

Creative Destruction and the Rational Evolution of Bubbles*

Daniel Andrei[†]
Bruce I. Carlin[‡]

January 26, 2017

ABSTRACT

Creative destruction not only involves bringing new technology to market, it imposes higher risk on the future of existing assets. We contrast the asset pricing implications of creative destruction in a setting in which non-cooperative agents compete for market share with a socially optimal benchmark (i.e., a one agent case). Compared to the one-agent case, non-cooperative behavior leads to over-investment in uncertain projects, higher asset prices and risk premia, and price reversals, all of which resemble a bubble. However, these pricing patterns solely arise from competitive behavior and do not require information asymmetry, behavioral biases, or financial frictions to arise. As such, the explanation that bubbles occur when the price of an asset exceeds the asset's fundamental value is sufficient, but not necessary, to rationalize the pricing patterns that we observe during technology booms.

*We are grateful to Tony Bernardo, Murray Carlson, Lorenzo Garlappi, Mark Huson, Francis Longstaff, Jeff Pontiff, and Avindhar Subramanyam for their comments and input.

[†]Anderson School of Management, University of California, Los Angeles, 110 Westwood Plaza, Los Angeles, CA 90095, daniel.andrei@anderson.ucla.edu.

[‡]Anderson School of Management, University of California, Los Angeles, 110 Westwood Plaza, Los Angeles, CA 90095, bruce.carlin@anderson.ucla.edu.

“There is much more to a bubble than a mere security price increase. There is innovation, displacement of existing firms, creation of new ones, and more generally a paradigm shift as entrepreneurs and investors rush toward a new Eldorado.” —*Greenwood, Shleifer, and You (2016)*

1 Introduction

According to Schumpeter (1934), invention and entrepreneurship are two distinct activities. And it is the entrepreneurs who are responsible for creative destruction because they are tasked with bringing innovations to market. Yet, there is a third party, *the financier*, who has been largely ignored, but plays an important role. The comparative advantage of the financier is to assess risk and allocate scarce resources efficiently in the market. Within Schumpeter’s framework, the financier’s responsibility would be to govern how much entrepreneurship comes to market, based on the risk of new projects and the uncertainty that creative destruction imposes on existing assets. This is the focus of this paper.

Typically, when we consider adding a new asset to a well-diversified portfolio, we take the variance-covariance matrix of the existing assets as given so that there is increased diversification with the new asset and less risk. But, creative destruction is different because it potentially makes the future of existing assets more risky (Kung and Schmid, 2015) and may endogenously change the variance-covariance matrix. Gârleanu, Kogan, and Panageas (2012) refer to this as “displacement risk” and show that it can rationalize both the existence of the growth-value factor in returns, as well as the equity premium. Because creative destruction may divert scarce resources away from existing assets or alter their growth options and capabilities, it is not surprising that there appears to be a risk premium associated with it (Grammig and Jank, 2015).

Another difference is that creative destruction involves learning by doing and has an observer effect, which is quite different than what has been explored to date in the asset pricing literature. When capital gets allocated to a new opportunity, learning occurs via *experimentation*, which affects expectations about existing assets in the rest of the market. This “perturbs” the system of asset prices and expectations, so learning has feedback effects to the rest of the market and is costly through its effect on risk. Such learning by doing contrasts with the standard learning processes that are typically analyzed, whereby agents receive a time-series of signals for free and update their beliefs with Kalman filtering.¹

In this paper, we characterize the equilibria and asset pricing implications that arise with

¹The incomplete information literature starts with Williams (1977), Detemple (1986), Dothan and Feldman (1986) and Gennotte (1986). A comprehensive survey is provided in Ziegler (2003).

creative destruction. We contrast a setting in which non-cooperative agents compete for market share with a socially optimal benchmark (i.e., a one agent case). As we show, compared to the one-agent case, non-cooperative behavior causes over-experimentation and leads to higher asset prices and risk premia, which resembles a technology bubble. However, these pricing patterns arise from competitive behavior, without requiring information asymmetry, behavioral biases (e.g., optimism or overconfidence), or financial markets frictions.

For decades, most researchers have defined a bubble to be a setting in which the price of an asset exceeds its fundamental value. This has spawned an extensive literature to characterize this distortion and its associated price process. With this definition in mind, bubbles may arise when investors have heterogeneous beliefs and traders have the option to resell the asset to more optimistic agents in the future (Harrison and Kreps, 1978; Morris, 1996; Scheinkman and Xiong, 2003) or when there are short-sale constraints that prevent pessimists in the market from countering the demand from optimists (Miller, 1977; Ofek and Richardson, 2003). Bubbles may also arise when some investors have overconfidence or excessive optimism, and may grow when arbitrageurs are differentially informed about the presence of the bubble and face financial constraints (Abreu and Brunnermeier, 2003). Brunnermeier and Oehmke (2013) provide an excellent review.

More recently, however, Pastor and Veronesi (2009) have shown that the price patterns associated with what we call bubbles do not require a (possibly irrational) wedge between fundamentals and prices. In their model, if a new technology becomes sufficiently promising, its expected cash flows rise, which initially pushes up the stock price. However, because the risk of the technology gradually shifts from idiosyncratic to systematic, the discount rate rises, ultimately leading to a drop in valuation of the asset. Along these lines, DeMarzo, Kaniel, and Kremer (2007) also provide an explanation of bubbles that does not require information asymmetry, behavioral biases, or other frictions. In their model with multiple agents, a “keep up with the Joneses” concern endogenously arises as no agent wants to be left behind. This leads to over-investment that is predictably unprofitable, which they argue is consistent with a bubble.

In our analysis, there is also no wedge between prices and fundamentals, and we focus on the fact that many bubbles arise when investors race to market during eras of technological change (e.g., the Tronics boom, 1959-1962; the Biotech bubble of the 1980’s; the Dotcom era, 1995-2001).² As such, it is natural to consider the interaction between investors as they compete for market share and endogenously affect asset prices and uncertainty in the market. In our model of competition, two agents make choices about experimentation in a Stackelberg

²See Section 1 of Xiong (2013) for a thorough review of historical bubbles.

game.³ Like standard Cournot models and settings in which entry strategy is analyzed, we find that the leader over-experiments compared to a one agent setting in order to prevent the rival from experimenting as well (i.e., entry deterrence). The disturbance effect and displacement risk caused by creative destruction are used by the leader to accomplish this and in equilibrium, the follower never enters. With higher experimentation compared to an economy that is governed by a social planner, both the volatility of future consumption and the uncertainty about the expected growth rate of the economy are magnified. In equilibrium, this results in over-valuation of the risky asset through Jensen’s inequality effects—loosely speaking, the asset has a *convex*, option-like payoff and becomes very valuable with over-experimentation. As time goes on, there is an aftermath in which prices reverse and return to normal. Increased experimentation leads to faster learning and resolution of uncertainty, which causes prices to fall quickly after the run up. This price process resembles what we usually consider to be an asset pricing bubble, but arises only from non-cooperative behavior. Marrying considerations from the industrial organization literature with those in asset pricing accomplishes this, and seems natural given the level of competition that typically occurs during technological change.

While we focus on the rational evolution of bubbles in this paper, our model does accommodate differences in beliefs as well. As we show, if followers are more optimistic about the new technology, they are more inclined to enter and make an investment. Anticipating this, the leader experiments even more in an effort to dissuade entry, which leads to a more acute run up in prices and a steep aftermath. As such, there is an interaction between competitive effects and a (possibly) behavioral bias that can exacerbate a bubble through the channel that we characterize.

The rest of the paper proceeds as follows. Section 2 characterizes optimal experimentation and equilibrium asset prices in an economy governed by a social planner. Section 3 considers strategic experimentation and asset pricing bubbles. Section 4 returns to the one agent setting and characterizes an extension in which the agent may exercise an option to expand or abandon experimentation at every instant in time. Section 5 concludes. All proofs are relegated to the Appendix.

³As will become evident, this setup makes our analysis tractable, but is not necessary to generate our results. Any non-cooperative game in which a Prisoner’s Dilemma gives rise to overinvestment will generate an asset pricing bubble.

2 Socially Optimal Experimentation

Consider a pure exchange economy defined over a continuous-time finite horizon $[0, T]$. In the status quo, one agent consumes the aggregate output

$$\frac{d\delta_t^S}{\delta_t^S} = \bar{f}dt + \sigma dW_t, \quad (1)$$

where the parameters \bar{f} and σ are known.

At time $t = 0$, the agent has the choice to re-allocate existing capital $x_0 \geq 0$ to an *experimental* asset in the market. This reallocation affects the consumption stream in two ways. First, the new asset has an unknown effect on the drift, which becomes $(\bar{f} + \tilde{\beta}x_0)$. The parameter β is unknown and captures both the adverse effect that creative destruction imposes on other assets and a possible benefit in higher future consumption for the agent. More x_0 increases the effect that β has on consumption. We assume that the agent has initial beliefs such that

$$\beta \sim N(\widehat{\beta}_0, \nu_0), \quad (2)$$

where $\widehat{\beta}_0 > 0$, so that the agent starts with an initial prior that investing in the experimental asset is a good idea.

Second, investment in the experimental asset amplifies the magnitude of the diffusion term to $(1 + kx_0)\sigma$, which captures the increased uncertainty that creative destruction introduces into the economy. What we have in mind is that economic risk may increase because the future of existing assets becomes more uncertain. Taken together, for any x_0 , the dynamics for the aggregate output stream in the new economy is⁴

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \beta x_0)dt + (1 + kx_0)\sigma dW_t. \quad (3)$$

The agent chooses x_0 to maximize aggregate utility (i.e., social welfare)

$$U(c) = \mathbb{E} \left[\int_0^T e^{-\rho s} \frac{c_s^{1-\gamma}}{1-\gamma} ds \right]. \quad (4)$$

We assume that $\gamma \geq 1$, which is consistent with empirical evidence (e.g., Friend and Blume

⁴While we take (3) as a starting point, Johnson (2007) provides a discrete-time microfoundation in a production economy, which is consistent with our choice of allowing experimentation to affect the drift. However, we depart in that experimentation affects the uncertainty regarding the interaction between the new technology and existing assets.

(1975); Dreze (1981)).⁵ At $t = 0$, the agent commits to an experimentation level x_0 that remains constant from time 0 to T . As such, the agent chooses how far to open Pandora's box at $t = 0$ and lives with the consequences. We consider a dynamic extension later in Section 4.

If the agent chooses $x_0 = 0$, then the economy remains in the status quo. Once the agent chooses $x_0 > 0$, she learns over time how the experiment impacts the expected growth rate of δ_t . But the experiment comes at a cost: it “disturbs” the process by increasing the magnitude of the diffusion. This implies that there is an observer effect, which is distinct from what is typically modeled in the asset pricing literature. Usually, agents observe signals about the drift of the dividend process for free and update their beliefs with Kalman filtering. Here, we depart from this approach and assume that our agent learns *by doing*, as in Grossman, Kihlstrom, and Mirman (1977). However, unlike Grossman et al. (1977), the cost of experimentation is the added disturbance introduced into the diffusion term.

2.1 Learning and Optimal Experimentation

Before deriving the socially optimal x_0 , let us consider how the agent learns for any given level of experimentation.

