Asset Pricing in the Quest for the New El Dorado*

Daniel Andrei†
Bruce I. Carlin‡

May 23, 2017

ABSTRACT

Creative destruction not only involves bringing new technology to market, it imposes higher risk on the future of existing assets. We characterize the asset pricing implications of creative destruction when investors compete for market share. Compared to the social optimum, the quest for oligopoly rents leads to over-investment in uncertain projects, spikes in asset prices and risk premia, and an aftermath in which prices fall steeply as uncertainty resolves. These pricing patterns resemble a bubble ex post, but arise solely from competitive behavior and do not require information asymmetry, behavioral biases, or financial frictions. Our analysis yields novel empirical predictions and we discuss how financial innovation might be used to predict bubbles ex ante.

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

†Anderson School of Management, University of California, Los Angeles, 110 Westwood Plaza, Los Angeles, CA 90095, www.danielandrei.info, 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 (2017)*

1 Introduction

According to Schumpeter (1934), invention and entrepreneurship are two distinct activities. And it is the entrepreneur who is tasked with bringing advancements to market. Naturally, the incentive to innovate depends on the rents available for extraction and creative destruction is often associated with wealth transfer. Much of the existing literature that has followed has further developed and tested Schumpeter's hypotheses, with a focus on competitive industry analysis (e.g., Williamson, 1965; Reinganum, 1983; Aghion, Bloom, Blundell, Griffith, and Howitt, 2005)[1] and quantifying whether creative destruction indeed increases social welfare (e.g., Witt, 1996; Aghion, Akcigit, Bergeaud, Blundell, and Hémous, 2015; Komlos, 2014).[2]

Yet, there is an unsung hero in this story, *the financier*, who has been largely neglected, 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 activity drives the evolution of asset prices in the market and in turn the ability to bear risk. This is the focus of this paper.

Creative destruction and entrepreneurship are different from simply adding new assets to a well-diversified portfolio because innovation 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

[1]See also Swan (1970), Loury (1979), Reinganum (1985), Aghion and Howitt (1992), Boldrin and Levine (2008), Aghion and Griffith (2008).

[2]See also Carree and Thurik (2005), Bampoky, Prieger, Blanco, and Liu (2016), Acemoglu and Restrepo (2017), and Akcigit, Grigsby, and Nicholas (2017).

capabilities, it is not surprising that there appears to be a risk premium associated with it (Grammig and Jank, 2015).

Moreover, 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* (e.g., Arrow, 1962; Grossman, Kihlstrom, and Mirman, 1977; Aghion, Bolton, Harris, and Jullien, 1991; Rob, 1991), 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 update their beliefs from a time-series of signals that they receive for free.[3]

In this paper, we characterize the equilibria and asset pricing implications that arise with creative destruction. We study an exchange economy in which investment in creative destruction has two effects on the consumption stream: it changes the growth rate of consumption—hopefully for the better—and it amplifies the magnitude of the diffusion. The former captures the fact that new technologies have uncertain benefits for economic growth; the latter captures the fact that new technologies have unintended consequences and make the future of existing assets more uncertain. We contrast a socially optimal benchmark with a setting in which non-cooperative investors compete for a share of the rents created by the new technology. We model the strategic interaction as an aggregative game (i.e., Cournot competition) in which competitors make simultaneous choices about how much to experiment with the new technology and the payoffs are a function of each investor's claim to an endogenous dividend stream.[4]

We find that when competitors fight for market share, this leads to overinvestment compared to what is socially optimal, and that aggregate experimentation grows with the number of investors. In turn, with socially excessive experimentation, both the volatility of future consumption and the uncertainty about the expected growth rate of the economy are magnified. In equilibrium, this results in a spike in the price-

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

[4] In the aggregative game, the payoff to each player is a function of the sum of their own choice and that of all competitors. This not only affects each player's consumption share, but affects the nature of the consumption stream they get to enjoy.

dividend ratio because of Jensen's inequality effects—loosely speaking, the asset has a *convex*, option-like payoff and becomes more valuable with over-experimentation. As uncertainty resolves over time, there is an aftermath in which prices reverse and return to normal. Higher experimentation leads to faster learning and resolution of uncertainty, which causes prices to fall quickly after the run up.

As such, competition among rent-seeking agents in markets for new technologies generates price patterns that resemble a "bubble". However, the price paths generated by creative destruction can be identified as a bubble only ex post. Ex ante, agents do not expect prices to fall, but learning by doing decreases uncertainty over time and stock prices adjust endogenously. Bayesian uncertainty (i.e., posterior variance) decays faster with over-experimentation, which accelerates the eventual decline in asset prices. Therefore, our model predicts that markets characterized by fierce competition for new technologies are not only prone to asset price inflation but also to steeper subsequent declines through learning by doing and faster resolution of uncertainty.

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 the valuation of the asset. We share their view that irational behavior is not a prerequisite for bubbles.

Our model is distinct, however, because over-valuation of the asset is generated by increased competition and over-investment in new technologies, and risk is systematic at all times. Moreover, it is the learning by doing feature of our model, coupled with excessive experimentation, that leads to steep subsequent declines in asset prices as uncertainty resolves over time.

DeMarzo, Kaniel, and Kremer (2007) also provide an explanation of bubbles that does not require a wedge between prices and fundamentals. 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. Such over-investment arises in our paper as well. But the focus of our paper is on the asset pricing implications of socially excessive experimentation. We show that competition magnifies uncertainty in the market and leads to over-valuation of asset prices, followed by an aftermath characterized by gradual resolution of uncertainty and price reversals. It is likely that a "keep up with the Joneses" concern would magnify the competitive forces in our model, and thus we see the two papers as complementary.

The pricing patterns that we generate arise from rational competitive behavior, without requiring information asymmetry, behavioral biases (e.g., optimism or overconfidence), or financial markets frictions. 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. Inded, 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).[5] However, while we provide a rational explanation for the evolution of bubbles, we do not take the position that irrational exuberance does not contribute to price fluctuations. Quite the opposite. It would be easy to envision how irrationality might exacerbate the boom-bust price patterns that we demonstrate.

Our model generates a unique prediction on how financial innovation might be used to detect bubbles. As Greenwood et al. (2017) document, sharp price increases predict a substantially heightened probability of an ensuing crash and attributes of the price run-up such as price volatility and asset turnover can help forecast the eventual aftermath. In our model, high price-dividend ratios may arise not only as a result of

[5]See Section 1 of Xiong (2013) for a thorough review of historical bubbles.

agents' expectations about future growth, but also as a result of uncertainty. But, it is not possible to infer from the price-dividend ratio alone both the expected growth and the degree of uncertainty. However, in our analysis expected growth enters in the price-dividend valuation as a linear function of the maturity of the dividend stream, while parameter uncertainty enters as a quadratic function. This suggests that if a market for dividend strips were to exist, as proposed by Brennan (1998b), bubbles might be detectable by plotting log values for a horizon of dividends. As we show in the paper, without competitive behavior, this plot is almost linear, whereas with competition and over-experimentation, it is non-linear.

Last, a testable prediction of our model is that, in markets characterized by intense creative destruction or by heightened uncertainty about the expected growth of new technologies, investors demand a higher premium per quantity of risk. Consistent with this prediction, recent empirical research by Grammig and Jank (2015) show that indeed invention activity creates risk which has to be compensated, and that it can account for the size and value premia observed in financial markets.

The rest of the paper proceeds as follows. Section 2 poses the model, characterizes learning and optimal experimentation, and contrasts the equilibrium behavior of competitive agents to what is socially optimal. Section 3 characterizes the asset pricing implications of creative destruction. Section 4 returns to a one agent setting in which the investor maximizes social welfare 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.

2 Experimentation and Learning

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

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

where the parameters $\bar{f}$ and σ are known. The economy is populated by an *active* investor who consumes a share θ of the output and a *passive* investor who consumes fraction $1 - \theta$. For now, we fix θ exogenously so that the active investor cannot modify her consumption share. Later, we endogenize θ, which will reflect rent seeking behavior and strategic considerations among competitive active investors.

