product market differentiation can effectively accomplish the same thing. Bresnahan et al. (1997) describe how rents were effectively protected in the personal computer market in the 1980s despite high rates of entry and imitation. The oft-cited network effects for internet innovators were, in fact, a strong determinant of profits (see Rajgopal et al., 2003).

The role of internalizing learning externalities in my argument suggests some valuable insight into the question raised in the introduction as to *when* technological revolutions lead to real bubbles. It is probably not a coincidence that many of the historical examples cited by commentators are of inventions with natural monopoly characteristics, such as electricity generation and transport networks. This paper's analysis highlights a possible reason for this.

3. Solutions

The situation envisioned by the model is this: at time $t = 0$ the new technology with its unknown production function is invented and agents have favorable enough beliefs about the likely returns that they invest a significant proportion of their wealth in it. Beliefs about α are characterized by the mean and variance m_0 and $v_0 \equiv 1/\pi_0$. The question of interest now is: how are agents' investment and valuation decisions affected by the second parameter, that is, by parameter uncertainty?

The results below will show that levels of investment and prices of installed capital are both rising in this uncertainty. The model's depiction of a bubble-like episode thus starts from the point where, in effect, the bubble is already inflated. The build-up occurs in a single step, from $t = -1$ (before the technology existed) to $t = 0$. The model then describes how the economy adjusts for $t > 0$ as knowledge accumulates and overinvestment dissipates.

This focus on the role of v_0 intentionally puts issues of bias in the background. Of course the model can describe episodes with much richer evolutionary dynamics depending on the initial bias $m_0 - \alpha$. For example, one case would be that agents initially underappreciate the new opportunity so that $m_0 < \alpha$ is biased downwards. Then, on average, positive surprises will follow because subjective expected output is below its true mean. With each positive surprise in $y^{(1)}$, $K^{(1)}$ will rise until the bias disappears, thus adding a longer build-up phase to the story.

Some of the simulations in Section 4 will exhibit this type of interaction with simultaneous changes in bias and uncertainty. But, for the present section, it will not be necessary to specify the true value of α at all. Instead, optimal allocations and prices will be compared with the levels that would prevail without parameter uncertainty. Hence, here

[12] These effects are analyzed in Chamley and Gale (1994) and Gul and Lundholm (1995). Bolton and Harris (1999) provide an example in which information free-riding can actually increase individual experimentation.

[13] The protection need not be perpetual, just long enough to learn the returns-to-scale parameter reasonably accurately. In the numerical results, this typically takes about 10 years.

Table 1
Parameters for numerical solutions

Parameter	Symbol	Value
Current wealth	W	100
Return on riskless technology	R	0.05
Subjective discount factor	β	1/1.05
Depreciation rate	δ	0.10
Production function (log) intercept	a	log(0.20)
Production function exponent	b	$-1/4$
Output volatility	σ	0.20
Minimum investment	$\underline{K}$	1
Investment interval	Δt	1 year
Coefficient of relative risk aversion	γ	1

"underinvestment" and "overinvestment" refer to the deviation of the optimal level $K^{(1)}$ (or ω) from what it would be at the same point in the state space if the returns-to-scale parameter were known with certainty to be equal to its current estimate m, or equivalently, if v were 0.

The full-information level of investment is the one that the economy will eventually converge to. Moreover it is also the benchmark that an observer of the economy with a long history—such as an econometrician—would deem rational *ex post*. Hence this is the appropriate comparison in seeking explanations for what appear to be anomalies in the data. In general, deviations from the full-information level can arise from both the active, strategic learning incentive and from the additional static parameter uncertainty that directly raises the subjective variance of next period's output. Another interesting question, then, is how much of the model's overinvestment is due to each effect. This can be answered by comparing the model's allocations to those that would be made by an agent who believes himself to face the same uncertain production distribution (given by Eq. (7)) but who myopically ignores the influence of his current decisions on future information. This will also permit a second characterization of overinvestment—the degree to which optimal active learning alters these myopic policies.

These topics will now be addressed in the context of a numerical analysis.

3.1. Results

To start, this subsection examines the optimal level of investment in the new sector. Section 3.2 investigates prices, and Section 4 explores expected returns to financial claims.

For the remainder of the paper, the computations will adopt the parameters shown in Table 1. The value function and optimal policies are then computed directly by dynamic programming over a discrete grid of points in the state space.[14] The computational demands of the problem make an exhaustive exploration of the parameter space infeasible. But the following discussion will indicate the extent to which the results are sensitive to the choices shown in the table.

[14] The algorithm is described in a technical appendix, available from the author upon request and included in the version of this paper available at http://www.london.edu/financeworkingpapers.html.

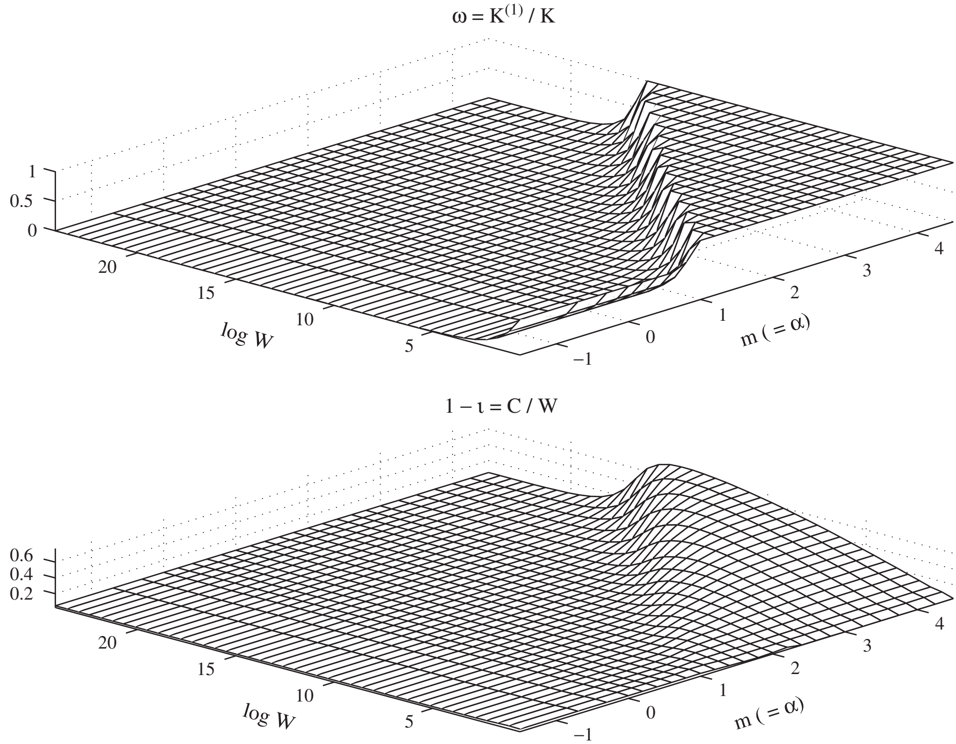

Fig. 2. Optimal policies. The figure shows optimal policies for the model of Section 2 when the returns-to-scale parameter α is known. The top panel shows the optimal share of capital allocated to the new technology. The bottom panel shows optimal consumption as a percentage of wealth. All parameter settings are as in Table 1.