Proposition 1 *From the agent's viewpoint, this partially observed economy is equivalent to a perfectly observed economy with consumption process*

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \hat{\beta}_t x_0)dt + \sigma(1 + kx_0)d\widehat{W}_t, \quad (5)$$

where

$$d\hat{\beta}_t = \frac{x_0}{\sigma(1 + kx_0)}\nu_t d\widehat{W}_t, \quad (6)$$

$$d\nu_t = -\frac{x_0^2}{\sigma^2(1 + kx_0)^2}\nu_t^2 dt, \quad (7)$$

and

$$d\widehat{W}_t \equiv dW_t + \frac{x_0(\beta - \hat{\beta}_t)}{\sigma(1 + kx_0)}dt \quad (8)$$

represents the “surprise” component of the change in output.

⁵Friend and Blume (1975) estimate an average coefficient of relative risk aversion well in excess of one and perhaps in excess of two. Dreze (1981) finds even higher values using an analysis of deductibles in insurance contracts. See also Mehra and Prescott (1985, p. 154).

The agent's expectation formation are "extrapolative" (Brennan, 1998) in that she revises her estimates of β in the direction of the output surprises she observes.⁶ We define ν_t as the Bayesian uncertainty about β at time t (i.e., posterior variance). The expression in (7) together with the initial condition ν_0 implies

$$\nu_t = \frac{1}{\frac{x_0^2}{\sigma^2(1+kx_0)^2}t + \frac{1}{\nu_0}}. \quad (9)$$

As Equation (9) shows, the posterior variance starts at ν_0 but then decays to zero as t goes to infinity. One benefit of experimentation is that the agent can learn about the new technology and lower the Bayesian uncertainty. However, through the term $(1+kx_0)$, experimentation also has a negative effect on learning because it disturbs the economy. In fact, consider the limit of the speed of learning in (9) as $x_0 \rightarrow \infty$

$$\frac{1}{k^2\sigma_s^2}. \quad (10)$$

For any $k > 0$, the speed of learning cannot go above (10) in any finite time. Because experimentation disturbs the economy, it indeed applies a brake to learning.

The agent's problem is to choose a level of experimentation that balances between information gains coupled with the chance of a good experiment and higher disturbance to future consumption coupled with the chance of a bad experiment. This tradeoff affects the expected value of future dividends and the future volatility of dividends, which we derive in Appendix A.1.2. There, we show that both $\mathbb{E}_0[\delta_t]$ and $\text{Var}_0[\ln(\delta_t)]$, for any $t > 0$, unambiguously increase with experimentation. In turn, the tradeoff affects the agent's value function. In any equilibrium, the agent's lifetime expected utility of consumption at time t is

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t) = \mathbb{E}_t \left[\int_t^T e^{-\rho s} \frac{\delta_s^{1-\gamma}}{1-\gamma} ds \right] = \frac{1}{1-\gamma} \int_t^T e^{-\rho s} \mathbb{E}_t [\delta_s^{1-\gamma}] ds. \quad (11)$$

In Appendix A.1.3, we derive the agent's value function to be

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t) = \frac{e^{-\rho t} \delta_t^{1-\gamma}}{1-\gamma} F(\widehat{\beta}_t, \nu_t, t). \quad (12)$$

where $F(\widehat{\beta}_t, \nu_t, t)$ is the price-dividend ratio in this economy. This value function unambigu-

⁶The extrapolative nature of learning is different from the "extrapolation bias," a pervasive phenomenon in human judgement and decision making. Extrapolation bias refers to the tendency to overweight recent events when making decisions about the future. In our case, the agent does not overweight recent events, but applies standard Bayesian updating rules. See Hirshleifer, Li, and Yu (2015) for an analysis of extrapolation bias in a production economy.

ously increases with $\mathbb{E}_0[\delta_t]$ and unambiguously decreases with $\text{Var}_0[\ln(\delta_t)]$, for any $t > 0$. This trade-off implies an optimal level of experimentation.

Proposition 2 *At time $t = 0$, the optimal level of experimentation that maximizes expected lifetime utility $J(\delta_0, \widehat{\beta}_0, \nu_0, 0)$ is given by*

$$x_0^* = \frac{(\widehat{\beta}_0 - \gamma k \sigma^2) \mathbb{D}_0}{\gamma k^2 \sigma^2 \mathbb{D}_0 + (\gamma - 1) \nu_0 \mathbb{C}_0}, \quad (13)$$

where \mathbb{D}_0 is the equity duration (the weighted average maturity of cash-flows) and \mathbb{C}_0 is the equity convexity (the weighted average squared maturity of cash-flows).⁷

The optimal level of experimentation in (13) resembles a mean-variance portfolio. The expression (13) is implicit in x_0^* , however, because the duration and convexity of the cash flows that experimentation induces are functions of x_0^* . Experimentation has more impact when the duration of cash-flows is higher (both in terms of higher growth and higher future consumption volatility) because the agent's choice has a longer-lasting impact on the economy. When the margin between the expected benefit of experimentation and the penalty from disturbing the economy is higher, the agent's incentive to experiment increases further with the duration of cash flows. On the other hand, because convexity magnifies the uncertainty of future consumption volatility, it lowers the agent's incentive to experiment.

Inspecting (13), for $\gamma = 1$ a closed form expression exists

$$x_0^* = (\widehat{\beta}_0 - k \sigma^2) / (k^2 \sigma^2). \quad (14)$$

In this case, experimentation is increasing in $\widehat{\beta}_0$ and decreasing in k and σ .

2.2 Implications for Asset Prices

We begin by characterizing the dynamics of the stochastic discount factor, the risk-free rate and the market price of risk in this economy.

Proposition 3 *The stochastic discount factor, $\xi_t \equiv e^{-\rho t} (\delta_t / \delta_0)^{-\gamma}$, follows*

$$\frac{d\xi_t}{\xi_t} = - \left(\rho + \gamma (\bar{f} + \widehat{\beta}_t x_0) - \frac{1}{2} \gamma (\gamma + 1) \sigma^2 (1 + k x_0)^2 \right) dt - \gamma \sigma (1 + k x_0) d\widehat{W}_t. \quad (15)$$

⁷These quantities, defined in Appendix A.1.3, are omitted here for ease of exposition.

The risk-free rate and the market price of risk are therefore given by

$$r_t^f = \rho + \gamma(\bar{f} + \widehat{\beta}_t x_0) - \frac{1}{2}\gamma(\gamma + 1)\sigma^2(1 + kx_0)^2 \quad (16)$$

$$\theta_t = \gamma\sigma(1 + kx_0). \quad (17)$$

As usual, the equilibrium risk-free rate increases with the expected growth rate of consumption and decreases with the volatility of aggregate output. However, by inspection of (16), the level of experimentation amplifies both of these well-known asset-pricing effects. Also, experimentation increases the market price of risk because the agent disturbs the economy when she invests in creative destruction.

The equilibrium price of the risky asset is

$$P_t \equiv \frac{1}{\xi_t} \mathbb{E}_t \left[\int_t^T \xi_s \delta_s ds \right] = e^{\rho t} (1 - \gamma) \delta_t^\gamma J(\delta_t, \widehat{\beta}_t, \nu_t, t), \quad (18)$$

which is negatively related to the value function because $\gamma \geq 1$. This has two implications. First, when the agent chooses an optimal experimentation level to maximize her value function, the equilibrium price reaches a minimum. This implies that any level of experimentation away from the optimum *always* increases the share price. Second, the product of the state price density and consumption, $\xi_s \delta_s$, is *convex* in δ_s . Jensen's inequality implies that the share price must increase with the path of future variance. Because experimentation increases the path of future variance, over-experimentation generates over-valuation of the asset through the convexity of discounted future cash flows.

Defining the function $\kappa(x_0, \widehat{\beta}_t)$ as

$$\kappa(x_0, \widehat{\beta}_t) \equiv (1 - \gamma) \left(\bar{f} + x_0 \widehat{\beta}_t - \gamma \frac{\sigma^2(1 + kx_0)^2}{2} \right) - \rho, \quad (19)$$

the following proposition further characterizes stock prices in this economy.

Proposition 4 *For an initial experimentation level x_0 , the stock price at time t equals*

$$P(\widehat{\beta}_t, \nu_t, t; x_0) = \delta_t \int_t^T \exp \left[\kappa(x_0, \widehat{\beta}_t)(s - t) + \frac{(1 - \gamma)^2}{2} x_0^2 \nu_t (s - t)^2 \right] ds, \quad (20)$$

with expected return and volatility given by

$$\mu_{P,t} = \bar{f} + \widehat{\beta}_t x_0 - \kappa(x_0, \widehat{\beta}_t) - \frac{1}{F(\widehat{\beta}_t, \nu_t, t)} + \gamma(1 - \gamma)x_0^2 \nu_t \mathbb{D}_t \quad (21)$$

$$\sigma_{P,t} = \sigma(1 + kx_0) \left| 1 + (1 - \gamma) \frac{x_0^2 \nu_t}{\sigma^2(1 + kx_0)^2} \mathbb{D}_t \right|. \quad (22)$$

Denoting by $\mu_{P,t}^C$ the cum-dividend expected return on the stock, the equilibrium risk premium in the economy is then

$$\mu_{P,t}^C - r_t = \gamma\sigma^2(1 + kx_0)^2 + \gamma(1 - \gamma)x_0^2 \nu_t \mathbb{D}_t. \quad (23)$$

The stock price valuation (20) highlights an additional impact of experimentation, which arises through the Bayesian uncertainty ν_t . Because in this economy the price-dividend ratio is convex in the expected growth rate, Jensen’s inequality implies a higher valuation in presence of uncertainty (as shown also in Pástor and Veronesi (2003) and Pástor and Veronesi (2006)). In our setting, this effect is further amplified by experimentation.

Experimentation has two effects on the volatility of stock returns. It first increases the volatility by disturbing the economy and therefore amplifying macroeconomic fluctuations. Second, experimentation decreases the volatility of stock returns as long as there is uncertainty about the expected growth of the new technology. This arises from agent’s “extrapolative” expectations.⁸

Figure 1 provides an example that illustrates the effect of experimentation on the asset price, the risk premium, and volatility. The thin vertical lines in each of the panels represent the optimal level of experimentation (Proposition 2). As expected, the share price attains its minimum at this optimal level (left panel). Beyond x_0^* , more experimentation increases both the variance of future consumption and the uncertainty of expected growth. Through the convexity of the present value of asset’s payoffs, this makes the asset a more expensive bet, as discussed above.

The middle panel shows that the risk premium increases with experimentation, and ultimately decreases when experimentation is high enough. The initial increase is driven by the first, “disturbance” term in expression (23), whereas the subsequent decrease is driven by the second, “uncertainty” term. A similar hump-shaped pattern emerges for the volatility

⁸To see this, consider a positive output surprise, which generates a higher expected output growth. The stock price increases because it pays off more consumption into the future. However, the hedging properties of the stock deteriorate, because it generates high consumption when needed less, which dampens the price increase when the investor is more risk-averse than the log-utility investor. This second effect—stronger when the equity duration is higher and when the risk aversion parameter is large—lowers the volatility of asset returns.