At time $t = 0$, the active investor has the choice to re-allocate existing capital $x_0 \geq 0$ to an *experimental* asset in the market (e.g., a new technology). This new technology is sufficiently important to affect the entire economy. It affects the aggregate consumption stream in two ways. First, the new asset has an unknown effect on the drift, which becomes $(\bar{f} + \beta x_0)$. 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. Thus, innovation can both fuel economic growth (Aghion and Howitt, 1992; Romer, 1990) and be destructive (e.g., Acemoglu and Restrepo, 2017), and β is similar to a Schumpeter "creativity ratio" as defined in Komlos (2014). We assume that the active investor has initial beliefs such that

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

where $\widehat{\beta}_0 > 0$, so that the active investor 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 economic risk that creative destruction introduces into the economy (Kung and Schmid, 2015). Taken together, for any x_0, the dynamics for the aggregate output stream in the new economy is

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

In Appendix A.1.1, we provide a microfoundation for (3) in an economy with two differentiated goods. Here, we take (3) as a starting point for our analysis.[6]

The agents have preferences over lifetime consumption

$$U(c_i) = \mathbb{E}\left[\int_0^T e^{-\rho t}\frac{c_{i,t}^{1-\gamma}}{1-\gamma}dt\right], \tag{4}$$

where $c_{a,t} = \theta\delta_t$ for the active investor and $c_{p,t} = (1-\theta)\delta_t$ for the passive investor. Both agents have the same coefficient of risk aversion, $\gamma \geq 1$. This ensures that

[6] 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.

the agents' expected utility is higher under complete information than under incomplete information.[7] Furthermore, this is consistent with empirical evidence[8], and is frequently used in analyses like these (e.g., Pastor and Veronesi, 2009).

Only the active investor can choose x_0. At $t = 0$, the active investor commits to an experimentation level x_0 that remains constant from time 0 to T. As such, the active investor chooses how far to open Pandora's box at $t = 0$ and then both investors live with the consequences.[9] Because θ is fixed and investors are otherwise identical, the choice of the active investor is the socially optimal choice (it simultaneously maximizes the lifetime utility of both investors—see Proposition 2).

If the active investor chooses $x_0 = 0$, then the economy remains in the status quo. Once the active investor chooses $x_0 > 0$, both investors observe the total output δ_t and learn over time how the new asset impacts future expected growth. 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 with incomplete information. Usually, agents update their beliefs by observing signals about the drift of the dividend process for free. Instead, in our case learning occurs *by doing*, as in Arrow (1962), Grossman et al. (1977), and Rob (1991). However, unlike these papers, the cost of experimentation is the added disturbance introduced into the diffusion term through the parameter $k > 0$ and the added uncertainty about the true expected growth of the experimental asset.

Both investors observe the aggregate output stream δ_t, whose changes are informative about the unknown parameter β. Since the choice of x_0 is common knowledge, both investors form expectations under the same probability measure.

Proposition 1 ***(Learning)*** *From investors' viewpoint, this partially observed econ-*

[7] More precisely, for any time-additive utility function which is increasing and concave in current consumption, better information will increase expected utility whenever the second derivative of the utility with respect to the natural logarithm of consumption is negative. In the case of power utility, this condition is satisfied when $\gamma > 1$. See Chapter 2 in Ziegler (2003) for a discussion.

[8] Friend 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).

[9] In Section 4, we consider a dynamic extension in which the active investor chooses the experimentation level at every instant in time.

omy is equivalent to a perfectly observed economy with consumption process

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

where

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

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

and $d\widehat{W}_t \equiv dW_t + \frac{x_0(\beta - \widehat{\beta}_t)}{\sigma(1+kx_0)}dt$ *represents the "surprise" component of the change in total consumption.*

Investors revise their estimate of β in the direction of the output surprises they observe (Brennan, 1998a). 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 a deterministic path for ν_t:

$$\nu_t = \frac{1}{\frac{x_0^2}{\sigma^2(1+kx_0)^2}t + \frac{1}{\nu_0}}. \tag{8}$$

The posterior variance starts at ν_0 but then decays to zero as t goes to infinity. One benefit of experimentation is that investors can learn about the new technology and lower the future Bayesian uncertainty. Moreover, uncertainty decreases faster when x_0 is high. However, when $k > 0$, experimentation also has a negative effect on learning because it disturbs the economy. In fact, consider the limit of the speed of learning (the term multiplying time t in the denominator of (8)) as $x_0 \to \infty$:

$$\frac{1}{k^2\sigma^2}. \tag{9}$$

For any $k > 0$, the speed of learning cannot go above (9) in any finite time. Because experimentation disturbs the economy, it indeed applies a brake to learning.

2.1 Socially Optimal Experimentation

The active investor'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 (i.e., $\mathbb{E}_0[\delta_t]$) and the future volatility of dividends (i.e., $\mathrm{Var}_0[\ln \delta_t]$), which we derive in Appendix A.1.3. There, we show that both $\mathbb{E}_0[\delta_t]$ and $\mathrm{Var}_0[\ln \delta_t]$, for any $t > 0$, unambiguously increase with experimentation, consistent with the idea that creative destruction increases expected growth but also makes the future more uncertain.

The above tradeoff affects the active investor's lifetime expected utility of consumption at time t, which is defined as

$$J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) = \mathbb{E}_t\left[\int_t^T e^{-\rho s}\frac{(\theta\delta_s)^{1-\gamma}}{1-\gamma}ds\right] = \frac{\theta^{1-\gamma}}{1-\gamma}\int_t^T e^{-\rho s}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]ds. \tag{10}$$

In Appendix A.1.4, we show that

$$J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) = \frac{e^{-\rho t}(\theta\delta_t)^{1-\gamma}}{1-\gamma}F(\widehat{\beta}_t, \nu_t, t), \tag{11}$$

where $F(\widehat{\beta}_t, \nu_t, t)$ is the price-dividend ratio in this economy (to be defined and fully characterized in Proposition 5). We further show that the value function unambiguously increases with $\mathbb{E}_0[\delta_t]$ and unambiguously decreases with $\mathrm{Var}_0[\ln \delta_t]$, for any $t > 0$. This tradeoff implies a socially optimal level of experimentation.

Proposition 2 ***(Socially Optimal Experimentation)*** *At $t = 0$, the active investor chooses an optimal level of experimentation that maximizes social welfare. This level is*

$$x_S^* = \begin{cases} \frac{(\widehat{\beta}_0 - \gamma k\sigma^2)\mathbb{D}_0}{\gamma k^2\sigma^2\mathbb{D}_0 + (\gamma-1)\nu_0\mathbb{C}_0}, & \text{if } \widehat{\beta}_0 - \gamma k\sigma^2 > 0 \\ 0, & \text{otherwise,} \end{cases} \tag{12}$$

where $\mathbb{D}_0$ is the equity duration (i.e., the weighted average maturity of discounted cash-flows) and $\mathbb{C}_0$ is the equity convexity (i.e., the weighted average squared maturity of discounted cash-flows).[10]

[10]These positive quantities, defined in Appendix A.1.5, are omitted here for ease of exposition.

The optimal level of experimentation in (12) resembles a mean-variance portfolio. The expression (12) is implicit in x_0^*, however, because the duration and convexity of the cash flows 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 active investor'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 active investor's incentive to experiment increases further with the duration of cash flows. On the other hand, uncertainty lowers the active investor's incentive to experiment. This effect is stronger when the convexity of future cash flows is high.

Going forward, we assume that the socially optimal level of experimentation x_S^* is strictly positive, i.e., the condition $\widehat{\beta}_0 - \gamma k \sigma^2 > 0$ is satisfied.

2.2 Experimentation with Rent Seeking Behavior

Schumpeterian innovation is typically associated with a wealth transfer from the owners of existing assets to people who innovate, and arises most commonly when (potential) monopoly rents are high. Investors who engage in creative destruction often tradeoff between the wealth gains they enjoy from new knowledge and the losses they endure because the existing assets they own will be rendered obsolete. This mirrors Schumpeter's idea that creative destruction, "incessantly revolutionizes the economic structure *from within*, incessantly destroying the old one, incessantly creating a new one" (Schumpeter, 1942, p. 83).[11] In what follows, we analyze the amount of entrepreneurship that arises based on these considerations.