The assumptions about the production parameters are supposed to describe a hypothetical emerging industry, but otherwise are meant to be standard. They imply that a unit investment yields expected payoff 0.104, after depreciation, with a 1 standard deviation shock in either direction giving a range of 0.064–0.144. A risk-neutral investor with full information would invest about 2.7 for $\alpha = 1$, rising to 6.4 and then 16.8 for $\alpha = 1.2$ and 1.4.

While the focus of the section is on the relative changes in policies as uncertainty varies, a look at the shape of the actual policies themselves reveals some noteworthy features of the model.

The top panel of Fig. 2 plots the optimal share of capital in the new sector when the posterior uncertainty v is 0.[15] For any fixed value of the returns-to-scale parameter m (which is also α in this case), the fraction declines with wealth because marginal returns decline with the allocation, $K^{(1)}$. On the other axis, as m increases, the new capital share rapidly hits 1.0 and the non-negativity constraint on the old sector's capital binds.

The lower panel of the figure shows optimal consumption as a fraction of wealth. This fraction stays fixed at about 0.05 for most of the state space. As m rises, the agent shifts

[15]The policy surfaces with positive levels of uncertainty are qualitatively very similar to the ones shown.

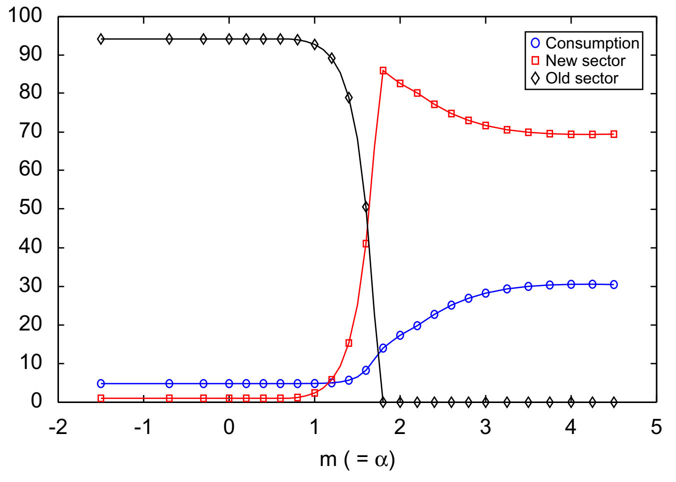

Fig. 3. Optimal quantities. The figure shows optimal policies for the model of Section 2 when the returns-to-scale parameter α is known and the level of wealth is fixed at $W = 100$. All parameter settings are as in Table 1.

investment to the new sector (see the previous panel), but consumption does not have to adjust until the old sector's capital hits 0. Once it does, the agent responds to increases in m by increasing consumption, thus *lowering* total capital and hence $K^{(1)}$.[16] Thus there are two quite different regions in terms of the response of $K^{(1)}$ to m. For present purposes, the behavior in the constrained region is of less interest as the type of innovations under consideration (e.g. the internet sector) never remotely approached anything like 100% of the capital stock.

Fig. 3 shows the optimal policies in quantity units for fixed wealth $W = 100$. For the parameter values employed, wealth will actually be stationary (since $\beta = 1/(1+R)$), and the steady state of the model is essentially described by these policies.[17] Once α is learned, the shares of each sector are fixed, and all quantities are perfectly correlated with the technology shocks of the new sector.

What happens before the agent has learned α? In the long-run, expectations about this parameter in the range of 1.0–1.4 will imply from 3% to 17% of wealth optimally allocated to the new sector. This is the order of magnitude of the sort of technological innovation that might be considered historically important. So the analysis will focus on that range of m values. For the appropriate initial level of uncertainty, the prior standard deviation for α (that is, $\sqrt{v_0}$) will be assumed to be below 0.20. For a risk-neutral agent believing $m_0 = 1.0$, this degree of uncertainty translates into a prior distribution over the optimal allocation to the new sector having a mean of 4.0 (or 4% of initial wealth); a 99th percentile of 24.0; and a 7% probability mass on 0 as the optimal allocation. These values describe a level of uncertainty about the long-run scale of a new sector which seems plausible.

[16]Also visible in the figure, the increase in consumption with m in the constrained region is smaller at smaller wealth levels. This is an artifact of the locally convex nature of the production function $f(K^{(1)})$, which entails a locally declining marginal product of capital for low values of $K^{(1)}$ and high values of α.

[17]As discussed above, these are literally the optimal policies as a percentage of wealth that characterize the homogeneous version of the model.

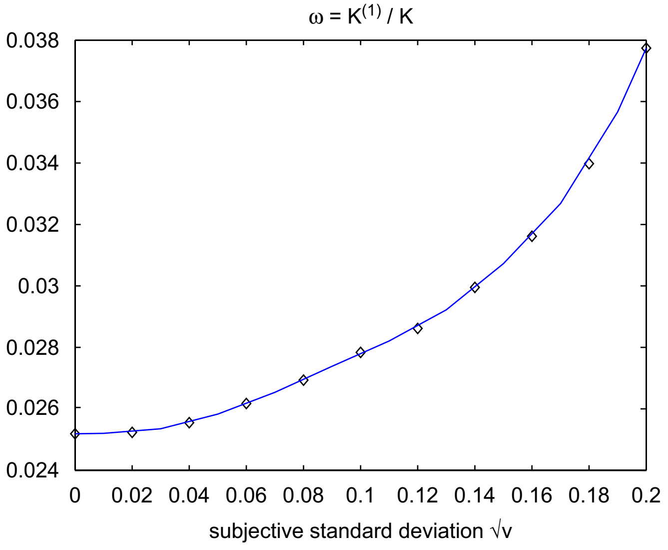

Fig. 4. Optimal investment when $m = 1.0$. The figure shows the optimal level of investment in the new sector $K^{(1)}$, as a function of the subjective standard deviation for the returns-to-scale parameter. The graph scales $K^{(1)}$ by total investment K. The subjective mean is set to $m = 1.0$. Other parameter settings are as in Table 1.

To examine the effect of returns-to-scale uncertainty, I now solve the dynamic program for successively higher levels of the subjective standard deviation over α up to 0.20. Fig. 4 plots the optimal level of new sector investment for an investor with wealth $W = 100$ whose current estimate of α is 1.0. The horizontal axis is the standard deviation $\sqrt{v}$ which, reading to the right, goes from 0—corresponding to perfect information—to 0.20. So, reading from right to left, one can interpret the curve as the slope down which investment would descend to its long-run equilibrium level as information accumulates.