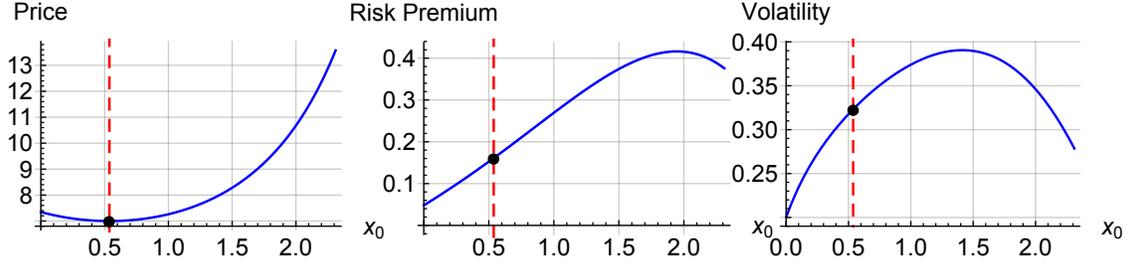


Figure 1: **Experimentation and Asset Prices.** The effect of experimentation on the asset price, the risk premium, and volatility. The vertical lines represent the optimal level of experimentation, and the black dots represent the corresponding asset price, risk premium and volatility. The calibration used is: $\gamma = 1.2$, $\bar{f} = 0.1$, $\hat{\beta}_0 = 0.3$, $\nu_0 = 0.15$, $\sigma = 0.2$, $k = 2$, $\rho = 0.05$, $T = 10$, and $\delta_0 = 1$.

of stock returns, which we illustrate in the right panel of Figure 1.

Experimentation has further implications for the term structure of risk. To see this, in Corollary 5 we decompose the asset into dividend strips (i.e., assets that pay the aggregate consumption only at time $s > t$):

Corollary 5 *The price of a dividend strip with maturity $s > t$ is*

$$P_{t,s} \equiv \frac{1}{\xi_t} \mathbb{E}_t[\xi_s \delta_s] = \delta_t \exp \left[\kappa(x_0, \hat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s-t)^2 \right], \quad (24)$$

whereas the risk premium and the volatility for each maturity are given by

$$\mu_{P,t,s} - r_t = \gamma \sigma^2 (1 + kx_0)^2 + \gamma (1-\gamma) x_0^2 \nu_t (s-t) \quad (25)$$

$$\sigma_{P,t,s} = \sigma (1 + kx_0) \left| 1 + (1-\gamma) \frac{x_0^2 \nu_t}{\sigma^2 (1 + kx_0)^2} (s-t) \right|. \quad (26)$$

The effect of experimentation on the term structure of risk is shown in Figure 2. The red dashed lines show the term structure of risk premia (left panel) and volatilities (right panel) with no experimentation, whereas the blue solid lines show the optimal experimentation case. There are two effects caused by experimentation. First, by amplifying the volatility of the consumption output, more experimentation produces a level increase in risk premia and volatilities at all maturities. Second, uncertainty dampens this increase less in the near future and more in the far future and thus experimentation produces a *downward sloping* term structure of risk premia and volatilities (in Equations (25)-(26), uncertainty amplifies the maturity for each dividend strip). In contrast, a one-agent setup with no experimentation (red dashed lines) produces a flat term structure of risk premia and volatilities.

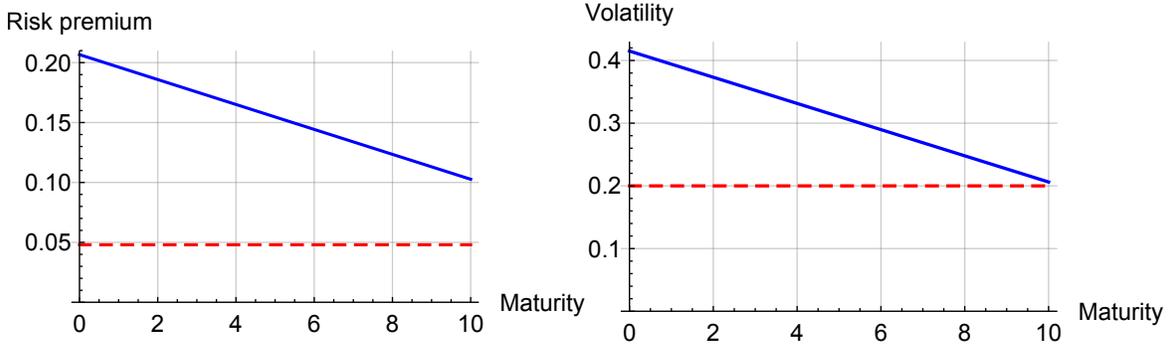


Figure 2: **Experimentation and the Term Structure of Risk.** The term structure of risk premia and volatility with no experimentation (red dashed lines) and with experimentation at the optimal level x_0^* (blue solid lines). Calibrated parameters are provided in Figure 1.

3 Strategic Experimentation and Asset Price Bubbles

3.1 Competition for New Technology

Consider now a strategic setup where there are two players who compete for market share. Assume a Stackelberg game in which a leader (x) chooses first how much to experiment and then a follower (y) responds. Suppose that each receive the following output streams:

$$\frac{d\delta_{x,t}}{\delta_{x,t}} = [\bar{f} + \beta(x_0 - y_0)] dt + \sigma (1 + k_x x_0 + k_y y_0) dW_t \quad (27)$$

$$\frac{d\delta_{y,t}}{\delta_{y,t}} = [\bar{f} + \beta(y_0 - x_0)] dt + \sigma (1 + k_x x_0 + k_y y_0) dW_t. \quad (28)$$

This structure is meant to capture that the two players compete and each has adverse effects on the other. First, each player loses a fraction of growth which is proportional to the experimentation level of the other player. Second, both players disturb the economy through experimentation, with two (possibly distinct) disturbance parameters, k_x and k_y , that appear in the diffusion of both output processes.

The focus of our analysis is the strategic action of the leader. Having the first-mover advantage, the leader is able to influence follower's choice. For simplicity, we assume that player two is a log-utility agent, but we leave the risk aversion coefficient of the leader to be any number greater or equal to one. As shown in Proposition 2, for a log-utility investor optimal experimentation has a closed-form solution. This leads to a simple “best response”

of the follower for a given experimentation level x_0 of the leader:

$$y_0^*(x_0) = \frac{\widehat{\beta}_0 - k_y \sigma^2 (1 + k_x x_0)}{k_y^2 \sigma^2}, \quad (29)$$

which can be substituted into the output process of the leader:

$$\frac{d\delta_{x,t}}{\delta_{x,t}} = \left(\bar{f} - \beta \frac{\widehat{\beta}_0 - k_y \sigma^2}{k_y^2 \sigma^2} + \frac{k_x + k_y}{k_y} \beta x_0 \right) dt + \frac{\widehat{\beta}_0}{k_y \sigma} dW_t. \quad (30)$$

The diffusion of this process does not depend on experimentation (even when $k_x \neq k_y$). As (29) shows, more experimentation from the leader induces less experimentation from the follower. This results in a diffusion term that is now insensitive to experimentation in leader's point of view. Since leader's choice of x_0 has now an impact only on the drift of the output process (30), she can experiment without feeling the negative consequences of disturbing the economy, which are exactly offset by the benefits of driving the follower out of the market.

Given this, the leader finds it optimal to choose an experimentation level at which the follower reaches $y_0^* = 0$ and is driven out of the market. If this level is above the optimal experimentation choice in a one-agent economy—in the “over-experimentation” region—then there is no benefit of experimenting beyond this point, because once the follower is out of the market, the force that offsets the disturbance effect in leader's output process (30) vanishes (negative experimentation for the follower is not possible). For any k_x and k_y , the leader can prevent the follower from experimenting by choosing an x_0 that solves

$$\frac{\widehat{\beta}_0 - k_y \sigma^2 (1 + k_x x_0)}{k_y^2 \sigma^2} = 0. \quad (31)$$

Proposition 6 *Suppose the leader and the follower have log-utility and receive the cash-flow streams defined by (27) and (28). Then, the leader's optimal level of experimentation is*

$$x_0^* = \begin{cases} x_{0,\text{bm}}^* & \text{if } 0 < k_x \leq k_y, \\ \frac{k_x}{k_y} x_{0,\text{bm}}^* + \frac{k_x - k_y}{k_x k_y}, & \text{if } k_y < k_x, \end{cases} \quad (32)$$

where $x_{0,\text{bm}}^*$ is the optimal level of experimentation in the benchmark case of a one-agent economy with log utility:

$$x_{0,\text{bm}}^* \equiv \frac{\widehat{\beta}_0 - k_x \sigma^2}{k_x^2 \sigma^2}. \quad (33)$$

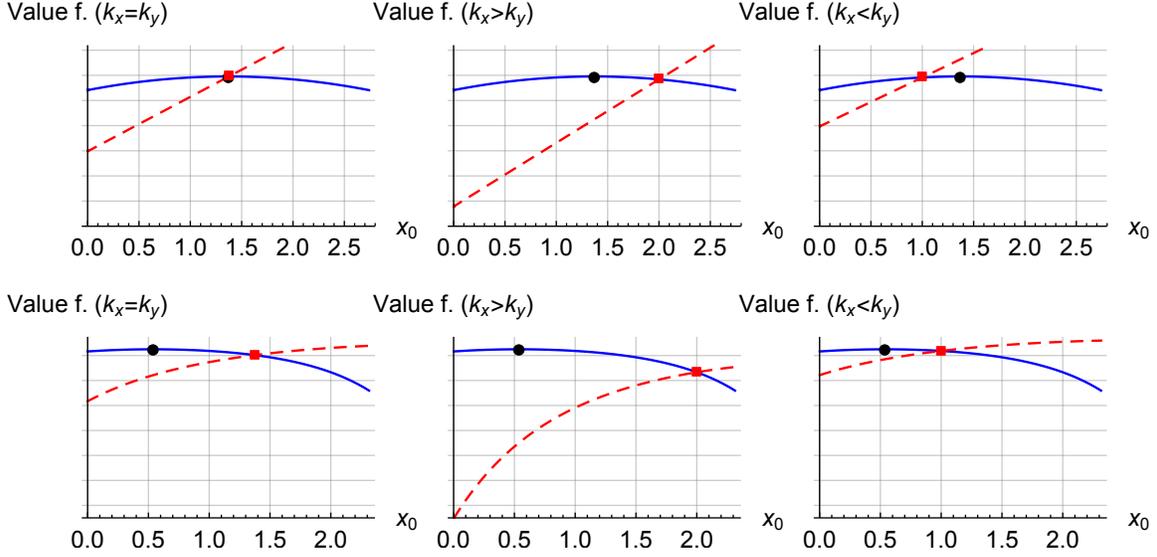


Figure 3: **Strategic Experimentation.** Value functions of the leader in two situations: in a one-agent economy (blue solid line) or in presence of a follower (red dashed line). The black dot is the optimal level of experimentation in the one-agent economy. The intersection of the two lines (red square) is the point where the follower is driven out of the market ($y_0^* = 0$). In the first row, both the leader and the follower have log utility. In the second row, the leader is more risk averse than log, whereas the follower has log utility. The left panel represents the situation $k_x = k_y$, the middle panel is for $k_x > k_y$ and the right panel $k_x < k_y$. Calibrated parameters are provided in Figure 1.

The first row of plots in Figure 3 illustrates three possible situations. The left panel considers the case when both agents have the same disturbance parameter $k_x = k_y$ and the right depicts when the follower is less efficient ($k_x < k_y$). The blue solid line shows the value function of the leader in the benchmark one-agent economy, with the black dot representing the corresponding optimal level of experimentation. The red dashed line draws the utility of the leader in presence of a follower. When $k_x = k_y$, the result of Proposition 6 simplifies exactly to $x_{0,\text{bm}}^*$ and the two lines cross at the benchmark optimal level. This is the required level of experimentation implemented by the leader to preempt entry of the follower. When $k_x < k_y$, experimenting at the benchmark optimal level is more than sufficient to preempt entry of the follower. So, in both cases, experimentation by the leader is the same as in the one-agent economy. The leader's value function in each case is formed by the left segment of the dashed red line coupled with the right segment of the solid blue line, joined together where they intersect.