Suppose that $N \geq 1$ active investors compete for market share and simultaneously choose experimentation levels $x_i \geq 0$, for $i \in \{1, ..., N\}$ at $t = 0$ to maximize (4). Each active investor i's consumption share depends not only on her experimentation decision but also on the aggregate level of experimentation in the economy

$$\theta_i(x_i, \tilde{x}_C) = \frac{1}{N} e^{-\varphi + x_i(w - \tilde{x}_C)}, \tag{13}$$

where $\varphi > 0$ and $\tilde{x}_C \equiv \sum_{i=1}^N x_i$ is the total level of experimentation in the economy. The passive investor receives $c_{p,t} = (1 - \theta_C)\delta_t$, where $\theta_C \equiv \sum_{i=1}^N \theta_i(x_i, \tilde{x}_C)$.

[11] See also Aghion and Howitt (1992), where this tradeoff is made explicit by showing that more future research discourages current research by threatening to make it obsolete.

By construction, the fraction of the dividend stream claimed by agent i is increasing in wx_i, which is meant to capture the increase in the share of aggregate wealth that agent i can claim (monopoly rents) from the new technology.[12] Going forward, we assume that

$$w > \frac{2(\widehat{\beta}_0 - k\sigma^2)}{k^2\sigma^2} > 0, \tag{14}$$

so that this incentive is sufficiently large. Indeed, we are interested in the eras with technology booms, in which investors do in fact race to the "new El Dorado." However, the wealth share is also decreasing in the product $x_i\tilde{x}_C$, which captures how much wealth is lost to competition and also because agent i's existing assets are cannibalized or rendered obsolete (Reinganum, 1983; Aghion and Howitt, 1992). As such, this is an aggregative game that resembles Cournot competition: the market share for each investor is growing in their own experimentation, but is decreasing in the sum.

Based on this, the dynamics of consumption are driven by the aggregate level of experimentation:

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \widehat{\beta}_t\tilde{x}_C)dt + \sigma(1 + k\tilde{x}_C)d\widehat{W}_t. \tag{15}$$

Each agent experiments less when their competitors experiment and disturb the economy, which increases the output volatility that everyone faces. Competition also lowers the share of consumption that each agent earns by experimenting. Last, experimentation by competitors increases the overall uncertainty in the economy. Given this, each agent chooses an optimal level of experimentation, taking into account the simultaneous choices of the other players.

Proposition 3 ***(Cournot Experimentation)*** *There exists a unique Nash equilibrium in which the aggregate level of experimentation under competition among N*

[12] Innovation and top income inequality in the US and other developed countries tend to follow a parallel evolution. According to Aghion et al. (2015), 11 out of the 50 wealthiest individuals across US states in 2015 *"are listed as inventors in a US patent and many more manage or own firms that patent."*

active investors is

$$\tilde{x}_C^* = \frac{\widehat{\beta}_0 \mathbb{D}_0 - \gamma k \sigma^2 \mathbb{D}_0 + w}{\gamma k^2 \sigma^2 \mathbb{D}_0 + (\gamma - 1)\nu_0 \mathbb{C}_0 + \frac{N+1}{N}}, \tag{16}$$

where each active agent experiments at a symmetric level $x_i^* = \tilde{x}_C^*/N$*. The quantity* $\tilde{x}_C^*$ *is strictly increasing in* N *and* $\tilde{x}_C^* > x_S^*$ $\forall N$*.*

Proposition 3 also shows that rent seeking increases experimentation in the economy, and that competition exacerbates this. As competition intensifies, the aggregate level of experimentation moves further away from the socially optimal level x_S^*, reaching a maximum when $N \to \infty$.

Such creative destruction and competition for rents causes the lifetime utility from the future consumption stream to deteriorate and lowers the welfare for the passive investor,

$$U(c_i) = \frac{(1-\theta_C)^{1-\gamma}}{1-\gamma} \mathbb{E}\left[\int_0^T e^{-\rho t} \frac{\delta_t^{1-\gamma}}{1-\gamma} dt\right]. \tag{17}$$

This is because rent seeking by active investors not only reduces the passive investors' consumption share, it exposes them to too much risk and lowers the expected utility which they derive from the future consumption stream. Further, higher competition exacerbates this because experimentation rises and hurts residual claimants in the economy.

We illustrate the effect of rent seeking behavior and creative destruction on the aggregate level of experimentation in Figure 1. The parameters that we choose for this numerical example (provided in the caption of the Figure) are economically plausible and will be the same when we explore the implications for asset prices. The plot depicts three equilibria (**M**onopoly, **D**uopoly, and **P**erfect **C**ompetition). These equilibria are the fixed-point solutions of Eq. (16) for three different values of N. We add the socially optimal level of experimentation (point **S**) for comparison. Confirming the results of Proposition 3, there is higher experimentation as competition rises in the market.

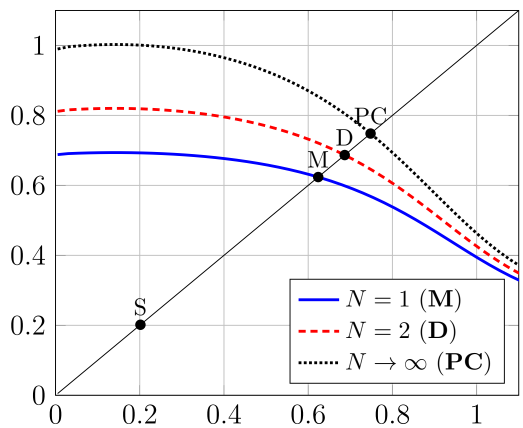

Figure 1: **Experimentation with Rent Seeking Behavior.** The fixed point solution for Eq. (16), for three different values of N: $N = 1$ (monopoly, **M**), $N = 2$ (duopoly, **D**), and $N \to \infty$ (perfect competition, **PC**). The point **S** represents the socially optimal level of experimentation (Proposition 2). The calibration used is: $\gamma = 2$, $\bar{f} = 0.03$, $\widehat{\beta}_0 = 0.03$, $\nu_0 = 0.03^2$, $\sigma = 0.05$, $k = 2$, $\rho = 0.03$, $T = 100$, $\delta_0 = 1$, $k = 3$, $e^{-\varphi} = 0.1\%$, and $w = 2$.

3 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 for a generic level of experimentation x_0.

Proposition 4 ***(Stochastic Discount Factor)*** *The stochastic discount factor, defined as $\xi_t \equiv e^{-\rho t}(\delta_t/\delta_0)^{-\gamma}$, follows*

$$\frac{d\xi_t}{\xi_t} = -\left[\rho + \gamma(\bar{f} + \widehat{\beta}_t x_0) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kx_0)^2\right] dt - \gamma\sigma(1+kx_0)d\widehat{W}_t. \tag{18}$$

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

$$r_t^f = \rho + \gamma(\bar{f} + \widehat{\beta}_t x_0) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kx_0)^2 \tag{19}$$

$$\theta_t = \gamma\sigma(1+kx_0). \tag{20}$$

The equilibrium risk-free rate increases with the expected growth rate of consumption and decreases with the volatility of aggregate consumption. The level of experimentation amplifies both of these well-known asset pricing effects. Furthermore, experimentation increases the market price of risk. This arises because the process of creative destruction disturbs the economy.