And the slope is downward. Substantial overinvestment is indeed initially optimal, and this is the heart of the paper's results. In percentage terms, the plots show a decline of about 50%. The overinvestment shows no sign of being bounded above, suggesting that still larger levels would be observed were prior uncertainty increased further.[18]

Looking now across values of the prior mean, the left panel of Fig. 5 shows how the degree of overinvestment varies with m. The lines plotted are for successively larger levels of the subjective standard deviation, $\sqrt{v} = 0.00, 0.02, 0.04, \ldots, 0.20$, and now the new sector allocation is plotted relative to the full-information case. This shows that, in percentage term, the biggest effect is near $m = 0.8$. The effect falls to 0 as m declines, since for low α the risky asset is unattractive at all levels of precision.

More surprisingly, the overinvestment effect also vanishes for high m. Once the borrowing constraint binds, and the fraction of new sector capital hits 1.0, the agent can only increase investment by forgoing consumption. But with very high m, $K^{(1)}$ is a large fraction of wealth, and risk aversion curbs investment. In addition, for large levels of $K^{(1)}$, incremental learning from investment diminishes. Mathematically, the function $x(K^{(1)})$

[18]The denominator, total capital K, does not perceptibly change with $\sqrt{v}$ here. So the graph of $K^{(1)}$ itself versus $\sqrt{v}$ looks the same as the plot shown.

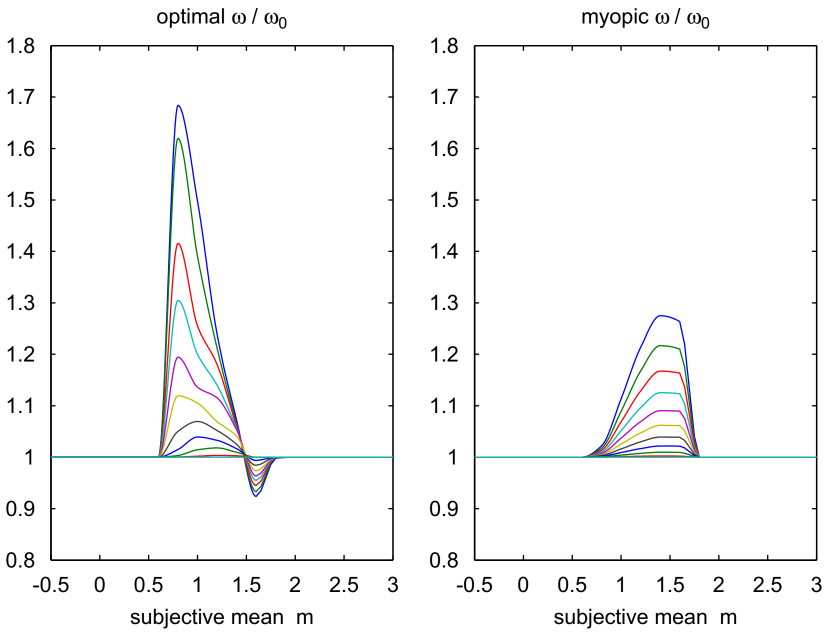

Fig. 5. Standardized optimal investment. The figure shows optimal new sector allocation as a fraction of the corresponding full-information level. The horizontal axis is the current expectation, m, about returns-to-scale. Each line is for a fixed value of the current uncertainty $\sqrt{v}$, with the outermost line being for $\sqrt{v} = 0.20$. The left graph shows the full active learning solution. The right graph shows the myopic solution that accounts for parameter uncertainty but not learning. Parameter settings are as in Table 1.

which governs gains in posterior precision (cf. Eq. (5)) approaches an upper bound. Intuitively, then, when the new sector becomes very large, the investor is unwilling to risk a little more for the sake of a vanishing benefit of learning.

While this graph demonstrates that overinvestment is a prominent feature of the optimal policies, it also exhibits a small region of underinvestment. This underscores the complexity of the solutions and the difficulty of obtaining general results. It also naturally raises the question of whether, with other parameter sets, the underinvestment regions might become more prominent and overinvestment disappear. Based on limited numerical experimentation, this appears to be rare. The overinvestment effect is still present if any of the main parameters σ, A, R, δ, or γ is varied by at least a factor of 2. Lowering σ or the subjective discount factor β do eventually kill the effect, but they do not reverse it.

Two main forces drive overinvestment in this model: the active learning incentive and the pure Jensen's inequality effect by which α risk raises expected output. To assess the relative contribution of these, we may eliminate the first motivation and solve the problem of a myopic agent who, when computing "optimal" policies, ignores the effect of current investment on future values of m and v.[19] The right panel of Fig. 5 plots the relative investment curve for this case. The myopic agent, who has only one incentive to overinvest, still does so, but by significantly less than the agent who has both incentives. The Jensen effect

[19]Mathematically, the agent sets the active learning terms $\mathrm{d}m_{t+1}/\mathrm{d}K_t^{(1)}$ and $\mathrm{d}\pi_{t+1}/\mathrm{d}K_t^{(1)}$ to 0 in the investment first-order condition.

is definitely part of the story, and is even the larger part between about $m = 1.3$ and 1.6. For m values in the range 0.8–1.2, however, the contribution of convexity to the total effect is marginal, and the active learning effect by itself induces about twice as much overinvestment (i.e. 60% vs 30%) as the most that convexity generates. This demonstrates that the active learning facet of the problem is, indeed, crucial to understanding the economy's behavior.

The graphical results presented so far give a static picture of investment as a function of the state variables. It is also easy to see how the new industry bubble would evolve dynamically as information precision increases. At $t = 0$ the agent with wealth 100 learns of the new opportunity and investment leaps from nothing to about 3.6 if $m_0 = 1.0$. Suppose this belief is unbiased and α is also 1.0. Then a simulation of the next few years would see $K^{(1)}$ decline, stochastically but systematically, to about 2.4. Referring back to Fig. 4, the rate of disinvestment can also be calculated. For $K^{(1)} \approx 3.6$ precision increases by $\Delta\pi = [x(K^{(1)})/\sigma]^2 \approx 30$. This suggests a decay from $\pi = 1/0.20^2$ to $1/0.06^2$, for example, in about 8 years. Hence the model implies what seems like a plausible duration for adjustment of the real side of the economy following a "technological revolution".

Clearly the overinvestment in this model does not occur in response to signals from overheated asset markets: allocations are determined directly from the economic primitives. Nevertheless, the next section shows that the model can still account for the *appearance* of a link between investment and the cost of capital since claims to the new technology appear "overpriced" at the same time as the overinvestment is occurring.