The middle panel depicts the case when the follower is more efficient and disturbs the economy less than the leader ($k_x > k_y$). In this situation the leader faces the threat of entry

of a follower who will optimally choose to experiment relatively more. This has the potential to hurt leader's future growth and thus she optimally decides to experiment more than in the benchmark case and deter entry of the (*more efficient*) follower. The value function of the leader is formed by the left segment of the dashed red line coupled with the right segment of the solid blue line, crossing at the red square and at the left of the black dot, in the over-experimentation region. The red square is the optimal level of experimentation which allows the leader to maximize utility and preempt entry of the follower.

The discussion so far has analyzed the case when the leader has log utility. But the result of Proposition 6 can be made more general. In the second row of Figure 3 we analyze the case when the leader is more risk-averse than log (while we maintain log-utility for the follower) and illustrate similar patterns, with an important distinction. Although the leader is more risk-averse than log, she chooses the same level of experimentation as if she would be a log-utility investor. Consequently, the leader now over-experiments in all cases, even when $k_x < k_y$.⁹ This arises because the follower, being less risk-averse, experiments relatively more than the leader. The optimal response of the leader is to over-experiment and in this way to preempt entry of the follower. The solution selected by market forces is once again not welfare improving.

Before characterizing how this affects asset prices, we should note that it is also possible to consider differences in beliefs in our model. Indeed, differences in beliefs are often advanced as a theory of bubbles.¹⁰ In the presence of short-sale constraints, asset buyers hold an option to resell the asset to more optimistic agents. This biases the price towards the optimists' valuation, causing over-valuation. This can arise in both static and dynamic settings.¹¹

Our setup can also accommodate this consideration. Suppose that the follower has log utility and the leader is more risk averse than log, and that $k_x = k_y = k$. Suppose also that the follower is more optimistic about the future growth of the new technology:

$$\widehat{\beta}_{x,0} < \widehat{\beta}_{y,0}. \quad (34)$$

In this case, the output process for the leader becomes

$$\frac{d\delta_{x,t}}{\delta_{x,t}} = \left(\bar{f} - \beta \frac{\widehat{\beta}_{y,0} - k\sigma^2}{k^2\sigma^2} + 2\beta x_0 \right) dt + \frac{\widehat{\beta}_{y,0}}{k\sigma} dW_t, \quad (35)$$

⁹If the leader is much more efficient ($k_x \ll k_y$), then she finds optimal to experiment at the benchmark level. More precisely, there is a threshold k_y^* for follower's disturbance parameter above which the leader does not need to over-experiment.

¹⁰Comprehensive surveys are provided by Hong and Stein (2007), Brunnermeier and Oehmke (2013) and Xiong (2013).

¹¹Miller (1977); Harrison and Kreps (1978); Ofek and Richardson (2003); Scheinkman and Xiong (2003).

which again shows that the leader benefits from experimenting without further disturbing the output process. The prior of the follower, $\widehat{\beta}_{y,0}$, which by assumption (34) is higher than the prior of the leader, now appears in the output process (35) further reducing its expected growth and increasing its volatility. Facing these less favorable prospects, the leader optimally responds by over-experimenting:

$$x_0^* = \frac{\widehat{\beta}_{y,0} - k\sigma^2}{k^2\sigma^2} > x_{0,\text{bm}}^*. \quad (36)$$

When the follower holds higher expectations about the promise of the a technology, she has stronger incentives to experiment if she enters the market. The leader again responds by over-experimenting to preempt entry by the follower. The more disparate the two agents' prior beliefs are, the greater the degree to which over-experimentation will occur.

3.2 Asset price bubbles

As described in Section 2, the equilibrium price of the risky asset is negatively related to an agent's value function, which has two implications. First, when the agent acts as a monopolist and chooses an optimal experimentation level to maximize her value function, the equilibrium price reaches a minimum. Any level of experimentation away from the optimum *always* increases the share price. Second, the asset valuation is convex in the value of future dividends and in their expected growth. Jensen's inequality implies that the share price must increase with the path of future variance and with the uncertainty about future growth. Because experimentation increases both the future variance and the growth uncertainty, over-experimentation generates over-valuation of the asset thorough the convexity of discounted future cash flows.

As such, when competition is present and the leading agent preempts entry through over-experimentation, this generates a price pattern that resembles a bubble. This is illustrated in Figure 4, which describes an example when the leader is more risk-averse than log and $k_x > k_y$. The red dotted line in the left panel of Figure 4 represents the new level of experimentation implemented by the leader, and the red square shows the new higher equilibrium price.

Over-experimentation also effects the equity risk-premium and the volatility of asset returns, which are illustrated in the two other panels of Figure 4. Both the risk-premium and the volatility increase to levels higher than the ones obtained in the one agent economy, although the relationship is hump-shaped overall (as observed from Equations (22) and (23)).

So, far over-experimentation can be shown to cause higher asset prices, increased risk

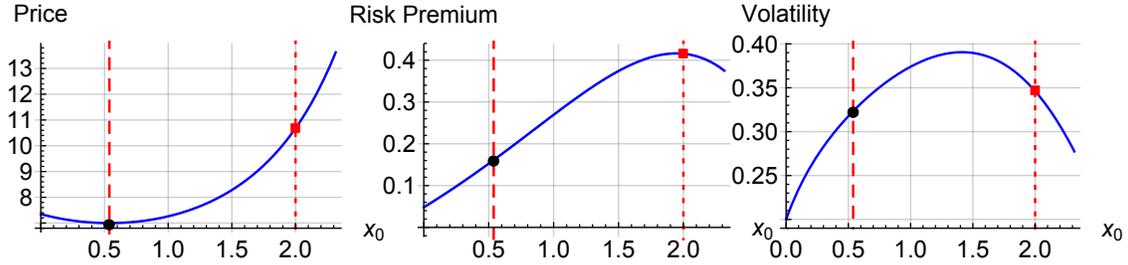


Figure 4: **Strategic Experimentation and Asset Price Bubbles.** The effect of experimentation on the asset price, the risk premium, and volatility. The vertical lines represent the optimal level of experimentation in the one-agent economy (dashed lines) and in the two-agents economy (dotted lines). Calibrated parameters are provided in Figure 1.

premia and higher volatility of asset returns. Now, we can consider the consequences on asset prices in the aftermath of the decision to experiment. We investigate this in Figure 5. There, we suppose that $\hat{\beta}_0 > 0$, but the true β is zero.¹² As soon as experimentation begins, the agent observes the evolution of the output process and updates her estimate of β . But, because learning depends on how much experimentation takes place, asset prices will behave in different ways when the agent experiments optimally or when she over-experiments due to competitive pressures. Figure 5 illustrates the averages of asset prices over 2,500 simulations of two years at daily frequency. The starting points at time zero (the black dot and the red square) are the same points from the left panel in Figure 4 and correspond respectively to optimal and over-experimentation. We emphasize here that the only difference in the two situations is the level of experimentation, while the prior $\hat{\beta}_0$ and the true β are the same in each case. Nevertheless, the average price paths are substantially different: whereas in the optimal experimentation case (blue solid line) the average price path is increasing over the two-year period, the over-experimentation case depicts a strong downward trend. This is the cost that over-experimentation imposes on asset prices, which are inflated initially, but then deflate over several weeks or months. In contrast with the optimal experimentation case, over-experimentation not only brings too much capital into a non-productive technology, but also strongly disturbs the economy by cannibalizing existing assets and creating additional uncertainty.

It is important to mention that the situation depicted by the red dashed line can be identified as a “bubble” only ex post; the agent in our model actually never expects prices to fall. But, because learning by doing reveals the true value of β over time, stock prices adjust endogenously to take into account agent’s new expectations. The pace of this adjustment process depends on the rate at which uncertainty about the new technology is

¹²Technically, β can be any number lower than the prior $\hat{\beta}_0$, without qualitatively affecting these results.

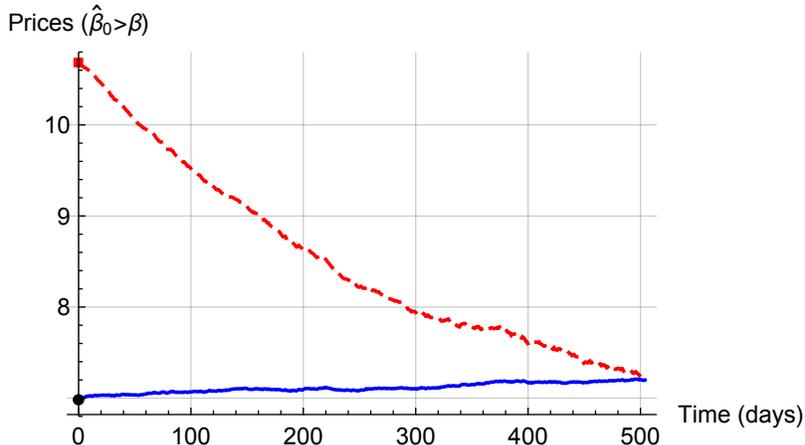


Figure 5: **The Aftermath of Experimentation.** The lines depict average asset prices over 2,500 simulations of two years at daily frequency, when $\hat{\beta}_0 > \beta = 0$. The black dot and the red square are the same points as in the left panel of Figure 4. The blue solid line corresponds to optimal experimentation in a one-agent economy. The red dashed line corresponds to the over-experimentation in a strategic setup. Calibrated parameters are provided in Figure 1.

resolved. More specifically, Eq. (9) shows that the posterior variance decays faster with over-experimentation,¹³ which accelerates the eventual decline in asset prices. Our model therefore predicts that markets characterized by higher competition for new technologies are not only prone to stronger asset price inflation but also to steeper subsequent declines through faster revelation of the promise of a new technology.

Finally, we highlight here the importance of learning for the evolution of bubbles. The aftermath depicted in Figure 1 can only take place through agent’s updating about the true value of β . If no uncertainty (and thus learning) about the true expected growth of the new technology were present, then prices would never fall and we would not observe an aftermath.

4 Dynamic experimentation

Now, we return to the one agent case and consider that the agent can choose the experimentation level x_t at any time t , in order to maximize her expected lifetime utility. Thus, she can alter the level of experimentation dynamically and retains the option to expand or abandon her investment at every instant. The agent’s expected lifetime utility of consumption J

¹³In Eq. (9), the coefficient multiplying time in the denominator can be interpreted as the *speed of learning*. It is straightforward to show that the speed of learning increases with experimentation.

satisfies the following partial differential equation at any time t :

$$0 = \max_x \left[\mathcal{D}J(\delta_t, \widehat{\beta}_t, \nu_t, t) + e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} \right], \quad (37)$$

with boundary condition $J(\delta_T, \widehat{\beta}_T, \nu_T, T) = 0$ and subject to $x_t \geq 0, \forall t$. In equilibrium, the CRRA utility conjecture

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t} \frac{\delta_t^{1-\gamma}}{1-\gamma} F(\widehat{\beta}_t, \nu_t, t) \quad (38)$$

results in a partial differential equation for the price-dividend ratio $F(\cdot)$, which we relegate to Appendix A.3 for sake of brevity. The optimal level of experimentation then follows from the first order condition on x_t .