Proposition 5 ***(Asset Prices)*** *For any experimentation level x_0, the equilibrium price-dividend ratio at time t, $F(\widehat{\beta}_t, \nu_t, t; x_0)$, is*

$$\frac{P(\widehat{\beta}_t, \nu_t, t; x_0)}{\delta_t} = \int_t^T \exp\left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s-t)^2\right] ds, \tag{21}$$

where the function $\kappa(x_0, \widehat{\beta}_t)$ is defined as

$$\kappa(x_0, \widehat{\beta}_t) \equiv -\rho - (\gamma - 1)(\bar{f} + x_0 \widehat{\beta}_t) + \frac{\gamma(\gamma-1)}{2}\sigma^2(1+kx_0)^2. \tag{22}$$

The equilibrium stock market volatility is given by

$$|\sigma_{P,t}| = \sigma(1+kx_0)\left|1 + (1-\gamma)\frac{x_0^2 \nu_t}{\sigma^2(1+kx_0)^2}\mathbb{D}_t\right|, \tag{23}$$

and the equilibrium risk premium in the economy is

$$RP_t = \gamma\sigma^2(1+kx_0)^2 + \gamma(1-\gamma)x_0^2\nu_t\mathbb{D}_t. \tag{24}$$

A key result of Proposition 5 is that the equilibrium price-dividend ratio increases with the uncertainty about β. This effect arises because, as uncertainty about the true expected growth of the new technology increases, the expected value of future income streams rises. To see this, suppose that β were known. Then, an application of Leibniz' integral rule confirms that the price-dividend ratio is convex in β:

$$\frac{\partial^2 F(\beta, t; x_0)}{\partial \beta^2} = \int_t^T x_0^2(\gamma-1)^2(s-t)^2 e^{\kappa(x_0,\beta)(s-t)} ds > 0. \tag{25}$$

Since β is a random variable, Jensen's inequality implies that that the price-dividend ratio under incomplete information must be *greater* than the one under complete information in the presence of uncertainty about growth rates (and not about levels).[13]

More important, in our setup experimentation amplifies this over-valuation effect in two ways. First, higher experimentation brings more uncertainty to the economy. Second, experimentation amplifies the convexity of the price-dividend ratio (as shown in Eq. 25), which causes over-valuation of the risky asset.

[13] See also Pástor and Veronesi (2003) and Pástor and Veronesi (2006).

Experimentation has two effects on the volatility of stock returns. First, it increases volatility by disturbing the economy and 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 investors' learning and can be understood as follows. Consider a positive output surprise. Through learning, this 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. This second effect dominates when investors are more risk averse than log-utility, lowering the volatility of asset returns.

The risk premium admits a similar interpretation. It increases due to the disturbance effect of experimentation, but decreases due to the uncertainty about the new technology.

Figure 2 provides an example that illustrates the effect of experimentation on the price-dividend ratio, the risk premium, and the volatility of stock returns. Panel (a) shows that the price-dividend ratio reaches a minimum at the socially optimal level of experimentation of Proposition 2 (the dot labeled **S** on the graph). To understand why this is the case, re-write Eq. (11) under the following form:

$$F(\widehat{\beta}_0, \nu_0, 0) = (1-\gamma)(\theta\delta_0)^{\gamma-1} J_a(\delta_0, \widehat{\beta}_0, \nu_0, 0), \tag{26}$$

This shows that the price-dividend ratio is directly related to the value function of the active investor, with a negative sign when the coefficient of risk aversion is higher than one. Because the active investor chooses an experimentation level that maximizes the value function, it follows that the equilibrium price-dividend ratio reaches a minimum when experimentation is socially optimal.

The direct implication of this inverse relationship is that any level of experimentation that deviates from the optimum *always* increases the price-dividend ratio. Panel (a) shows that the price-dividend ratio increases under monopoly (Proposition 3, $N = 1$, dot labeled **M**), further increases under duopoly ($N = 2$, dot labeled **D**), and reaches a maximum under perfect competition ($N \to \infty$, dot labeled **PC**).

Panels (b) and (c) of Figure 2 show that the risk premium and the volatility are hump-shaped in experimentation. The initial increase is driven by the first terms in (23)-(24), whereas the subsequent decrease is driven by the second, uncertainty

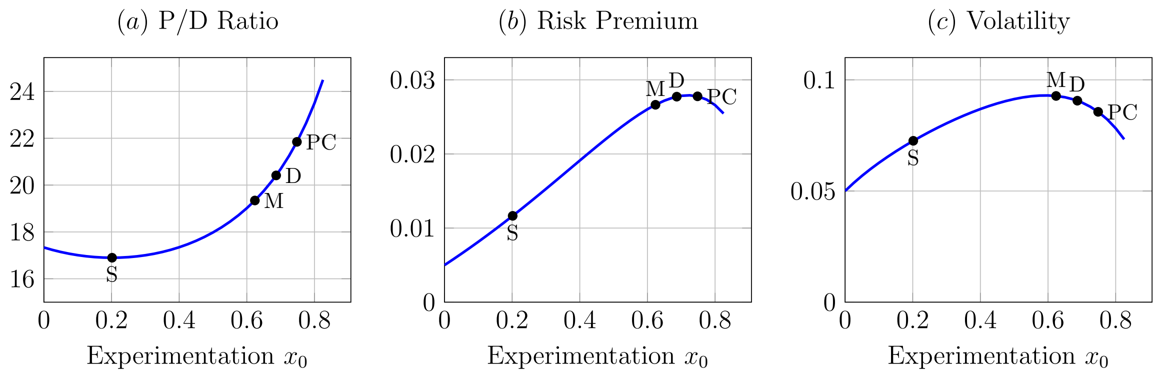

Figure 2: **Experimentation and Asset Prices.** The effect of experimentation on the price-dividend ratio, the risk premium, and stock market volatility. The dot labeled **S** represents the socially optimal amount of experimentation, x_S^* (Proposition 2). The dot **M** represents the optimal amount of experimentation under monopoly, x_M^* (Proposition 3, $N = 1$). The dots **D** and **PC** represent the total amount of experimentation under Cournot competition, with $N = 2$ and $N \to \infty$ respectively (Proposition 3). All parameters are provided in Figure 1.

terms.

Based on our analysis, there appears to be a higher risk premium associated with creative destruction, but this admittedly does depend on the parameters specified in the model. For most reasonable specifications, however, this will indeed be the case. To explore this further, we define the *normalized risk premium* (i.e., risk premium per unit of variance) in the economy as[14]

$$\frac{RP_t}{\sigma_{P,t}^2} = \frac{\gamma}{1 - (\gamma - 1)\frac{x_0^2 \nu_t}{\sigma^2(1+kx_0)^2}\mathbb{D}_t}. \tag{27}$$

This expression conveniently isolates in the denominator the impact of experimentation x_0 and uncertainty ν_t. Experimentation leads to an increase in the equilibrium normalized risk premium above the value γ that would prevail in the market without experimentation. Thus, comparing a market characterized by innovation and creative destruction with a market without creative destruction yields a higher risk premium per unit of risk for the former. This is consistent with the fact that agents with risk-aversion $\gamma > 1$ require a risk premium when more uncertainty is present in the economy (see Footnote 7). Furthermore, Eq. (27) suggests a testable prediction of

[14]The normalized risk premium, divided by the risk aversion, represents the tangency portfolio in a portfolio choice problem with random investment opportunities (Merton, 1973; Brennan, 1998a).

our model, in that sectors characterized by stronger invention activity are riskier.[15]

Experimentation has further implications for the term structure of risk. To see this, in Corollary 5.1 we decompose the asset into dividend strips (i.e., assets that pay the aggregate consumption only at time $s > t$).

Corollary 5.1 *For any experimentation level x_0, the price of a dividend strip with maturity $s > t$ is*

$$P_{t,s} \equiv \frac{1}{\xi_t}\mathbb{E}_t[\xi_s \delta_s] = \delta_t \exp\left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s-t)^2\right]. \tag{28}$$

The risk premium and the volatility for each maturity $s \geq t$ are given by

$$RP_{t,s} = \gamma\sigma^2(1+kx_0)^2 + \gamma(1-\gamma)x_0^2\nu_t(s-t) \tag{29}$$

$$|\sigma_{P,t,s}| = \sigma(1+kx_0)\left|1 + (1-\gamma)\frac{x_0^2\nu_t}{\sigma^2(1+kx_0)^2}(s-t)\right|. \tag{30}$$

Experimentation has two effects on the term structure of risk premia and volatilities. First, by increasing the disturbance associated with the introduction of new technologies, experimentation increases the risk premia and volatilities at all maturities. This can be seen from the first terms in (29)-(30). Second, experimentation introduces uncertainty about the expected growth of the new technology. This dampens the risk premia and volatilities (as elaborated above). Corollary 5.1 shows that the dampening effect positively depends on maturity. The volatility of dividend strips with longer maturities is lowered through learning, which also induces lower risk premia. As a result, the term structures of risk premia and volatilities become downward sloping with experimentation.

The effect of experimentation on the term structure of risk is shown in Figure 3. Without experimentation (solid lines), the term structure of risk premia and volatilities is flat. Then, consistent with Corollary 5.1, experimentation generates downward sloping term structures. 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, more so in the far future. Competition among active investors amplifies this steepening effect. The slopes

[15] This prediction has been tested empirically by Grammig and Jank (2015), who show that cross-sectional differences across size and book-to-market sorted portfolios can be explained as premia for bearing creative destruction risk.