3.2. Asset prices

No financial markets are involved in solving the model. However, following standard practice, asset prices can be characterized by the shadow prices of claims to different cash-flows. In particular, the stock price of the new industry is the value placed on its flow of net dividends $D_t^{(1)} = Y_t^{(1)} - K_t^{(1)} + (1-\delta)K_{t-1}^{(1)}$. This interpretation imagines a portfolio problem faced by the representative agent (or a collection of households with identical preferences) at time t but after the time t allocations have been made. Since the equilibrium allocation policies are optimal for this agent, these policies can be regarded as exogenous for purposes of the portfolio problem. Hence the pricing calculation proceeds as if the same agent were living in an endowment economy.

To understand the resulting prices, it is important to realize that the same representative agent applies different first-order conditions to the portfolio problem from those that apply to the production decision. This statement is not true in many production-based asset pricing models, and so deserves comment.

In the portfolio problem, with exogenous dividends, the situation is standard, and the investor's first-order condition can be written as

$$q_t^{(1)} \equiv \frac{P_t^{(1)}}{K_t^{(1)}} = \mathrm{E}_t\left[\beta\frac{u'(C_{t+1})}{u'(C_t)}\left\{\frac{P_{t+1}^{(1)}}{K_t^{(1)}} + \frac{Y_{t+1}^{(1)}}{K_t^{(1)}} - \frac{K_{t+1}^{(1)}}{K_t^{(1)}} + (1-\delta)\right\}\right], \tag{8}$$

where $P_t^{(1)}$ is the dividend claim's price. Similarly, without any uncertainty about the production parameters, standard arguments (using the envelope theorem and the consumption first-order condition) imply that the optimal time t allocation $K_t^{(1)}$ must satisfy

$$\mathrm{E}_t\left[\beta\frac{u'(C_{t+1})}{u'(C_t)}\frac{\mathrm{d}Y_{t+1}^{(1)}}{\mathrm{d}K_t^{(1)}}\right] = \frac{(R+\delta)}{(1+R)}, \tag{9}$$

where $\mathrm{d}Y^{(1)}_{t+1}/\mathrm{d}K^{(1)}_t$ denotes $Af'(K^{(1)}_t)\exp(\varepsilon_{t+1})$ (cf. Eq. (3)). In a standard production problem, these two conditions could be combined with a simple backward induction argument to prove that $q^{(1)}_t = 1$ for all t. Now that argument would fail for several reasons.

One reason is the form of the production function, for which the marginal return to new sector capital is not equal to its average return, a necessary condition to keep $q^{(1)} = 1$.

Second, the active learning aspect of the problem implies that there will be a difference between portfolio investor's and producer's marginal rates of substitution in this economy. The actual production first-order condition is

$$\mathrm{E}_t\left[\frac{\partial J_{t+1}}{\partial W_{t+1}}\frac{\mathrm{d}W_{t+1}}{\mathrm{d}K^{(1)}_t}+\frac{\partial J_{t+1}}{\partial m_{t+1}}\frac{\mathrm{d}m_{t+1}}{\mathrm{d}K^{(1)}_t}+\frac{\partial J_{t+1}}{\partial \pi_{t+1}}\frac{\mathrm{d}\pi_{t+1}}{\mathrm{d}K^{(1)}_t}\right]=0, \tag{10}$$

which does not reduce to (9) due to the second and third terms. Moreover these learning terms have no counterparts in the portfolio problem, as financial asset decisions do not affect the rate at which information accumulates.

Finally, the shadow price $P^{(1)}$ is determined without imposing portfolio constraints, whereas the production decision is made under non-negativity quantity requirements.

For these three reasons, the behavior of the price of risky capital is complicated in this model and cannot be deduced analytically. Instead, I perform the iterated integration in (8) numerically over the state space, using the optimal quantities computed above.

Before going to the results, one other price computation should be noted. In the usual endowment economy fashion, one can imagine a net-zero-supply riskless one-period claim, and immediately deduce its price

$$\mathrm{E}_t\left[\beta\frac{u'(C_{t+1})}{u'(C_t)}\right]\equiv\frac{1}{1+r_t},$$

which defines the risk-free rate r_t. This will not always be equal to R, the one-period return on the riskless technology, however, since once allocations are made, claims on that technology are in fixed supply. Whenever the production decision is constrained such a claim's price can be less than $1/(1+R)$. In thinking about financial returns, then, r_t is the relevant riskless rate and it is stochastic.

How is the price of the risky claim affected by returns-to-scale uncertainty? The answer is shown in Fig. 6. Here the price–capital ratios for non-zero levels of the subjective standard deviation $\sqrt{v}$ are plotted as a fraction of the full-information value at the same point in the state space. Hence the ratios measure asset prices relative to what, with hindsight, would look like the "right" level. The left-hand graph is the case of optimal active learning. The outermost line (furthest from 1.0) corresponds to a standard deviation $\sqrt{v} = 0.20$, with lower levels of uncertainty being progressively closer to unity. If initial beliefs about the new technology are characterized by a standard deviation 0.20 and a mean of between 0.9 and 1.7, the graph shows that asset prices will indeed be inflated, with the inflation exceeding 30% at $m = 1.2$. This range of m values aligns closely with the range of overinvestment shown in Fig. 5, which illustrates that the model does deliver significant bubbles in asset markets that coincide with periods of seemingly excessive investment.

Apparent inflation of asset prices relative to "fundamental" levels is a common feature of models with parameter uncertainty. This is due to another Jensen's inequality result. Cash-flows are convex in the unobserved parameter, so uncertainty enhances the option-like value of these limited liability claims. Here, we can assess the degree to which the price

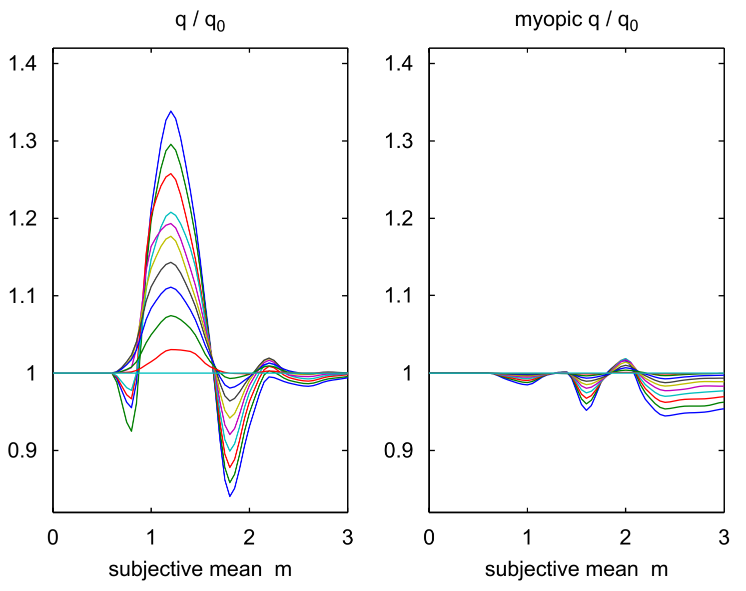

Fig. 6. Standardized price–capital ratios. The figure shows lines of constant subjective standard deviation for the price–capital ratio, q, of the new sector as a function of the subjective mean. The q values are standardized by the corresponding perfect-information level q_0. The outermost line corresponds to $\sqrt{v} = 0.20$. The left panel corresponds to the full optimizing model. The right panel corresponds to the myopic solution. All parameter settings are as in Table 1.

inflation is similarly due to convexity effects by solving for prices in the myopic case introduced above. As before, the myopic investor is valuing payoffs under the same distribution for future output as in the full model, but ignoring variation in his future information set. So asset prices for him capture the convexity of output, without the added effects of active learning. The right-hand panel of Fig. 6 shows the results.