Proposition 7 *If the problem (37) has an interior maximum, then the optimal level of experimentation at time t solves*

$$x_t^* = \frac{\widehat{\beta}_t - \gamma k \sigma^2}{\gamma k^2 \sigma^2} - \frac{\nu_t}{\gamma k^2 \sigma^2} \left(\frac{F_\beta}{F} - \frac{x_t^* \nu_t}{(\gamma - 1) \sigma^2 (1 + k^2 x_t^*)^3} \frac{F_{\beta\beta} - F_\nu}{F} \right). \quad (39)$$

The solution (39) constitutes an implicit form since the control x appears on the right hand side of the equation. Nevertheless, it highlights two main components of the optimal level of experimentation. The first is a “mean-variance” component which increases when the agent expects a higher growth for the new technology $\widehat{\beta}_t$ and decreases with the risk aversion coefficient and the disturbance parameter k . The second is a “hedging” component, which vanishes when there is no uncertainty about the new technology or in the log-utility case.¹⁴ This term results from agent’s desire to hedge variations in the filter $\widehat{\beta}_t$ but also from agent’s ability to exert control through her experimentation choice on the evolution of both $\widehat{\beta}_t$ and ν_t .

4.1 Asset prices with dynamic experimentation

Proposition 7 shows that in the dynamic case the optimal experimentation level fluctuates as new information becomes available and affects agent’s expectations. This has further impact on asset prices in the economy, which we characterize below.

¹⁴All the partial derivatives in brackets are zero when $\gamma = 1$.

Proposition 8 *The risk-free rate and the market price of risk in an economy with dynamic experimentation are given by*

$$r_t^f = \rho + \gamma(\bar{f} + x_t^* \hat{\beta}_t) - \frac{1}{2} \gamma(\gamma + 1) \sigma^2 (1 + kx_t^*)^2 \quad (40)$$

$$\theta_t = \gamma \sigma (1 + kx_t^*), \quad (41)$$

whereas the aggregate risk premium and the volatility of stock returns are

$$\mu_{S,t}^C - r_t^f = \gamma \sigma^2 (1 + kx_t^*)^2 + \gamma(1 - \gamma)(x_t^*)^2 \nu_t \tilde{\mathbb{D}}_t \quad (42)$$

$$\sigma_{P,t} = \sigma(1 + kx_t^*) \left| 1 + (1 - \gamma) \frac{(x_t^*)^2 \nu_t}{\sigma^2 (1 + kx_t^*)^2} \tilde{\mathbb{D}}_t \right|. \quad (43)$$

Expressions (40)-(43) have a similar structure with the ones in the static case (Section 2.2), with two main differences. First, the market price of risk now fluctuates and increases with the level of experimentation. Second, in (42)-(43), the equity duration from the static case has been adjusted to account for time variation in x_t . We define this adjusted duration in Appendix A.3. As such, the intuition from the static case applies here as well: by increasing the volatility of aggregate consumption, experimentation magnifies both the equity risk premium and the volatility of asset returns; however, too much experimentation can lower risk premia and volatility through the extrapolative expectations channel, which dampens asset price fluctuations. Therefore, as in the static case, both the risk premium and the volatility will feature a hump-shaped pattern in experimentation.

We compare the risk premium and the volatility in the static versus the dynamic case in Figure 6, first as a function of the filter $\hat{\beta}_t$ in the upper panels, then as a function of the uncertainty ν_t in the lower panels. The static experimentation case is depicted with blue solid lines, whereas the dynamic experimentation case is depicted with red dashed lines.

When the filter $\hat{\beta}_0$ is sufficiently low, neither the “static” or the “dynamic” agent decide to experiment. The “dynamic experimenter” has a lower threshold above which she decides to start experimenting (this threshold is around $\hat{\beta}_t = 0.1$ in the graph) and she experiments more than the “static experimenter” for any level of $\hat{\beta}_t$ beyond this threshold. The dynamic experimenter always has the option to stop or decrease later on. In contrast, the static experimenter is more cautious when fixing an initial experimentation level.

Because of this the risk premium and the volatility are generally higher with dynamic experimentation than with static experimentation. Indeed, the option to abandon lowers risk for the dynamic experimenter. But, internalizing this she experiments more, which raises the risk premium and volatility. This is depicted in lower panels of Figure 6. There, the risk premium and the volatility are plotted for different values of the uncertainty ν_t . In

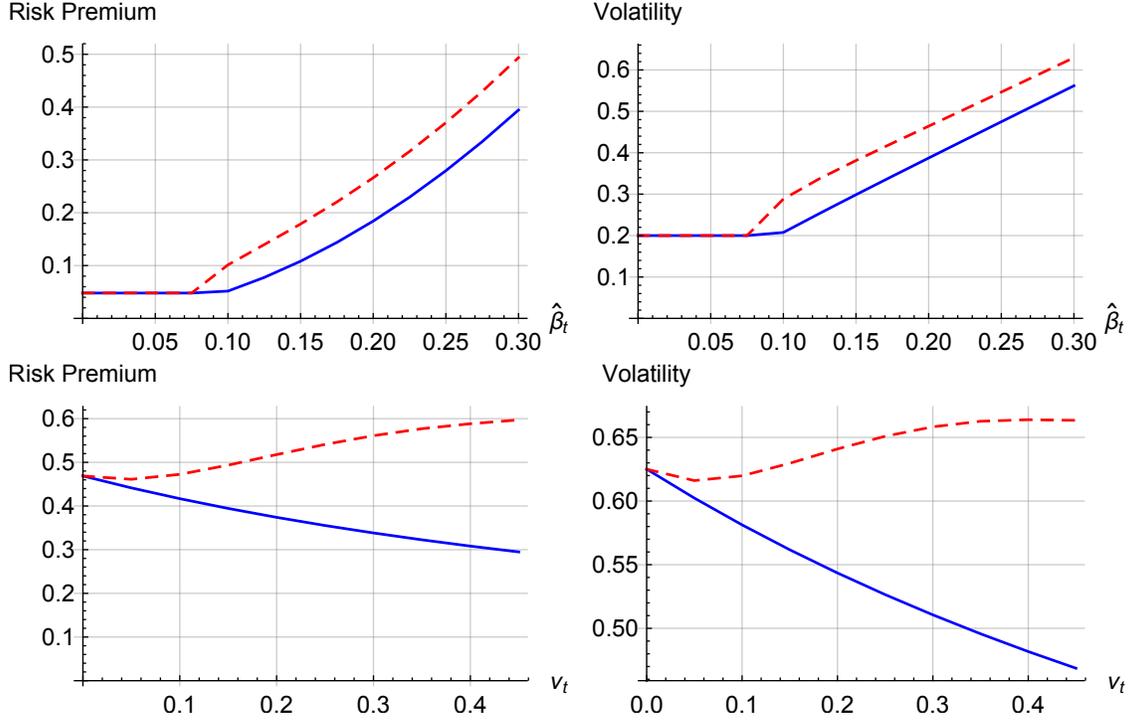


Figure 6: **Asset Pricing with Dynamic Experimentation.** The four panels depict the risk premium and the volatility of asset returns in the static experimentation (blue solid lines) versus the dynamic experimentation (red dashed line) settings. The upper panel plot functions of the prior $\hat{\beta}_t$. The lower panels plot functions of the uncertainty ν_t . Calibrated parameters are provided in Figure 1, except for the maturity, which is fixed at $T = 1$ to improve numerical accuracy.

the static case, the risk premium and the volatility decrease with uncertainty. The “static experimenter” always chooses a lower level of experimentation if uncertainty is high, which decreases the risk premium and the volatility. In the dynamic case, however, both the risk premium and the volatility increase in ν_t —and eventually decrease once the last terms in (42)-(43) dominate. We can understand this by analyzing two extreme cases. First, if uncertainty is negligible, then both the “static” and “dynamic” agent know the parameter β and thus their choices over x_0 are the same (as both panels show). Second, if uncertainty is unusually high, both the “static” and “dynamic” agent give up experimenting. Consequently, only intermediate values of ν_0 are beneficial for the “dynamic experimenter”: because she has the flexibility to stop experimenting later, she is experimenting more aggressively today and thus both the risk premium and the volatility are higher for intermediate values of uncertainty.

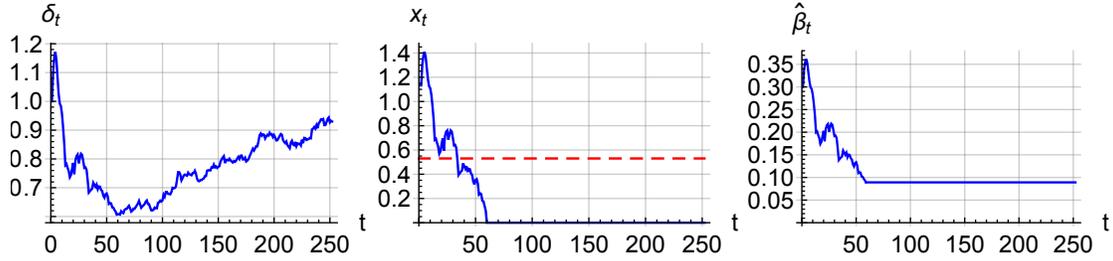


Figure 7: **Experimentation and the Abandonment of a New Technology.** The left panel shows a simulated path of consumption over 252 days. The middle panel shows agent’s optimal experimentation choice, given the observed history of consumption (the red dashed line shows the optimal experimentation level in the static case of Section 2). The right panel shows the filter $\hat{\beta}_t$ which results from agent’s updating. Calibrated parameters are provided in Figure 1, except for the maturity, which is fixed at $T = 1$ to improve numerical accuracy.

4.2 Impact on the real economy

Dynamic experimentation provides the agent with an additional option to abandon the new technology at any point in the future, which results in more aggressive experimentation than in the static case. Although this additional option is an improvement in agent’s set of choices, it can have adverse consequences on the economy. A primary consequence pertains to the volatility of aggregate consumption, which in the dynamic case is given by $\sigma(1 + kx_t)$ and thus behaves stochastically. Experimentation, thus, amplifies the volatility of consumption through the observer effect that it imposes on the economy. This is a particular aspect of our model, in that learning-by-doing transforms the economy *from within*, and stands in contrast with standard models of learning studied in the literature, where learning takes place through observation of signals occurring exogenously and does not have a direct impact on the real economy.

Furthermore, the abandonment option embedded in dynamic experimentation has the power to decide the future of a new technology. This is best illustrated by the example in Figure 7, where we consider a situation in which the agent stops experimenting just because of an unusually bad stream of dividends has occurred. In this example, the agent starts with a positive prior $\hat{\beta}_0 > 0$, which also happens to be equal to the true β . As such, this is a perfectly viable technology that can improve agent’s welfare if adopted. However, lack of perfect knowledge about its productivity and learning by doing leads the agent to conclude after the unusually strong downward trend observed in the left panel that the technology is not productive. In the middle panel, the agent stops experimenting after about 50 days. This also stops the learning process: in the right panel, the estimated $\hat{\beta}_t$ remains constant at a relatively low value once experimentation stops.

5 Conclusion

This paper proposes a financial markets perspective on Schumpeter (1934)'s evolutionary economics ideas, according to which introduction of new technologies disturbs the flow of economic life and forces existing means of production to lose their position within the economy. It is then the task of the financier to decide how much of the new technology the economy will be willing to take.

From the financier's viewpoint, an optimum exist. This optimum balances the gains of economic development associated with new productive technologies against the disturbance imposed on the status quo. The process of reaching such an optimum involves learning by doing (i.e., experimentation), which has an observer effect and creates uncertainty in financial markets. Fierce competition for new markets can force technology leaders to over-experiment, inflating in this way asset prices and generating high risk premia and high volatility. In hindsight, asset prices exhibit familiar boom and bust patterns observed during technological revolutions.

A worthwhile direction for future research would be to analyze the consequences of experimentation when the financier does not fully internalize its consequences on the economy. As we show in this paper, experimentation can lead to market crashes. This can have severe consequences not only for the financier, but for all economic agents, including those who are not necessarily involved in entrepreneurial activity.