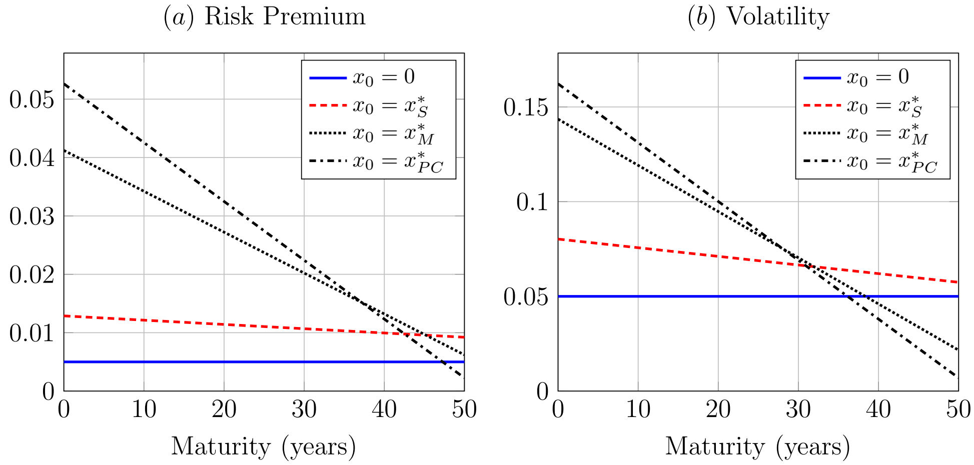

Figure 3: **Experimentation and the Term Structure of Risk.** The term structure of risk premia and volatilities in several cases. The solid lines depict the flat term structure that arise in the "status-quo" economy with no experimentation. The dashed lines depict the case of socially optimal experimentation (Proposition 2). The remaining two lines (dotted and dash-dotted) depict term structures that arise with monopoly or with perfect competition (Proposition 3). All parameters are provided in Figure 1.

are steepest in the perfect competition case, when the total experimentation in the economy equals $x_0 = x_{PC}^*$.

3.1 The Aftermath of Experimentation

We now consider the consequences on asset prices in the aftermath of the decision to experiment. We illustrate how the rent seeking behavior of active investors engaged in creative destruction amplifies the effect that competition has on stock price valuations and generates patterns that ex post can be identified as booms and busts.

To this end, we consider an economy with perfect competition ($N \to \infty$) and assume a relatively high value for the rent seeking parameter, $w = 5$. The reason for this choice is twofold. First, it implies an equilibrium price-dividend ratio well above 30, which is often considered as a warning sign that stock markets are in a bubble (Shiller, 2015). Second, it implies an aggregate consumption share of roughly 5% for the mass of active investors, who initially start from a share of 0.1%. This large increase in consumption share highlights the power of rent-seeking behavior for wealth redistribution and is also in line with widely acknowledged facts about top income

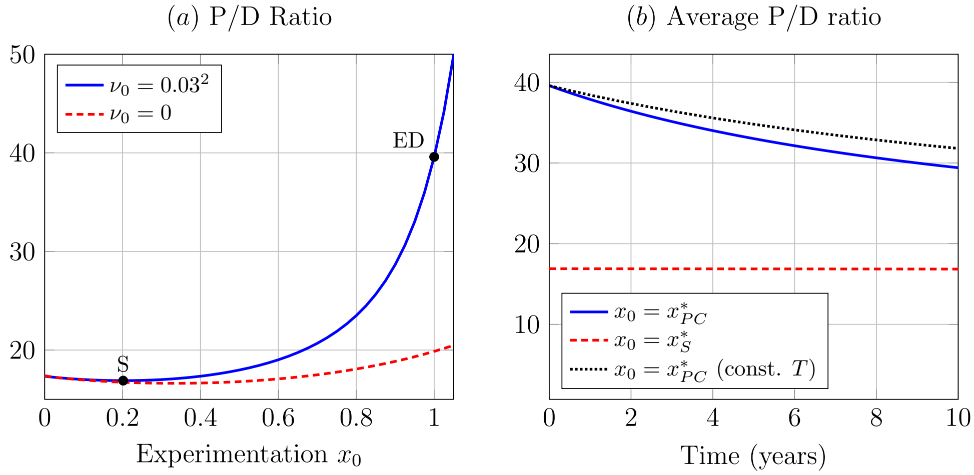

Figure 4: **Booms and Busts in Asset Prices.** Panel (a) shows the price-dividend ratio as a function of the aggregate experimentation level in the economy, x_0. The solid line includes uncertainty about β. Uncertainty is set at zero for the dashed line. The dot labeled **ED** ("El Dorado") corresponds to the equilibrium aggregate level of experimentation in perfect competition. The dot labeled **S** corresponds to the socially optimal level of experimentation. Panel (b) depicts average price paths over time starting from **ED** (solid and dotted lines) and from **S** (dashed line), starting from $\widehat{\beta}_0 = \beta$. The dotted line maintains the maturity of the asset constant at $T = 100$ years. With the exception of $w = 5$, all parameters are provided in Figure 1.

inequality worldwide (e.g., Aghion et al., 2015).

Panel (a) in Figure 4 depicts the price-dividend ratio in this economy as a function of the aggregate experimentation (solid line). The dot labeled **ED** ("El Dorado") represents the equilibrium price-dividend ratio with perfect competition. In the same plot, the dashed line depicts the price-dividend ratio when uncertainty about β is fixed at $\nu_0 = 0$. The gap between the two lines provides a measure of the overvaluation due to uncertainty (see Proposition 5 and its discussion). We also show on the plot the point **S**, which corresponds to the socially optimal experimentation level (Proposition 2). Rent seeking behavior and intense competition lead to over-experimentation. In turn, this adds uncertainty to the economy and increases the the gap between the two lines. This gap grows from being almost negligible at **S** to roughly 100% at **ED**.

It is instructive to compare the dynamics of the price-dividend ratio once the points **ED** and **S** have been reached. To do so, let us assume that the prior $\widehat{\beta}_0$ is exactly equal to the true value of β. That is, we are assuming that agents are initially right—but still uncertain—about the expected growth of the new technology, and thus we are not imposing any optimistic/pessimistic bias in the future paths of prices. As it turns out, this assumption is extremely convenient, because it ensures that simulations of the economy over time will result in an average $\widehat{\beta}_t$ which is equal to the true value of β.[16] Furthermore, because the remaining state variable which enters in the valuation of the price-dividend ratio, the uncertainty ν_t, decreases deterministically over time, we can directly plot the average price-dividend ratio over time by simply using Eq. (8) for the dynamics of ν_t, without resorting to simulations.

Panel (b) in Figure 4 depicts these dynamics. The solid line corresponds to the average path of the price-dividend ratio starting from **ED**, whereas the dashed line shows the average path of the price-dividend ratio starting from **S**. Because the remaining life of the asset diminishes, we also add to the plot the dotted line, which isolates the effect of the decrease in uncertainty by holding the remaining asset life constant at $T = 100$ years (this adjustment is unnecessary in the **S** case, where the decrease in the price-dividend ratio is almost imperceptible). The panel shows that, after reaching the point **ED**, asset prices have the tendency to decrease as uncertainty about the new technology resolves. In contrast, after experimentation at the socially optimal level (point **S**), the average decrease in asset price due to resolution of uncertainty is negligible.

As such, competition among rent-seeking agents in markets for new technologies generates price patterns that resemble a "bubble." It is important to mention that the situation depicted in panel (b) can be identified as a bubble only ex post; ex ante, agents do not expect prices to fall. But, because learning by doing decreases the uncertainty about β over time, stock prices adjust endogenously. The pace of this adjustment process depends on the rate at which uncertainty about the new technology is resolved. More specifically, Eq. (8) shows that the posterior variance decays faster with over-experimentation,[17] which accelerates the eventual decline in asset prices. Our model therefore predicts that markets characterized by fierce competition for new technologies are not only prone to stronger asset price inflation but also to steeper

[16]This can be seen from Proposition 1. At time $t = 0$, $d\widehat{W}_0 = dW_t$ (because $\beta - \widehat{\beta}_0 = 0$). Technically, at time $t = 0$ the filter $\widehat{\beta}$ is a martingale. This ensures that the average of its future simulated values one step ahead, $\widehat{\beta}_{0+dt}$, is exactly β. Then apply the same reasoning at time $t = 0 + dt$.