Now there is almost no variation in $q^{(1)}$ with the level of uncertainty. The myopic investor is *not* willing to pay more for the risky claim as v rises. (In fact, for high m the reverse is true. The additional uncertainty carries a positive risk premium.) The inflated prices in the left-hand panel are not, then, attributable to the mechanical effect of increased variance. Intuitively, the investor in the myopic model does not foresee the cash-flow benefits that active learning will deliver in the future, and so is unwilling to pay more for them.

The active learning model does not imply that asset prices increase with uncertainty everywhere. The left panel of the figure also features prominent underpricing intervals at the edges of the overpricing region. These occur as investment policies approach one of the constraints. For m between 0.5 and 0.9, new sector investment $K^{(1)}$ is declining toward its floor of 1.0. Here the pricing function is locally concave in m, so m-risk is undesirable. Similarly, for m above 1.7, as ω rises to its limit of 1.0, m-risk begins to accrue a large risk premium since consumption rises sharply with m. These underpricing zones are thus more limiting cases of the model than general features. As argued above, m values in the overpricing range produce values of real quantities that are the most plausible depictions of a "technological revolution". In this range, $K^{(1)}$ is unconstrained and significant, yet not so large as to dominate the entire economy.[20]

The calculations in this section have demonstrated that the learning problem modeled here may account for at least part of the overinvestment and inflated asset prices that seem to characterize the early years of important new industries. At this point it is hard to gauge exactly how realistic the numbers are, because our empirical knowledge of technological revolutions is limited. Even aside from the infrequency of such episodes, it is extremely difficult in any industry at any time to assess the "correct" levels of investment or prices, with respect to which excesses can be compared. In the next section, the question is approached from a different angle by asking whether the model as parameterized here is consistent with documented patterns of asset returns for new and high-growth stocks.

4. Returns

This section will examine how expected returns to shares in the new sector vary along three dimensions. First, in analogy with the literature on initial public offerings (IPOs), it looks at expected returns as a function of age. Next, it considers how expected returns are related to q over time.[21] Here the link is to the literature on market predictability. Last, it examines the relations between size, book-to-market, and returns under the model to see how they compare to the corresponding well-known patterns in the cross-section of equity returns.

The implicit argument underlying these comparisons is that the appearance of predictably negative excess returns following actual technological revolutions—which itself is hard to document rigorously—is likely to be driven by the same underlying mechanism producing the apparent overpricing documented in these three empirical settings. So, to the extent that the model can help account for familiar anomalies, its claim to describe "new economy" bubbles can be said to pass an important test.

This logic applies most naturally to the case of IPOs. Although the model is meant to depict the emergence of a sector, not a firm, and although new firms seldom introduce truly revolutionary technologies, nevertheless the data on IPOs do include several waves of industries that were genuinely novel and were thus likely to be characterized by the particular type of uncertainty modeled here: uncertainty about their economies of scale.

Suppose an observer had a sample of such issues for which agents' *ex ante* beliefs were characterized by the mean $m_0 = 1.0$ and variance $v_0 = 0.04$. What excess returns for these stocks would be expected in event time as investors learn about the industry's characteristics?

The answer to this question depends on the initial errors $m_0 - \alpha$ in the sample. For any given value of the true α, the economy can be simulated by drawing from the true distribution of output shocks ε_t, and tracing the evolution of prices, dividends, and the riskless rate. The top panel of Fig. 7 performs this computation under the assumption that the value of α is what it is believed to be initially: 1.0.

This figure shows that average returns are large and negative for at least the first 10 years of event time. As time passes and investment produces information, the excess returns

[20]The underpricing regions raise the interesting possibility that adjustment paths which involve large revision of beliefs may exhibit seemingly irrational oscillation about "fundamentals", giving the appearance of excess volatility while the bubble deflates.

[21]In this section the superscript on q and P will be suppressed for simplicity.

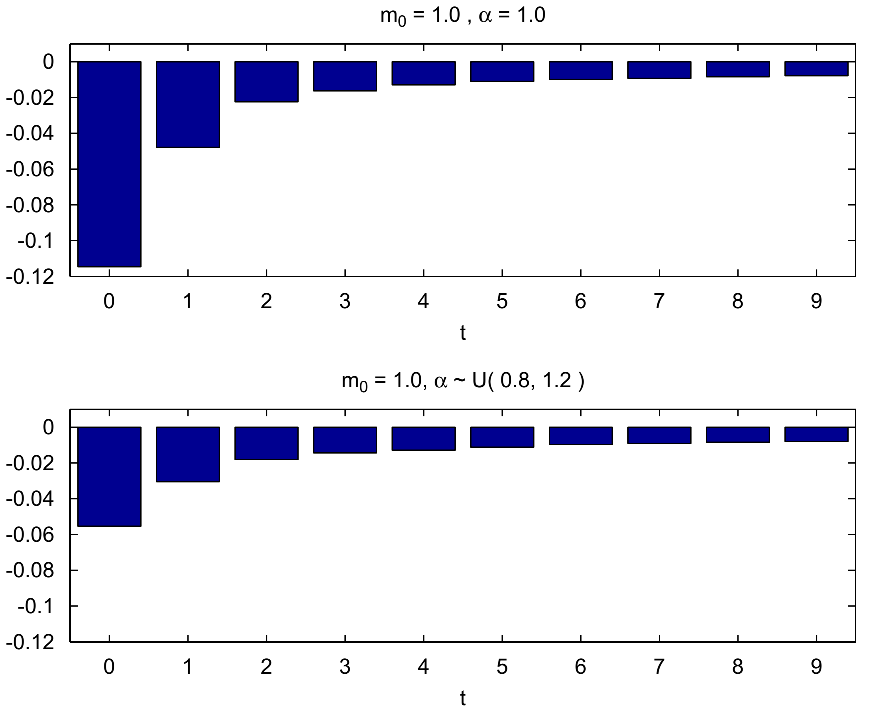

Fig. 7. Expected excess returns in event time. The figure shows annual excess returns for the risky asset as a function of time, starting from $m_0 = 1.0$ and $\sqrt{v_0} = 0.20$ averaged over 100 000 paths. In the top panel, the true value of α is also 1.0. In the bottom panel, α is uniformly distributed on [0.8 1.2]. All other parameter settings are as in Table 1.

approach 0. Hence, the apparent underperformance is entirely due to the presence of parameter uncertainty. The pattern is consistent with findings in the empirical IPO literature. While measurement issues are somewhat contentious, typical results broadly corroborate the conclusions of Ritter (1991) who reports abnormal excess buy-and-hold returns in the order of -5% to -10% per year for the first 3–5 years. The model's numbers in this simulation are not as large, and perhaps not as persistent, but have also not been specifically calibrated.