A Appendix

A.1 One-Agent Economy

A.1.1 Proof of Proposition 1

Start with the initial process:

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \beta x_0)dt + \sigma(1 + kx_0)dW_t. \quad (44)$$

Apply Ito's lemma on (44), integrate from 0 to t , then write

$$\ln(\delta_t) - \ln(\delta_0) = \left(\bar{f} + \beta x_0 - \frac{1}{2}\sigma^2(1 + kx_0)^2 \right) t + \sigma(1 + kx_0) \int_0^t dW_s^\delta. \quad (45)$$

Assume that the agent only has diffuse prior information on β , i.e. $\nu_0 \rightarrow \infty$. Then, agent's best estimate of β at time t is

$$\hat{\beta}_t = \frac{\ln(\delta_t) - \ln(\delta_0)}{x_0 t} + \frac{1}{x_0} \left(\frac{1}{2}\sigma^2(1 + kx_0)^2 - \bar{f} \right). \quad (46)$$

The mean square error (the Bayesian uncertainty) of the estimate is

$$\nu_t = \mathbb{E} \left[(\hat{\beta}_t - \beta)^2 \right] = \mathbb{E} \left[\left(\frac{\ln(\delta_t) - \ln(\delta_0)}{x_0 t} + \frac{1}{x_0} \left(\frac{1}{2}\sigma^2(1 + kx_0)^2 - \bar{f} \right) - \beta \right)^2 \right] \quad (47)$$

$$= \mathbb{E} \left[\left(\frac{\sigma(1 + kx_0)}{x_0 t} \int_0^t dW_s^\delta \right)^2 \right]. \quad (48)$$

This is the variance of a stochastic process and equals

$$\nu_t = \frac{\sigma^2(1 + kx_0)^2}{x_0^2 t^2} \mathbb{E} \left[\int_0^t ds \right] \quad (49)$$

$$= \frac{\sigma^2(1 + kx_0)^2}{x_0^2 t} \quad (50)$$

This formula holds only for $\nu_0 \rightarrow \infty$. In order to obtain the general formula for a positive and finite ν_0 , we need the dynamics of ν_t . Differentiate this with respect to t to get

$$d\nu_t = -\frac{\sigma^2(1 + kx_0)^2}{x_0^2 t^2} dt \quad (51)$$

$$= -\frac{x_0^2}{\sigma^2(1 + kx_0)^2} \nu_t^2 dt. \quad (52)$$

which is equation (7) in the text. This dynamic equation is valid for any starting value ν_0 , including $\nu_0 \rightarrow \infty$. Solving this partial differential equation with boundary condition $\nu(0) = \nu_0$ yields the

general formula for ν_t :

$$\nu_t = \frac{\sigma^2(1+kx_0)^2}{x_0^2 t + \frac{\sigma^2(1+kx_0)^2}{\nu_0}}. \quad (53)$$

Notice that if $\nu_0 \rightarrow \infty$ then we obtain (50). Also, when $t = 0$ then we get the initial condition $\nu(0) = \nu_0$. Finally, if experimentation is positive ($x_0 > 0$) then uncertainty is decreasing with time.

Finally, apply Ito's lemma on (46) to find $d\widehat{\beta}_t$:

$$d\widehat{\beta}_t = \frac{d \ln(\delta_t)}{x_0 t} - \frac{\ln(\delta_t) - \ln(\delta_0)}{x_0 t^2} dt \quad (54)$$

$$= \frac{d \ln(\delta_t)}{x_0 t} - \frac{\bar{f} + x_0 \widehat{\beta}_t - \frac{1}{2} \sigma^2 (1+kx_0)^2}{x_0 t} dt \quad (55)$$

$$= \frac{x_0(\beta - \widehat{\beta}_t) dt + \sigma(1+kx_0) dW_t}{x_0 t} \quad (56)$$

$$= \frac{\sigma(1+kx_0)}{x_0 t} d\widehat{W}_t = \frac{x_0}{\sigma(1+kx_0)} \nu_t d\widehat{W}_t \quad (57)$$

which is equation (6) in the text. □

A.1.2 Conditional moments of future consumption

Proposition 9 *The conditional mean and variance of $\ln(\delta_s)$ are given by*

$$\mathbb{E}_t [\ln(\delta_s)] = \ln(\delta_t) + \left(\bar{f} + x_0 \widehat{\beta}_t - \frac{\sigma^2(1+kx_0)^2}{2} \right) (s-t), \quad (58)$$

$$\text{Var}_t [\ln(\delta_s)] = \sigma^2(1+kx_0)^2 (s-t) + x_0^2 \nu_t (s-t)^2. \quad (59)$$

Therefore

$$\mathbb{E}_t [\delta_s] = \delta_t \exp \left[\left(\bar{f} + x_0 \widehat{\beta}_t \right) (s-t) + \frac{x_0^2 \nu_t}{2} (s-t)^2 \right]. \quad (60)$$

Proof Equation (58) results from application of Ito's lemma on $\ln(\delta_t)$, with the process of δ_t provided in (5). For equation (59), consider the following random vector:

$$Y_s \equiv \begin{bmatrix} y_s \\ \widehat{\beta}_s \end{bmatrix} \equiv \begin{bmatrix} \ln(\delta_s) - \left(\bar{f} - \frac{\sigma^2(1+kx_0)^2}{2} \right) (s-t) \\ \widehat{\beta}_s \end{bmatrix} \quad (61)$$

It follows that

$$dY_s = \begin{bmatrix} dy_s \\ d\widehat{\beta}_s \end{bmatrix} = \begin{bmatrix} 0 & x_0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} y_s \\ \widehat{\beta}_s \end{bmatrix} + \begin{bmatrix} \sigma(1+kx_0) \\ \frac{x_0 \nu_s}{\sigma(1+kx_0)} \end{bmatrix} d\widehat{W}_t \quad (62)$$

Let the conditional variance of Y_s , $\text{Var}_t[Y_s]$, be defined as:¹⁵

$$\text{Var}_t[Y_s] \equiv \begin{bmatrix} z_{11} & z_{12} \\ z_{12} & z_{22} \end{bmatrix} \quad (63)$$

We have

$$\begin{bmatrix} z'_{11} & z'_{12} \\ z'_{12} & z'_{22} \end{bmatrix} = \begin{bmatrix} 0 & x_0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} z_{11} & z_{12} \\ z_{12} & z_{22} \end{bmatrix} + \begin{bmatrix} z_{11} & z_{12} \\ z_{12} & z_{22} \end{bmatrix} \begin{bmatrix} 0 & 0 \\ x_0 & 0 \end{bmatrix} + \begin{bmatrix} \sigma(1+kx_0) \\ \frac{x_0\nu_s}{\sigma(1+kx_0)} \end{bmatrix} \begin{bmatrix} \sigma(1+kx_0) & \frac{x_0\nu_s}{\sigma(1+kx_0)} \end{bmatrix} \quad (64)$$

$$= \begin{bmatrix} 2x_0z_{12} + \sigma^2(1+kx_0)^2 & x_0z_{22} + x_0\nu_s \\ x_0z_{22} + x_0\nu_s & \frac{x_0^2\nu_s^2}{\sigma^2(1+kx_0)^2} \end{bmatrix} \quad (65)$$

The solutions for the z 's are found by integration of (65). Start by integrating the (2,2) term, then the (1,2) term, and finally the (1,1) term.

$$z_{22}(s) = \int_t^s \frac{x_0^2\nu_u^2}{\sigma^2(1+kx_0)^2} du \quad (66)$$

$$= \nu_t - \nu_s \quad (67)$$

We get this because

$$d\nu_t = -\frac{x_0^2}{\sigma^2(1+kx_0)^2} \nu_t^2 dt \quad (68)$$

$$\nu_s - \nu_t = -\int_t^s \frac{x_0^2\nu_u^2}{\sigma^2(1+kx_0)^2} du. \quad (69)$$

Now, move to z_{12} :

$$z_{12}(s) = x_0 \int_t^s [z_{22}(u) + \nu_u] du = x_0\nu_t(s-t) \quad (70)$$

Finally,

$$z_{11}(s) = \int_t^s [2x_0z_{12}(u) + \sigma^2(1+kx_0)^2] du \quad (71)$$

$$= \sigma^2(1+kx_0)^2(s-t) + \int_t^s [2x_0^2\nu_t(u-t)] du \quad (72)$$

$$= \sigma^2(1+kx_0)^2(s-t) + x_0^2\nu_t(s-t)^2 \quad (73)$$

¹⁵This is a direct approach to compute the conditional variance $\text{Var}_t[Y_s]$ in one step. See Ziegler [2003], page 180, and Bryson and Ho [1975], Section 11.4. An alternative approach is to proceed in several steps: compute first $\text{Var}_t[\hat{\beta}_s]$, then $\text{Cov}_t[y_s, \hat{\beta}_s]$, and finally $\text{Var}_t[y_s]$. Note that there is a change in the drift from δ_s to y_s , but this has no effect on the conditional variance of $\ln(\delta_s)$, i.e. $\text{Var}_t[\delta_s] = \text{Var}_t[y_s]$.

which is equation (59).

Equations (58) and (59) are now proved. Then, we can write

$$\mathbb{E}_t[\delta_s] = \mathbb{E}_t \left[e^{\ln(\delta_s)} \right] \quad (74)$$

$$= \exp \left[\mathbb{E}_t[\ln(\delta_s)] + \frac{1}{2} \text{Var}_t[\ln(\delta_s)] \right] \quad (75)$$

$$= \delta_t \exp \left[\left(\bar{f} + x_0 \hat{\beta}_t \right) (s - t) + \frac{x_0^2 \nu_t}{2} (s - t)^2 \right] \quad (76)$$

which proves equation (60). \square

A.1.3 Value Function

Proposition 10 *The value function unambiguously increases with expected future dividends $\mathbb{E}_t[\delta_s]$ and unambiguously decreases with the future variance $\text{Var}_t[\ln(\delta_s)]$, for any $s \geq t$. Furthermore, the value function can be written*

$$J(\delta_t, \hat{\beta}_t, \nu_t, t) = \frac{e^{-\rho t} \delta_t^{1-\gamma}}{1-\gamma} F(\hat{\beta}_t, \nu_t, t), \quad (77)$$

where

$$F(\hat{\beta}_t, \nu_t, t) \equiv \int_t^T \exp \left[\kappa(x_0, \hat{\beta}_t)(s - t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s - t)^2 \right] ds \quad (78)$$

and

$$\kappa(x_0, \hat{\beta}_t) \equiv (1-\gamma) \left(\bar{f} + x_0 \hat{\beta}_t - \gamma \frac{\sigma^2(1+kx_0)^2}{2} \right) - \rho. \quad (79)$$

Proof In any equilibrium, the agent consumes the entire output δ_t and thus her lifetime expected utility of consumption can be computed as

$$J(\delta_t, \hat{\beta}_t, \nu_t, t) = \mathbb{E}_t \left[\int_t^T e^{-\rho s} \frac{\delta_s^{1-\gamma}}{1-\gamma} ds \right] = \frac{1}{1-\gamma} \int_t^T e^{-\rho s} \mathbb{E}_t[\delta_s^{1-\gamma}] ds, \quad (80)$$

where the second equality results from application of Fubini's theorem.