[17]In Eq. (8), the coefficient multiplying time in the denominator can be interpreted as the *speed of learning*. It is a matter of algebra to show that the speed of learning increases with experimentation.

subsequent declines through learning by doing and faster resolution of uncertainty.

Finally, we highlight here the importance of learning for the evolution of bubbles. The aftermath depicted in Panel (b) of Figure 4 can only take place through learning. 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.

3.2 A Possible Ex Ante Diagnosis of "Bubbles"

Bubble peaks are notoriously hard to identify (Fama, 2014; Greenwood et al., 2017). In this section, we try to address this issue by proposing a plausible test for bubbles, that depends on future financial innovation.

In our model, high stock valuations and price-dividend ratios may arise not only because of agents' expectations about future growth, but also because of uncertainty coupled with rent seeking behavior and intensified competition. According to Proposition 5, the price-dividend ratio in this economy depends on both $\widehat{\beta}_t$ and ν_t:

$$F(\widehat{\beta}_t, \nu_t, t; x_0) = \int_t^T \exp\left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2} x_0^2 \nu_t (s-t)^2\right] ds. \qquad (31)$$

It is therefore not possible to infer both the expected growth and the degree of uncertainty from the price-dividend ratio alone. However, notice that the expected growth enters this valuation equation as a linear function of the maturity $(s-t)$, whereas parameter uncertainty enters as a quadratic function. This suggests that if a market for dividend strips were established (Brennan, 1998b), defined as claims to each year's dividend on an index, this might help detect bubbles ex ante by identifying when high stock price valuations are primarily driven by uncertainty and competition.

To illustrate this, consider Figure 5, where we use Corollary 5.1 to plot the log values of dividend strips for the remaining life of the asset. We consider two situations. The solid line shows the log values of dividend strips when the price-dividend ratio is at the point **ED** ("El Dorado"). The dashed line pertains to the socially optimal experimentation case (point **S** in Figure 4).

The solid line shows a clear non-linear pattern, which arises from the quadratic term in Eq. (31). This effect is almost nil with the dashed line, which is indistinguishable from a straight line. If a market for dividend strips were present, a statistical

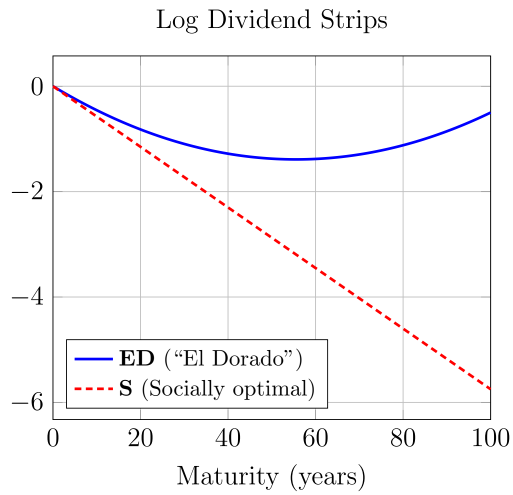

Figure 5: **An Ex Ante Diagnosis of "Bubbles."** Log prices of dividend strips (Corollary 5.1) in two situations. The solid line corresponds to the dot **ED** ("El Dorado") on Figure 4. The dashed line corresponds to the dot **S** (Socially optimal experimentation) on Figure 4. With the exception of $w = 5$, all parameters are provided in Figure 1.

test could be devised to help reject the null hypothesis of linear dependence on maturity. We emphasize that this test is different than a standard test of the no-bubble condition (Giglio, Maggiori, and Stroebel, 2016). In our case, prices of strips at *all* maturities offer a "term-structure" pattern whose shape becomes non-linear if asset prices are overvalued due to uncertainty.

4 Dynamic Experimentation

So far, we have assumed that the decision to experiment is made only once, at $t = 0$. In reality, however, innovation is not a one-time decision, but often occurs in waves (Gort and Klepper, 1982). Now, we return to the socially optimal case of Section 2.1 and consider that the active 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.[18]

The active agent's expected lifetime utility of consumption J_a satisfies the follow-

[18]See also Rob (1991) for a sequential model of entry in which uncertainty gradually resolves over time. In our case, the decision to alter the experimentation level is made by a single agent.

ing partial differential equation at any time t:

$$0 = \max_x \left[\mathcal{D} J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) + e^{-\rho t} \frac{(\theta c_t)^{1-\gamma}}{1-\gamma} \right], \tag{32}$$

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

$$J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t} \frac{(\theta \delta_t)^{1-\gamma}}{1-\gamma} F(\widehat{\beta}_t, \nu_t, t) \tag{33}$$

results in a partial differential equation for the price-dividend ratio $F(\cdot)$, which we relegate to Appendix A.3.2 for sake of brevity. The optimal level of experimentation then follows from the first order condition on x_t.

Proposition 6 ***(Optimal Dynamic Experimentation)*** *If the problem* (32) *has an interior maximum, then the optimal level of experimentation at time t solves*

$$x_t^* = \frac{\widehat{\beta}_t - \gamma k \sigma^2}{\gamma k^2 \sigma^2} + \frac{\nu_t}{\gamma k^2 \sigma^2} \left(\frac{F_\beta}{F} - \frac{x_t^* \nu_t}{(\gamma - 1)\sigma^2 (1 + k x_t^*)^3} \frac{F_{\beta\beta} - F_\nu}{F} \right). \tag{34}$$

The solution (34) constitutes an implicit form since the control x_t 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 active 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. 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 6 shows that in the dynamic case the optimal experimentation level fluctuates as new information becomes available and affects the agent's expectations. This has further impact on asset prices in the economy, which we characterize below.

Proposition 7 ***(Asset Prices with Dynamic Experimentation)*** *In an economy with dynamic experimentation, the risk-free rate and the market price of risk are given by*

$$r_t^f = \rho + \gamma(\bar{f} + x_t^* \widehat{\beta}_t) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kx_t^*)^2 \tag{35}$$

$$\theta_t = \gamma\sigma(1+kx_t^*), \tag{36}$$

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

$$\mu_{S,t}^C - r_t^f = \gamma\sigma^2(1+kx_t^*)^2 + \gamma(1-\gamma)(x_t^*)^2\nu_t\widetilde{\mathbb{D}}_t \tag{37}$$

$$|\sigma_{P,t}| = \sigma(1+kx_t^*)\left|1 + (1-\gamma)\frac{(x_t^*)^2\nu_t}{\sigma^2(1+kx_t^*)^2}\widetilde{\mathbb{D}}_t\right|. \tag{38}$$

Expressions (35)-(38) have a similar structure with the ones in the static case (Section 3), with two main differences. First, the market price of risk now fluctuates and increases with the level of experimentation. Second, in (37)-(38), 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.3.2 . 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 learning channel, which dampens asset price fluctuations. Therefore, in the dynamic case as in the static case, both the risk premium and the volatility will feature a hump-shaped pattern in experimentation.

We compare the optimal level of experimentation, the risk premium and the volatility in the static versus the dynamic case in Figure 6, first as a function of the filter $\widehat{\beta}_t$ in the upper panels, and then as a function of the uncertainty ν_t in the lower panels.[19] The static experimentation case is depicted with blue solid lines, whereas the dynamic experimentation case is depicted with red dashed lines.

When the filter $\widehat{\beta}_0$ is sufficiently low, neither the "static" or the "dynamic" agent decide to experiment, as shown in panel (a). When experimentation is positive, the "dynamic" experimenter allocates more capital to the new technology than the "static" experimenter for any level of $\widehat{\beta}_t$. This is because the dynamic experimenter

[19] We use a finite difference scheme to solve the partial differential equation associated with the problem (32).