The key to understanding the above result lies in the assumption about the sample properties of the error in the initial expectations $m_0 - \alpha$. Suppose the econometrician has selected a sample of firms (or sectors) whose long-run size turned out to be between 1% and 10% of aggregate wealth. In the model, this would imply that the true α's were in the interval [0.8 1.2]. The bottom panel in Fig. 7 repeats the previous exercise, still fixing $m_0 = 1.0$ and $v_0 = 0.04$, but now also drawing α uniformly from this range. The pattern in the upper panel is preserved, with significantly negative average excess returns still. However the magnitudes are reduced due to the asymmetric influence of positive surprises: cases where the technology turns out to yield large returns.

Were the dispersion of initial errors to be widened further, these extreme positive cases would gradually reduce the average underperformance, and eventually reverse it. In fact, if the simulation were to draw the initial value of $m_0 - \alpha$ from the *agent's prior* distribution, it would necessarily produce slightly positive average excess returns. This is because, for the agent in the model, there is nothing anomalous about the returns to new issues. The standard consumption CAPM holds in this economy, and expected excess returns exactly compensate for covariance with consumption. From the results in the last section, prices, dividends, and consumption all rise with m. So, if anything, subjective expected excess returns must be positive.

There is no particular reason why the distribution—across firms—of actual errors in a given IPO sample should coincide with the subjective distribution of agents concerning each individual firm. But even if they did, the odds would still favor observing underperformance in small samples. A typical study might use data spanning 50 years and include observations of 10 new, independent industries. Fig. 8 plots the distribution of realized 5-year average performance over 10 000 replications of samples of 10 observations when each individual sample draws α's from the agent's prior. Over 60% of the mass in the histogram is in the negative range. The "p-value" for an average underperformance of -20% is 0.17. These performance numbers do not include any risk adjustment. The simulated returns would look even worse if the econometrician used any benchmark higher than the riskless rate.

Next, consider the conditional properties one would expect to observe in equity returns from a single time-series realization in the model economy. A solid empirical case has been made by e.g. Campbell and Shiller (1998) that when the market appears inflated based on "fundamental" measures, aggregate returns over the intermediate term are low. Shiller (2000) directly links such inflation to periods of belief in a "new economy", which, in the context of the model, could be viewed as coinciding with the arrival of a new production

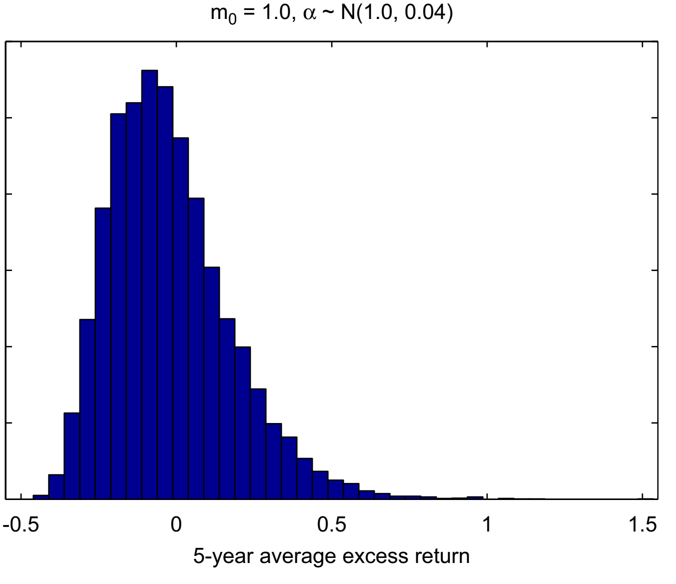

Fig. 8. Distribution of expected excess returns. The figure plots the sample histogram of mean 5-year excess returns from 10 000 simulated samples of 10 independent new industries, when $m_0 = 1.0$ and $\sqrt{v_0} = 0.20$ and the true value of α is drawn from the subjective prior distribution. All parameter settings are as in Table 1.

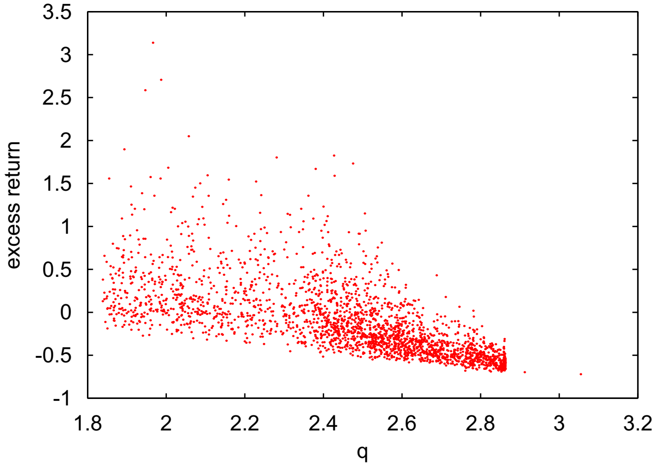

Fig. 9. Excess returns versus q. Realized excess returns are plotted against beginning of period q for 4000 simulated histories of length $T = 50$. The true value of α is 1.2, and the initial subjective belief has variance $v_0 = 0.04$. The initial expectational error $m_0 - \alpha$ is drawn randomly from a normal distribution with mean 0 and variance 0.04. All parameter settings are as in Table 1.

technology. It is interesting to ask whether the uncertainty that comes with this technology presents an alternative to "irrational exuberance" in explaining the predictability.[22] Here the model provides a natural measure of asset price inflation in q, the ratio of the market value of the new sector to the amount of installed capital in it.

Suppose this sector has true returns-to-scale $\alpha = 1.2$, for example, which implies that in the long-run the new sector will constitute about 10% of the capital stock. Given particular initial beliefs, the economy will undergo a single transient adjustment to the full-information steady-state as knowledge about α accumulates. While the model is not rich enough to capture the repeated arrival of new technologies or stochastic shocks to the scale parameter of the one technology, the dynamics of the single adjustment phase can still be compared to the historical patterns.