The expectation in (80) can be further expanded by using the property that for a normally distributed random variable $y = \ln(x)$, $\mathbb{E}[x^\alpha] = \exp(\alpha \mathbb{E}[y] + \alpha^2 \text{Var}[y]/2)$:

$$J(\delta_t, \hat{\beta}_t, \nu_t, t) = \frac{1}{1-\gamma} \int_t^T e^{-\rho s} \exp \left[(1-\gamma) \mathbb{E}_t[\ln \delta_s] + \frac{(1-\gamma)^2}{2} \text{Var}_t[\ln(\delta_s)] \right] ds, \quad (81)$$

$$= \frac{1}{1-\gamma} \int_t^T e^{-\rho s} \mathbb{E}_t[\delta_s]^{1-\gamma} \exp \left[-\frac{\gamma(1-\gamma)}{2} \text{Var}_t[\ln(\delta_s)] \right] ds, \quad (82)$$

where the second equality results from

$$\mathbb{E}_t[\ln(\delta_s)] = \ln(\mathbb{E}_t[\delta_s]) - \frac{1}{2} \text{Var}_t[\ln(\delta_s)]. \quad (83)$$

Replacing the results of Proposition 9 in Equation (83) yields the result.

From equation (82), it is straightforward to show that $J(\delta_t, \widehat{\beta}_t, t)$ increases in $\mathbb{E}_t[\delta_s]$ and decreases in $\text{Var}_t[\ln(\delta_s)]$, for any value of the risk aversion parameter $\gamma > 1$. Then, Equation (77) results from using Proposition 9 and equations (58) and (59) for $\mathbb{E}_t[\ln(\delta_s)]$ and $\text{Var}_t[\ln(\delta_s)]$ in (81). \square

A.1.4 Proof of Proposition 2

Write the first order condition

$$0 = \frac{\partial J(\delta_0, \widehat{\beta}_0, \nu_0, 0)}{\partial x_0} = \frac{\delta_0^{1-\gamma}}{1-\gamma} \int_0^T \left[(1-\gamma) \left(\widehat{\beta}_0 - \gamma k \sigma^2 (1 + k x_0) \right) t + (1-\gamma)^2 x_0 \nu_0 t^2 \right] \times \exp \left[\kappa(x_0, \widehat{\beta}_0) t + \frac{(1-\gamma)^2}{2} x_0^2 \nu_0 t^2 \right] dt, \quad (84)$$

and define the functions G and H as

$$G(\widehat{\beta}_t, \nu_t, t) \equiv \int_t^T (s-t) \exp \left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s-t)^2 \right] ds \quad (85)$$

$$H(\widehat{\beta}_t, \nu_t, t) \equiv \int_t^T (s-t)^2 \exp \left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s-t)^2 \right] ds. \quad (86)$$

After replacing G and H and dividing by F , the first order condition becomes

$$\widehat{\beta}_0 \mathbb{D}_0 - \gamma k \sigma^2 (1 + k x_0) \mathbb{D}_0 + (1-\gamma) x_0 \nu_0 \mathbb{C}_0 = 0, \quad (87)$$

where the two quantities \mathbb{D}_t and \mathbb{C}_t represent the equity duration and the equity convexity:

$$\mathbb{D}_t \equiv \frac{G(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)} \quad (88)$$

$$\mathbb{C}_t \equiv \frac{H(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)}. \quad (89)$$

The value \mathbb{D}_t represents the weighted average maturity and is therefore interpreted as equity duration, whereas the value \mathbb{C}_t represents the weighted average squared maturity and is interpreted as equity convexity. Both values are positive. Solving in (87) for x_0 yields (13). \square

A.1.5 Proof of Proposition 3

Consider the following standard results in asset pricing (Duffie, 2010). Suppose that the dynamics of consumption can be written as

$$\frac{dc_t}{c_t} = \mu_{ct}dt + \sigma_{ct}^\top dz_t \quad (90)$$

where μ_{ct} is the expected relative growth rate of consumption at σ_{ct} is the vector of sensitivities of consumption growth to the exogenous shocks of the economy. The variance of consumption growth is given by $\|\sigma_{ct}\|^2$.

Assuming time-additive expected utility we can define a stochastic discount factor from the optimal consumption plan of any individual as

$$\xi_t = e^{-\rho t} \frac{u'(c_t)}{u'(c_0)} \quad (91)$$

If a one agent exists, the equation holds for aggregate consumption.

Application of Ito's lemma on the function $\xi(c, t)$ implies

$$\frac{\partial \xi}{\partial t} = -\rho \xi_t \quad (92)$$

$$\frac{\partial \xi}{\partial c} = - \left(-c_t \frac{u''(c_t)}{u'(c_t)} \right) \frac{\xi_t}{c_t} \quad (93)$$

$$\frac{\partial^2 \xi}{\partial c^2} = \left(c_t^2 \frac{u'''(c_t)}{u'(c_t)} \right) \frac{\xi_t}{c_t^2} \quad (94)$$

The first term in (93) is the relative risk aversion of the individual. The first term in (94) is related to the relative prudence (which dictates precautionary savings); this term is positive when the absolute risk aversion is decreasing in the level of consumption (very plausible assumption, valid for the CRRA case). Let's denote these two terms by $\gamma(c_t)$ and $\eta(c_t)$ respectively. Then, the dynamics of the stochastic discount factor can be expressed as

$$\frac{d\xi_t}{\xi_t} = - \left(\rho + \gamma(c_t)\mu_{ct} - \frac{1}{2}\eta(c_t)\|\sigma_{ct}\|^2 \right) dt - \gamma(c_t)\sigma_{ct}^\top dz_t \quad (95)$$

The continuously compounded risk-free rate is (the negative of the drift of the stochastic discount factor):

$$r_t^f = \rho + \gamma(c_t)\mu_{ct} - \frac{1}{2}\eta(c_t)\|\sigma_{ct}\|^2 \quad (96)$$

The market price of risk process is (the negative of the diffusion of the stochastic discount

factor):

$$\theta_t = \gamma(c_t)\sigma_{ct} \quad (97)$$

which implies that the excess expected return (cum-dividend) on asset i over the instant following time t can be written as

$$\mu_{it} + \delta_{it} - r_t^f = \gamma(c_t)\sigma_{it}^\top \sigma_{ct} \quad (98)$$

$$= \gamma(c_t)\rho_{ict} \|\sigma_{it}\| \|\sigma_{ct}\| \quad (99)$$

where δ_{it} is the dividend yield of the asset. In our case, we have $\delta_{it} = \delta_t/S_t$. Furthermore, with CRRA utility, we have $\gamma(c_t) = \gamma$ and $\eta(c_t) = \gamma(\gamma + 1)$. Replace these in the equations above to obtain the results of Proposition 3. \square

A.1.6 Proof of Proposition 4 and Corollary 5

Recall that the state variables in this economy evolve according to (5)-(7) and that $P_t = \delta_t F(\widehat{\beta}_t, \nu_t, t)$. We first obtain F_t , F_β , F_ν , and $F_{\beta\beta}$, where $F(\widehat{\beta}_t, \nu_t, t)$ is defined in (78) and $\kappa(x_0, \widehat{\beta}_t)$ is defined in (19). Leibniz's integral rule yields

$$F_t = \int_t^T \left[-\kappa(x_0, \widehat{\beta}_t) - (1 - \gamma)^2 x_0^2 \nu_t (s - t) \right] \exp \left[\kappa(x_0, \widehat{\beta}_t)(s - t) + \frac{(1 - \gamma)^2}{2} x_0^2 \nu_t (s - t)^2 \right] ds - 1 \quad (100)$$

$$= -\kappa(x_0, \widehat{\beta}_t) F(\widehat{\beta}_t, \nu_t, t) - (1 - \gamma)^2 x_0^2 \nu_t G(\widehat{\beta}_t, \nu_t, t) - 1 \quad (101)$$

with $G(\widehat{\beta}_t, \nu_t, t)$ defined in (85). Then,

$$F_\beta = (1 - \gamma)x_0 G(\widehat{\beta}_t, \nu_t, t) \quad (102)$$

$$F_{\beta\beta} = (1 - \gamma)^2 x_0^2 H(\widehat{\beta}_t, \nu_t, t) \quad (103)$$

$$F_\nu = \frac{(1 - \gamma)^2}{2} x_0^2 H(\widehat{\beta}_t, \nu_t, t) \quad (104)$$

where $H(\widehat{\beta}_t, \nu_t, t)$ is defined in (86). Apply Ito's formula to $P_t = \delta_t F(\widehat{\beta}_t, \nu_t, t)$:

$$dP_t = \delta_t F \frac{d\delta_t}{\delta_t} + \delta_t F_\beta d\widehat{\beta}_t + \delta_t F_\nu d\nu_t + \delta_t F_t dt + \frac{1}{2} \left[\delta_t F_{\beta\beta} (d\widehat{\beta}_t)^2 + 2F_\beta (d\delta_t)(d\widehat{\beta}_t) \right] \quad (105)$$

to obtain

$$\frac{dP_t}{P_t} = \mu_{P,t} dt + \sigma_{P,t} d\widehat{W}_t \quad (106)$$

where

$$\mu_{P,t} = \bar{f} + \widehat{\beta}_t x_0 - \kappa(x_0, \widehat{\beta}_t) - \frac{1}{F(\widehat{\beta}_t, \nu_t, t)} + \gamma(1 - \gamma)x_0^2 \nu_t \frac{G(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)} \quad (107)$$

$$\sigma_{P,t} = \sigma(1 + kx_0) \left(1 + (1 - \gamma) \frac{x_0^2 \nu_t}{\sigma^2(1 + kx_0)^2} \frac{G(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)} \right). \quad (108)$$

To obtain the risk premium as in (23), write the cum-dividend expected return on the stock:

$$\mu_{P,t}^C = \rho + \gamma(\bar{f} + \widehat{\beta}_t x_0) - \frac{1}{2}\gamma(\gamma - 1)\sigma^2(1 + kx_0)^2 - \gamma(\gamma - 1)x_0^2 \nu_t \mathbb{D}_t, \quad (109)$$

and use the risk-free rate from Equation (16) (alternatively, one can multiply the market price of risk, $\theta_t = \gamma\sigma(1 + kx_0)$ (Proposition 3, Equation 17) with the diffusion of stock returns). \square

A.2 Strategic Experimentation

A.2.1 Proof of Proposition 6

The two consumption streams (27)-(28) can be re-written:

$$\frac{d\delta_{x,t}}{\delta_{x,t}} = (\tilde{f}_x + \beta x_0)dt + \tilde{\sigma}_x(1 + \tilde{k}_x x_0)dW_t \quad (110)$$

$$\frac{d\delta_{y,t}}{\delta_{y,t}} = (\tilde{f}_y + \beta y_0)dt + \tilde{\sigma}_y(1 + \tilde{k}_y y_0)dW_t, \quad (111)$$

with parameters

$$\tilde{f}_x \equiv \bar{f} - \beta y_0, \quad \tilde{\sigma}_x \equiv \sigma(1 + k_y y_0), \quad \tilde{k}_x \equiv \frac{k_x}{1 + k_y y_0}, \quad (112)$$

$$\tilde{f}_y \equiv \bar{f} - \beta x_0, \quad \tilde{\sigma}_y \equiv \sigma(1 + k_x x_0), \quad \tilde{k}_y \equiv \frac{k_y}{1 + k_x x_0}. \quad (113)$$

The optimal experimentation level of the follower (Equation (29) in the text) can then be obtained directly from Proposition 2:

$$y_0^*(x_0) = \frac{\widehat{\beta}_0 - k_y \sigma^2(1 + k_x x_0)}{k_y^2 \sigma^2}, \quad (114)$$

which can be replaced in (110) to obtain Equation (30) in the text. Proposition 6 then follows from solving Equation (31). \square

A.3 Dynamic Experimentation

A.3.1 Proof of Proposition 7

The dynamics of consumption with experimentation at time t now depend on x_t :