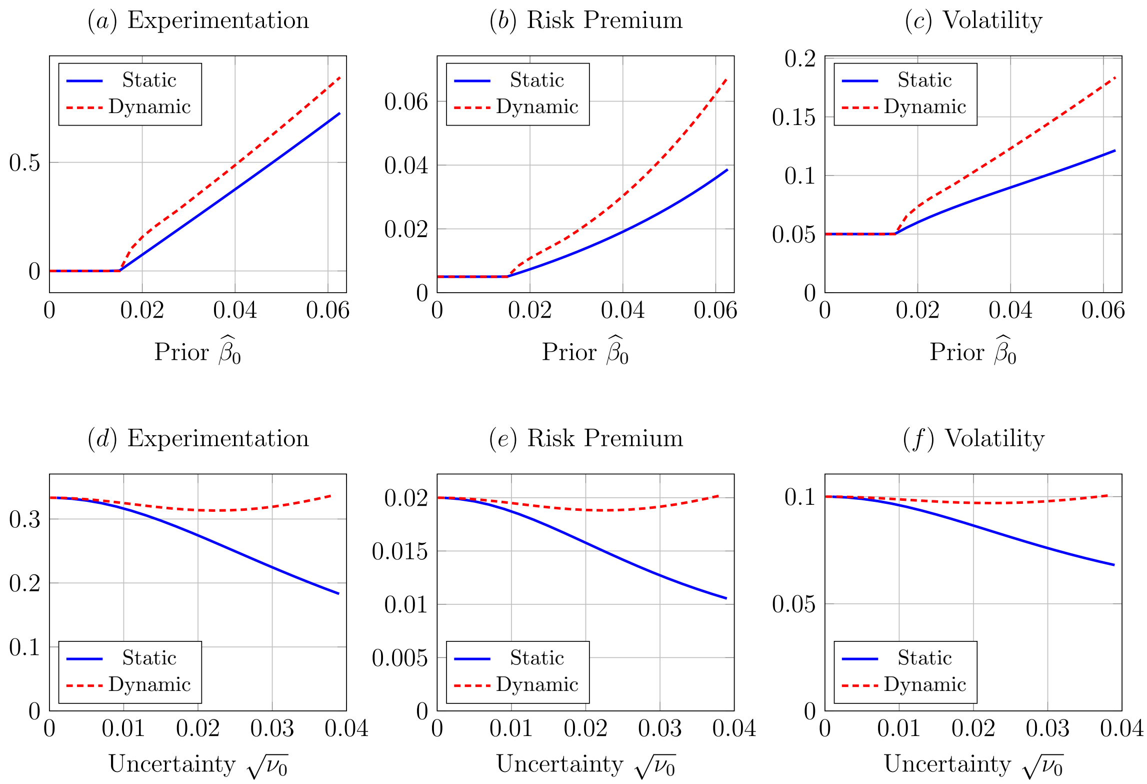

Figure 6: **Static Versus Dynamic Experimentation.** The six panels depict the optimal level of experimentation, the risk premium and the volatility of asset returns with static experimentation (solid lines) versus dynamic experimentation (dashed lines). The upper panels plot functions of the prior $\widehat{\beta}_0$. The lower panels plot functions of the square root of uncertainty, $\sqrt{\nu_0}$. All parameters are provided in Figure 1, except for the maturity, which is fixed at $T = 50$ in order to improve numerical accuracy.

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, as shown in panels (b) and (c). Indeed, the option to abandon lowers risk for the dynamic experimenter. But, internalizing this she experiments more, which raises the risk premium and volatility.

In the lower panels of Figure 6, the risk premium and the volatility are plotted for different values of the uncertainty ν_t. In 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 remain more or less constant in ν_t (eventually, they decrease once the last terms in (37)-(38) 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. 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 experiments more aggressively initially and thus both the risk premium and the volatility remain high with dynamic experimentation and for intermediate values of uncertainty.

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 should be willing to take.

From the financier's viewpoint, an optimum exists. 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. Rent-seeking behavior and intensified competition for new markets leads to socially excessive experimentation, inflating asset prices and generating high risk premia and high volatility. In hindsight, asset prices exhibit familiar boom and bust patterns observed during technological revolutions.

Rent seeking behavior by active investors has welfare consequences. It reduces the passive investors' (those who are not involved in entrepreneurial activity) consumption share, exposes them to too much risk and lowers the expected utility of their consumption stream. Higher competition exacerbates this because experimentation rises, which hurts the residual claimants in the economy.

A worthwhile direction for future research would be to link a proxy for experimentation (i.e., x_t^* in our model) with dynamic patterns in asset prices. Greenwood et al. (2017) find that, in conjunction with an sharp price increase, economies characterized by high stock issuance have a heightened probability of a subsequent crash.

Our model offers a justification for their finding and a theoretical argument in favor of using more information in order to identify bubbles.

We have assumed throughout this paper that investors share the same beliefs about the expected growth of a new technology. In reality, these markets might very well be characterized by strong divergence of opinion about the promise of new innovations. It might be an interesting exercise to study the interaction between competition and difference of beliefs. Our guess is that active agents will decide to experiment more in order to preempt entry of optimistic competitors, in this way exacerbating the effect of over-experimentation on asset prices.

Finally, it is important to further explore the implications of dynamic experimentation on asset prices. Because dynamic experimentation provides the agent an additional option to abandon a new technology at any point in the future, it can have adverse consequences for the valuation of firms involved in entrepreneurial activity. Coupled with lack of perfect knowledge, dynamic experimentation might lead the financier to conclude that a particular new technology is not productive and abandon it prematurely. This "abandonment risk" has consequences for asset prices.

A Appendix

A.1 Experimentation and Learning

A.1.1 Microfoundation

Consider two intermediate goods, a “status-quo” good and an “experimental” good:

$$\frac{d\delta_t^S}{\delta_t^S} = \bar{f}dt + \sigma dW_t \tag{39}$$

$$\frac{d\delta_t^E}{\delta_t^E} = \left(\Gamma - \frac{k^2\sigma^2}{2}x_0\right)dt + k\sigma dW_t. \tag{40}$$

The parameter Γ, which enters in the expected growth of the experimental good, is an unknown constant. The expected growth for the experimental good exhibits decreasing returns to scale: a higher initial level of experimentation x_0 decreases its expected growth by a term proportional to the instantaneous variance of δ_t^E, $k^2\sigma^2$. Define the consumption basket in the economy as[20]

$$\delta_t = \delta_t^S \left(\delta_t^E\right)^{x_0}. \tag{41}$$

If $x_0 = 0$, the economy remains in the status-quo and only the good δ_t^S is consumed. When $x_0 > 0$, the experimental good becomes a new variety in the consumption basket, with a weight given by $x_0/(1+x_0)$. Aplying Itô’s lemma on δ_t yields:

$$\frac{d\delta_t}{\delta_t} = \left(\bar{f} + \beta x_0\right)dt + \sigma(1+kx_0)dW_t, \tag{42}$$

where $\beta \equiv \Gamma + \frac{1}{2}(2-k)k\sigma^2$ is an unknown constant. This provides a microfoundation of Eq. (3) in the text.

A.1.2 Proof of Proposition 1 (Learning)

The proof of Proposition 1 follows from direct application of this standard result in filtering theory:

Theorem 8 ***(Liptser and Shiryayev, 1977)*** *Consider an unobservable process u_t and an observable process s_t with dynamics*

$$du_t = \left[a_0(t,s_t) + a_1(t,s_t)u_t\right]dt + b_1(t,s_t)dZ_t^u + b_2(t,s_t)dZ_t^s \tag{43}$$

$$ds_t = \left[A_0(t,s_t) + A_1(t,s_t)u_t\right]dt + B_1(t,s_t)dZ_t^u + B_2(t,s_t)dZ_t^s. \tag{44}$$

All the parameters can be functions of time and of the observable process. Then, the filter evolves according to (we drop the dependence of coefficients on t and s_t for notational

[20] This specification is commonly adopted in the international finance literature. See Helpman and Razin (2014), Cole and Obstfeld (1991), Zapatero (1995), and Pavlova and Rigobon (2007) among others.