These dynamics are readily calculable via simulation. Fig. 9 plots the relation between excess returns and q in 4000 realization of the model economy, each of 50 years. The simulations all start with initial uncertainty about the technology having standard deviation 0.20, or $v_0 = 0.04$. The initial expectation m_0 is then drawn from a normal distribution with the same variance, and mean equal to the true value 1.2. With this sampling procedure, the simulated distribution of the initial error $m_0 - \alpha$ is exactly what the agent believes it to be.

The figure shows a strong negative relationship between future returns and q. This relationship is also consistent *within* the histories. Regressing returns on lagged q for each one individually, more than 60% of the series yield OLS t-statistics below -2.

[22]While the risky asset in the model only represents the new sector, not the aggregate stock market, one might expect that the time-series predictability would be most apparent in the sector driving the "new economy". A fuller model of the non-new sector would be needed to directly address, for example, any spillover effects from the new industry to the prices of other assets.

How does this effect arise? As noted above, asset prices in this economy obey the consumption CAPM. Yet consumption risk is high in high m states, in which q is also high—which suggests if anything a positive relationship.

There are actually two distinct mechanisms underlying the results. One is that, as in the IPO example, the simulated economy does not include "enough" chance of a very high α. Indeed, along each history, α is the same. This source of predictability derives entirely from the sampling design across repeated histories. Yet it does not explain why the predictability relationship still holds along the *individual* histories themselves. This is due to a different effect that was not present in the IPO example.

The second reason why high returns tend to follow low q (and vice versa) is the correction of inferential errors about α, which are apparent *ex post* though not *ex ante*. Once a true value of α is fixed, higher errors (due to positive output surprises) produce higher prices along each history. Since q always converges to its full-information level, higher prices mean negative future returns. There is nothing irrational about this however. The agent does not know at the time what values of q are "low" or "high".

As a final point of comparison, consider the cross-sectional relation between expected returns and characteristics of the new sector. Since $1/q$ provides a natural proxy for the asset's book-to-market ratio, and its size can be proxied by P (the market value of its equity), we can ask whether the model throws any light on the well-known relation between these characteristics and expected returns. To the extent that these traits capture "value" or "glamour" (Lakonishok et al., 1994), the empirical evidence supports the notion of systematic bias in assessing growth prospects, which is one explanation for the linkage between technological revolutions and bubbles. Does the rational experimentation story here provide an alternative?[23]

Since, in the model, the characteristics are functions of the state, to make the desired comparison, a cross-sectional distribution over these states needs to be specified. The actual distribution that might be found in a sample of firms at any one time is hard to guess. Purely for illustrative purposes, and to ensure a reasonable spread in the sorting variables, state vector $\{m, v\}$ is chosen by drawing uniformly over the interval $[0.6\ 1.4] \times [0.00\ 0.04]$, fixing all other parameters as before. Doing this 250 000 times yields a population which is then sorted into quintiles of P and q, yielding "portfolios" of $N = 10\,000$ stocks. Table 2 shows the average characteristics of each portfolio. The effect of sorting on P and q is to produce an almost monotonic spread in, respectively, the first and second moments of beliefs about returns-to-scale.

The key issue in computing expected returns is the distribution of the true α. Here, beliefs across each portfolio will be assumed to be unbiased. That is, for each stock in the kth portfolio, 1-year expected returns will be computed under the assumption that $\alpha_k = 1/N\sum_i^N m^{(i)}$ where the sum runs over the N stocks in the portfolio. This choice isolates the component of the cross-section of returns that is independent of the cross-section of idiosyncratic estimation errors.

Table 3 shows the average expected excess return per portfolio. The overall pattern is precisely that of Fama and French (1992). Returns mostly decline with market value, and, holding size fixed, tend to be inversely proportional to q. The difference between the

[23] This is, again, not an issue the model can speak to entirely rigorously. There is no cross-section of risky assets in the model. Correctly speaking, all one can compare are characteristics and expected returns in a cross-section of *economies*, each with a single new technology. This is the actual exercise undertaken.

Table 2
Cross-sectional characteristics

Size (P)	Book-to-market ($1/q$)				
	Low	2	3	4	High
Subjective mean					
Small	0.6497	0.6447	0.6667	0.7050	0.7560
2	0.7806	0.8268	0.8792	0.8823	0.8322
3	1.0311	1.0213	1.0059	0.9764	0.9523
4	1.1828	1.1768	1.1631	1.1315	1.1337
Large	1.3610	1.3258	1.3097	1.3065	1.2908
Subjective variance					
Small	0.0130	0.0130	0.0236	0.0223	0.0152
2	0.0192	0.0153	0.0144	0.0183	0.0317
3	0.0285	0.0220	0.0185	0.0172	0.0171
4	0.0322	0.0233	0.0193	0.0192	0.0103
Large	0.0328	0.0272	0.0223	0.0159	0.0091

The table shows average characteristics for a pseudo cross-section of 250 000 risky assets sorted by size and book-to-market. The risky assets are generated by drawing state vectors $\{m, v\}$ uniformly over the range $[0.6\ 1.4] \times [0.00\ 0.04]$, and fixing $W = 100$ and all other parameters as in Table 1.

Table 3
Pseudo cross-section of expected returns

Size (P)	Book-to-market ($1/q$)				
	Low	2	3	4	High
Small	−0.0002	0.0002	−0.0007	0.0003	−0.0001
2	−0.0007	−0.0007	−0.0022	−0.0027	−0.0132
3	−0.0394	−0.0146	−0.0023	−0.0067	0.0004
4	−0.0858	−0.0419	−0.0170	−0.0123	−0.0021
Large	−0.0790	−0.0555	−0.0356	−0.0112	0.0011

The table shows expected annual excess returns, computed via numerical integration, for a pseudo cross-section of 250 000 risky assets sorted by size and book-to-market. The risky assets are generated by drawing state vectors $\{m, v\}$ uniformly over the range $[0.6\ 1.4] \times [0.00\ 0.04]$, fixing $W = 100$ and all other parameters as in Table 1.

extreme cells (upper right and lower left) is 7.9% a year, which corresponds to 64 basis points per month. Fama and French (1992) find about 70 basis point a month. Of course, in any real cross-section there is important variation in other characteristics (leverage, volatility, etc.) which is not captured here. Also missing is the variation due to bias. At any given time, a given sample of firms will include ones whose recent history contained good random shocks. These will have higher current P and q and lower expected returns. Likewise firms with recent unlucky shocks will have P and q biased down, and positive expected returns. Hence—importantly—bias in the model goes the same way as the pure information effect, and would reinforce the cross-sectional spread shown in the table.

As with the other exercises in this section, the apparent anomalies in the model data are attributable to differences between agents' probability distribution for α and the actual distribution induced in the construction of the sample. As in the IPO example, the portfolios here have a cross-sectional distribution of estimation errors that does not coincide with the subjective distribution for each constituent. Thus the calculations suggest that at least part of the size and book-to-market effect may be due to the addition of conditioning information about technology that is not observable *ex ante*.