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \widehat{\beta}_t x_t) dt + \sigma(1 + kx_t) d\widehat{W}_t, \quad (115)$$

with

$$d\widehat{\beta}_t = \frac{x_t}{\sigma(1 + kx_t)} \nu_t d\widehat{W}_t \quad (116)$$

$$d\nu_t = -\frac{x_t^2}{\sigma^2(1 + kx_t)^2} \nu_t^2 dt. \quad (117)$$

Note that now the agent can choose the experimentation level x_t at any time t , in order to maximize her expected lifetime utility. Thus, the agent's expected lifetime utility of consumption J satisfies the partial differential equation

$$0 = \max_x \left[\mathcal{D}J(\delta_t, \widehat{\beta}_t, \nu_t, t) + e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} \right] \quad (118)$$

with boundary condition $J(\delta_T, \widehat{\beta}_T, \nu_T, T) = 0$ and subject to

$$x_t \geq 0, \forall t \quad (119)$$

In equilibrium consumption equals total output and therefore

$$0 = \max_x \frac{e^{-\rho t} \delta_t^{1-\gamma}}{1-\gamma} + J_t + \delta_t(\bar{f} + \widehat{\beta}_t x) J_\delta + \frac{x^2 \nu_t^2}{2\sigma^2(1 + kx)^2} (J_{\beta\beta} - 2J_\nu) + \frac{1}{2} \delta_t^2 \sigma^2 (1 + kx)^2 J_{\delta\delta} + \delta_t \nu_t x J_{\delta\beta} \quad (120)$$

With CRRA utility, we make the usual assumption

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t} \frac{\rho_t^{1-\gamma}}{1-\gamma} F(\widehat{\beta}_t, \nu_t, t), \quad (121)$$

and thus the PDE (120) becomes

$$0 = \max_x F_t + \kappa(x, \widehat{\beta}_t) F + \frac{x^2 \nu_t^2}{2\sigma^2(1 + kx)^2} (F_{\beta\beta} - 2F_\nu) + (1 - \gamma) x \nu_t F_\beta + 1 \quad (122)$$

with boundary condition $F(\widehat{\beta}_T, \nu_T, T) = 0$ and $\kappa(x, \widehat{\beta}_t)$ defined as in (19). The first order condition

for x is

$$0 = \kappa'F + (1 - \gamma)\nu_t F_\beta + \frac{x\nu_t^2}{(1 + kx)^3\sigma^2}(F_{\beta\beta} - 2F_\nu) \quad (123)$$

$$= (\gamma - 1) \left(k(1 + kx)\gamma\sigma^2 - \widehat{\beta}_t \right) F + (1 - \gamma)\nu_t F_\beta + \frac{x\nu_t^2}{(1 + kx)^3\sigma^2}(F_{\beta\beta} - 2F_\nu) \quad (124)$$

This is a quartic equation in x , which we solve numerically. Re-arranging this equation yields Equation (39) in the text. \square

A.3.2 Proof of Proposition 8

The stochastic discount factor follows

$$\frac{d\xi_t}{\xi_t} = - \left(\rho + \gamma(\bar{f} + x_t^*\widehat{\beta}_t) - \frac{1}{2}\gamma(\gamma + 1)\sigma^2(1 + kx_t^*)^2 \right) dt - \gamma\sigma(1 + kx_t^*)d\widehat{W}_t^\delta \quad (125)$$

and thus the risk-free rate and the market price of risk from (40)-(41) follow.

The stock price at time t is $S_t = \delta_t F(\widehat{\beta}_t, \nu_t, t)$. The major change in this case with respect to the static case is that the dynamics of all state variables depend on the optimal level of experimentation at time t , x_t^* . The dynamics of the stock price can be written

$$\frac{dS_t}{S_t} = \left[\bar{f} + x_t^*\widehat{\beta}_t - \kappa(x_t^*, \widehat{\beta}_t) - \frac{1}{F} + \gamma x_t^*\nu_t \frac{F_\beta}{F} \right] dt + \sigma(1 + kx_t^*) \left[1 + \frac{x_t^*\nu_t}{\sigma^2(1 + kx_t^*)^2} \frac{F_\beta}{F} \right] d\widehat{W}_t, \quad (126)$$

which gives the volatility of stock returns. The risk premium is then given by

$$\mu_{S,t}^C - r_t^f = \gamma\sigma^2(1 + kx_t^*)^2 \left[1 + \frac{x_t^*\nu_t}{\sigma^2(1 + kx_t^*)^2} \frac{F_\beta}{F} \right] \quad (127)$$

The volatility and the risk premium depend on F_β/F . We know that

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t} \frac{\rho_t^{1-\gamma}}{1-\gamma} F(\widehat{\beta}_t, \nu_t, t), \quad (128)$$

and thus, at the optimal experimentation level x_t^* ,

$$F(\widehat{\beta}_t, \nu_t, t) = \mathbb{E}_t \left[\int_t^T \frac{e^{-\rho(s-t)} \delta_s^{1-\gamma}}{\delta_t^{1-\gamma}} ds \right] = \int_t^T \frac{e^{-\rho(s-t)} \mathbb{E}_t[\delta_s^{1-\gamma}]}{\delta_t^{1-\gamma}} ds \quad (129)$$

and

$$\mathbb{E}_t [\delta_s^{1-\gamma}] = e^{(1-\gamma)\mathbb{E}_t[\ln \delta_s] + \frac{(1-\gamma)^2}{2} \text{Var}_t[\ln \delta_s]} \quad (130)$$

Only $\mathbb{E}_t [\ln \delta_s]$ depends on $\widehat{\beta}_t$:

$$\mathbb{E}_t [\ln \delta_s] = \ln \delta_t + \mathbb{E}_t \left[\int_t^s \left(\bar{f} + \widehat{\beta}_\tau x_\tau^* - \frac{1}{2} \sigma^2 (1 + k x_\tau^*)^2 \right) d\tau \right] \quad (131)$$

and

$$\frac{\partial \mathbb{E}_t [\ln \delta_s]}{\partial \widehat{\beta}_t} = \mathbb{E}_t \left[\int_t^s x_\tau^* d\tau \right] \quad (132)$$

Therefore

$$\frac{F_\beta}{F} = \frac{(1 - \gamma) \int_t^T \mathbb{E}_t \left[\int_t^s x_\tau^* d\tau \right] \frac{e^{-\rho(s-t)\mathbb{E}_t[\delta_s^{1-\gamma}]}}{\delta_t^{1-\gamma}} ds}{\int_t^T \frac{e^{-\rho(s-t)\mathbb{E}_t[\delta_s^{1-\gamma}]}}{\delta_t^{1-\gamma}} ds} \quad (133)$$

$$= (1 - \gamma) x_t^* \frac{\int_t^T \mathbb{E}_t \left[\int_t^s \frac{x_\tau^*}{x_t^*} d\tau \right] \frac{e^{-\rho(s-t)\mathbb{E}_t[\delta_s^{1-\gamma}]}}{\delta_t^{1-\gamma}} ds}{\int_t^T \frac{e^{-\rho(s-t)\mathbb{E}_t[\delta_s^{1-\gamma}]}}{\delta_t^{1-\gamma}} ds} \quad (134)$$

$$= (1 - \gamma) x_t^* \widetilde{\mathbb{D}}_t, \quad (135)$$

where $\widetilde{\mathbb{D}}_t$ is a weighted average of discounted cash-flows, which are adjusted at each maturity τ by the term x_τ^*/x_t^* . $\widetilde{\mathbb{D}}_t$ can therefore be interpreted as an equity duration adjusted for time-variation in x_t . The stock return volatility (43) and the risk premium (42) follow by replacing (135) in Equations (126)-(127). \square

References

- Abreu, D. and M. K. Brunnermeier (2003). Bubbles and crashes. *Econometrica* 71(1), 173–204.
- Brennan, M. J. (1998). The role of learning in dynamic portfolio decisions. *European Finance Review* 1(3), 295–306.
- Brunnermeier, M. K. and M. Oehmke (2013). *Bubbles, Financial Crises, and Systemic Risk*, Volume 2 of *Handbook of the Economics of Finance*, Chapter 0, pp. 1221–1288. Elsevier.
- DeMarzo, P., R. Kaniel, and I. Kremer (2007). Technological innovation and real investment booms and busts. *Journal of Financial Economics* 85(3), 735–754.
- Detemple, J. (1986). Asset pricing in a production economy with incomplete information. *The Journal of Finance* 41(2), pp. 383–391.
- Dothan, M. U. and D. Feldman (1986, June). Equilibrium interest rates and multiperiod bonds in a partially observable economy. *Journal of Finance* 41(2), 369–82.
- Dreze, J. H. (1981). Inferring risk tolerance from deductibles in insurance contracts. *Geneva Papers on Risk and Insurance*, 48–52.
- Duffie, D. (2010). *Dynamic asset pricing theory*. Princeton University Press.
- Friend, I. and M. E. Blume (1975). The demand for risky assets. *The American Economic Review*, 900–922.
- Gârleanu, N., L. Kogan, and S. Panageas (2012). Displacement risk and asset returns. *Journal of Financial Economics* 105(3), 491–510.
- Gennotte, G. (1986, July). Optimal portfolio choice under incomplete information. *Journal of Finance* 41(3), 733–46.
- Grammig, J. and S. Jank (2015). Creative destruction and asset prices. *Journal of Financial and Quantitative Analysis*.
- Grossman, S. J., R. E. Kihlstrom, and L. J. Mirman (1977). A bayesian approach to the production of information and learning by doing. *The Review of Economic Studies*, 533–547.
- Harrison, J. M. and D. M. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 323–336.
- Hirshleifer, D., J. Li, and J. Yu (2015). Asset pricing in production economies with extrapolative expectations. *Journal of Monetary Economics* 76, 87–106.

- Hong, H. and J. C. Stein (2007). Disagreement and the stock market. *The Journal of Economic Perspectives* 21(2), 109–128.
- Kung, H. and L. Schmid (2015). Innovation, growth, and asset prices. *The Journal of Finance* 70(3), 1001–1037.
- Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of monetary Economics* 15(2), 145–161.
- Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of finance* 32(4), 1151–1168.
- Morris, S. (1996). Speculative investor behavior and learning. *The Quarterly Journal of Economics*, 1111–1133.
- Ofek, E. and M. Richardson (2003). Dotcom mania: The rise and fall of internet stock prices. *The Journal of Finance* 58(3), 1113–1137.
- Pástor, L. and P. Veronesi (2003). Stock valuation and learning about profitability. *The Journal of Finance* 58(5), 1749–1790.
- Pástor, L. and P. Veronesi (2006). Was there a nasdaq bubble in the late 1990s? *Journal of Financial Economics* 81(1), 61–100.
- Pastor, L. and P. Veronesi (2009). Technological revolutions and stock prices. *American Economic Review* 99(4), 1451–1483.
- Scheinkman, J. A. and W. Xiong (2003). Overconfidence and speculative bubbles. *Journal of political Economy* 111(6), 1183–1220.
- Schumpeter, J. A. (1934). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*, Volume 55. Transaction publishers.
- Williams, J. T. (1977). Capital asset prices with heterogeneous beliefs. *Journal of Financial Economics* 5(2), 219–239.
- Xiong, W. (2013). Bubbles, crises, and heterogeneous beliefs. *Handbook on Systemic Risk*, 663–713.
- Ziegler, A. C. (2003). *Incomplete Information and Heterogeneous Beliefs in Continuous-time Finance*. Springer-Verlag Berlin Heidelberg.