The goal of this section has been to show that, to the extent that there is empirical support for the linkage of technological revolutions and financial excess, the model is at least consistent with that evidence. Subtle differences in conditioning information, small sample effects, and the presence of idiosyncratic (i.e. not irrational) estimation errors can each lead to the appearance of anomalies whose sign and magnitude are surprisingly similar to findings in the empirical finance literature. It is perhaps worthwhile to stress that the examples presented here are not the result of a calibration exercise. I have not fully explored the flexibility afforded by the model's free parameters (e.g. σ, b, R) or the utility specification. So despite its obvious shortcomings as a depiction of the emergence of a new industry, it seems reasonable to conjecture that the patterns can be preserved in more general settings.

5. Conclusion

This study has been motivated by the intriguing parallels between the recent IT boom/bust and earlier historical technological revolutions. There appears to be widespread public acceptance of two ''stylized facts'' about these parallels: First, that there is something inevitable about the linkage between innovation and bubbles (exemplified by the paper's epigraph); and second, that the chain of causation runs from irrational financial overreaction to real overinvestment (as seen in the quotation from the IMF report in the introduction).

I propose a model which accounts for both perceptions, and yet suggests that they are fundamentally incomplete. The model implies that these bubble-like episodes are likely to accompany the emergence of new industries with particular characteristics: uncertainty about returns-to-scale and a competitive setting that protects production knowledge from free-riding. Given these conditions, predictable negative returns to financial claims and overshooting of real overinvestment can both result without the former having any role in determining the latter.

Most importantly, the model does not include market failures, sub-optimal learning, or irrationality. While these are not incompatible with its mechanism, they do have strikingly different implications about the consequences of such episodes. In the completely frictionless case examined here, bubbles are actually the most efficient way to achieve the right long-run level of investment. If that is so, trying to prevent, regulate, or even identify them *ex ante* may be misguided.

References

Allen, F., Gale, D., 2000. Bubbles and crises. Economic Journal 110, 1–20.

Åström, K.J., 1987. Adaptive feedback control. Proceedings of the IEEE 75, 185–217.

Balvers, R.J., Cosimano, T.F., 1990. Periodic learning about a hidden state variable. Journal of Economic Dynamics and Control 17, 805–827.

Bernanke, B., Gertler, M., Gilchrist, S., 1999. The financial accelerator in a quantitative business cycle framework. In: Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics. North-Holland, Amsterdam.

Bertocchi, G., Spagat, M., 1998. Growth under uncertainty with experimentation. Journal of Economic Dynamics and Control 23, 209–231.

Bolton, P., Harris, C., 1999. Strategic experimentation. Econometrica 67 (2), 349–374.

Bordo, M., 2003. Stock market crashes, productivity boom busts and recessions: some historical evidence. Rutgers University Working Paper.

Bresnahan, T.F., Stern, S., Trajtenberg, M., 1997. Market segmentation and the sources of rents from innovation: personal computers in the late 1980s. RAND Journal of Economics 28 (special issue), S17–S44.

Campbell, J.Y., Shiller, R.J., 1998. Valuation ratios and the long-run stock market outlook. Journal of Portfolio Management 24, 11–26.

Caplin, A., Leahy, J., 1994. Business as usual, market crashes, and wisdom after the fact. American Economic Review 83 (4), 548–565.

Carlstrom, C.T., Fuerst, T.S., 1997. Agency costs, net worth, and business fluctuations: a computable general equilibrium analysis. American Economic Review 87 (5), 893–910.

Chamley, C., Gale, D., 1994. Information revelation and strategic delay in a model of investment. Econometrica 62 (5), 1065–1085.

Chow, G.C., 1975. A solution to optimal control of linear systems with unknown parameters. The Review of Economics and Statistics 57 (3), 338–345.

Christensen, L.R., Jorgenson, D.W., Lau, L.J., 1973. Transcendental logarithmic production frontiers. The Review of Economics and Statistics 55, 28–45.

Datta, M., Mirman, L.J., Schlee, E.E., 2002. Optimal experimentation in signal-dependent decision problems. International Economic Review 43 (2), 577–607.

Economist, 2000. Bubble.com.

Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton.

Fama, E.F., French, K.R., 1992. The cross section of expected stock returns. Journal of Finance 47, 427–466.

Gilchrist, S., Himmelberg, C.P., Huberman, G., 2002. Do stock price bubbles influence corporate investment? Columbia University Working Paper.

Grossman, S.J., Kihlstrom, R.E., Mirman, L.J., 1977. A Bayesian approach to the production of information and learning by doing. Review of Economic Studies 44, 533–547.

Gul, F., Lundholm, R., 1995. Endogenous timing and the clustering of agents' decisions. Journal of Political Economy 103 (5), 1039–1066.

Holmstrom, B., Tirole, J., 1997. Financial intermediation loanable funds, and the real sector. Quarterly Journal of Economics 112 (3), 663–691.

International Monetary Fund, 2001. World Economic Outlook.

Jovanovic, B., 1997. Learning and growth. In: Kreps, D.M., Wallis, K.F. (Eds.), Advances in Econometrics: Theory and Applications, Seventh World Congress, vol. II. North-Holland, Amsterdam.

Lakonishok, J., Schleifer, A., Vishny, R.W., 1994. Contrarian investment, extrapolation and risk. Journal of Finance 49, 1541–1578.

MacRae, E.C., 1975. An adaptive learning rule for multiperiod decision problems. Econometrica 43 (5–6), 893–906.

Mirman, L.J., Samuelson, L., Urbano, A., 1993. Monopoly experimentation. International Economic Review 34, 549–564.

Perez, C., 2003. Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages. Edward Elgar, Northampton.

Philippon, T., 2003. Corporate governance over the business cycle. NYU-Stern School Working Paper.

Polk, C., Sapienza, P., 2002. The real effects of investor sentiment. Northwestern University Working Paper.

Prescott, E.C., 1972. The multiperiod control problem under uncertainty. Econometrica 40 (6), 1043–1058.

Rajgopal, S., Venkatachalam, M., Kotha, S., 2003. The value relevance of network advantages: the case of e-commerce firms. Journal of Accounting Research 41 (1), 136–162.

Ritter, J.R., 1991. The long-run performance of initial public offerings. Journal of Finance 46, 3–27.

Rothschild, M., 1974. A two-armed bandit theory of market pricing. Journal of Economic Theory 9, 185–202.

Scharfstein, D., Stein, J., 1990. Herd behavior and investment. American Economic Review 80 (3), 465–479.

Shiller, R.J., 2000. Irrational Exuberance. Princeton University Press, Princeton.

Wieland, V., 2000a. Learning by doing and the value of optimal experimentation. Journal of Economic Dynamics and Control 24, 501–534.

Wieland, V., 2000b. Monetary policy parameter uncertainty and optimal learning. Journal of Monetary Economics 46 (1), 199–228.