convenience):

$$d\widehat{u}_t = (a_0 + a_1 \widehat{u}_t)dt + [(b \circ B) + \nu_t A_1^\top](B \circ B)^{-1}[ds_t - (A_0 + A_1 \widehat{u}_t)dt] \quad (45)$$

$$\frac{d\nu_t}{dt} = a_1 \nu_t + \nu_t a_1^\top + (b \circ b) - [(b \circ B) + \nu_t A_1^\top](B \circ B)^{-1}[(b \circ B) + \nu_t A_1^\top]^\top, \quad (46)$$

where

$$b \circ b = b_1 b_1^\top + b_2 b_2^\top \quad (47)$$

$$B \circ B = B_1 B_1^\top + B_2 B_2^\top \quad (48)$$

$$b \circ B = b_1 B_1^\top + b_2 B_2^\top. \quad (49)$$

In the present setup, the unobservable variable is the constant β. Hence,

$$a_0 = a_1 = b_1 = b_2 = 0. \quad (50)$$

Furthermore, the observable process is δ_t. Applying Itô's lemma on $\ln \delta_t$ yields

$$A_0 = \bar{f} - \frac{1}{2}\sigma^2(1 + kx_0)^2, \quad A_1 = x_0, \quad B_1 = 0, \quad B_2 = \sigma(1 + kx_0). \quad (51)$$

Direct application of Theorem 8 yields Proposition 1. □

A.1.3 Conditional moments of future output

We prove the following proposition using the theory of affine processes (Duffie, Filipović, and Schachermayer, 2003):[21]

Proposition 9 *For any $s > t$, the expected value of future dividends and the future volatility of dividends are respectively given by*

$$\mathbb{E}_t[\delta_s] = \delta_t \exp\left[\left(\bar{f} + x_0 \widehat{\beta}_t\right)(s - t) + \frac{x_0^2 \nu_t}{2}(s - t)^2\right] \quad (52)$$

$$\mathrm{Var}_t[\ln \delta_s] = \sigma^2(1 + kx_0)^2(s - t) + x_0^2 \nu_t (s - t)^2. \quad (53)$$

Both quantities are increasing in the level of experimentation x_0.

Proof Apply first Itô's lemma on $\ln \delta_t$, with the process of δ_t provided in (5). Using the fact that $\widehat{\beta}_t$ is a martingale yields

$$\mathbb{E}_t[\ln \delta_s] = \ln \delta_t + \left(\bar{f} + x_0 \widehat{\beta}_t - \frac{\sigma^2(1 + kx_0)^2}{2}\right)(s - t). \quad (54)$$

[21] An alternative approach would be to follow Ziegler (2003, Appendix A) and Bryson and Ho (1975, Section 11.4) and compute $\mathbb{E}_t[\delta_s]$ and $\mathrm{Var}[\ln \delta_s]$ in one step. Both approaches are equally tedious.

Write the dynamics of the system of two state variables $\{\ln \delta_t, \widehat{\beta}_t\}$ under an affine form:

$$\begin{bmatrix} d\ln\delta_t \\ d\widehat{\beta}_t \end{bmatrix} = \left(\begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1+kx_0)^2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & x_0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \ln\delta_t \\ \widehat{\beta}_t \end{bmatrix} \right) dt + \begin{bmatrix} \sigma(1+kx_0) \\ A(t) \end{bmatrix} d\widehat{W}_t, \tag{55}$$

where $A(t)$ is a function of time (see Proposition 1):

$$A(t) \equiv \frac{\nu_0 x_0 \sigma(1+kx_0)}{\nu_0 x_0^2 t + \sigma^2(1+kx_0)^2}. \tag{56}$$

The following Lemma results immediately from differentiation of $A(t)$:

Lemma 10 *The function $A(t)$ satisfies*

$$A'(t) = -\frac{x_0}{\sigma(1+kx_0)} A(t)^2. \tag{57}$$

Notice that (55) is a *time-inhomogeneous* multifactor affine process (Filipović, 2005). This is because the diffusion of $\widehat{\beta}$ depends on time (but it does not depend on the two state variables). Define:

$$K_0 \equiv \begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1+kx_0)^2 \\ 0 \end{bmatrix} \tag{58}$$

$$K_1 \equiv \begin{bmatrix} 0 & x_0 \\ 0 & 0 \end{bmatrix} \tag{59}$$

$$H_0(t) \equiv \begin{bmatrix} \sigma(1+kx_0) \\ A(t) \end{bmatrix} \begin{bmatrix} \sigma(1+kx_0) \\ A(t) \end{bmatrix}' = \begin{bmatrix} \sigma^2(1+kx_0)^2 & \sigma(1+kx_0)A(t) \\ \sigma(1+kx_0)A(t) & A(t)^2 \end{bmatrix}. \tag{60}$$

In order to compute the expectation

$$\mathbb{E}_t[\delta_s] = \mathbb{E}_t\left[e^{\ln\delta_s}\right], \tag{61}$$

we conjecture an exponential-affine solution of the form

$$\mathbb{E}_t[\delta_s] = e^{\alpha_0(s-t)+\alpha_1(s-t)\ln\delta_t+\alpha_2(s-t)\widehat{\beta}_t}, \tag{62}$$

for some coefficient functions $\alpha_j(\cdot)$, $j = 0, 1, 2$. Since the system (55) is a multifactor affine diffusion, the coefficients $\alpha_j(\cdot)$, $j = 0, 1, 2$ satisfy a system of Riccati ODEs (Duffie et al., 2003). Defining $\tau \equiv s - t$, the system of ODE writes

$$\begin{bmatrix} \alpha_1'(\tau) \\ \alpha_2'(\tau) \end{bmatrix} = K_1^\top \begin{bmatrix} \alpha_1(\tau) \\ \alpha_2(\tau) \end{bmatrix} \tag{63}$$

$$\alpha_0'(\tau) = K_0^\top \begin{bmatrix} \alpha_1(\tau) \\ \alpha_2(\tau) \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \alpha_1(\tau) & \alpha_2(\tau) \end{bmatrix} H_0(t) \begin{bmatrix} \alpha_1(\tau) \\ \alpha_2(\tau) \end{bmatrix}, \tag{64}$$

with boundary conditions $\alpha_0(0) = 0$, $\alpha_1(0) = 1$, and $\alpha_2(0) = 0$. The first Riccati Eq. (63) has a straightforward solution:

$$\alpha_1(\tau) = 1 \tag{65}$$
$$\alpha_2(\tau) = x_0\tau, \tag{66}$$

which can be now inserted in the remaining Riccati Eq. (64):

$$\alpha_0'(\tau) = \begin{bmatrix}\bar{f} - \frac{1}{2}\sigma^2(1+kx_0)^2 & 0\end{bmatrix}\begin{bmatrix}1\\ x_0\tau\end{bmatrix} + \frac{1}{2}\begin{bmatrix}1 & x_0\tau\end{bmatrix} H_0(s-\tau)\begin{bmatrix}1\\ x_0\tau\end{bmatrix}, \tag{67}$$

which leads to

$$\begin{aligned}\alpha_0'(\tau) &= \bar{f} + \sigma(1+kx_0)x_0 A(s-\tau)\tau + \frac{x_0^2}{2}A(s-\tau)^2\tau^2 && (68)\\ &= \bar{f} + \sigma(1+kx_0)x_0\left[A(s-\tau)\tau + \frac{1}{2}\frac{x_0}{\sigma(1+kx_0)}A(s-\tau)^2\tau^2\right] && (69)\end{aligned}$$

and boundary condition $\alpha_0(0) = 0$. Using Lemma 10, we can write:

$$\begin{aligned}\alpha_0'(\tau) &= \bar{f} + \sigma(1+kx_0)x_0\left[A(s-\tau)\tau + \frac{1}{2}\frac{\partial A(s-\tau)}{\partial\tau}\tau^2\right] && (70)\\ &= \bar{f} + \sigma(1+kx_0)x_0\frac{\partial\left(\frac{1}{2}A(s-\tau)\tau^2\right)}{\partial\tau}, && (71)\end{aligned}$$

which we can now integrate to get

$$\begin{aligned}\alpha_0(s-t) &= \bar{f}(s-t) + \frac{1}{2}\sigma(1+kx_0)x_0\frac{x_0}{\sigma(1+kx_0)}\nu_t(s-t)^2 && (72)\\ &= \bar{f}(s-t) + \frac{1}{2}x_0^2\nu_t(s-t)^2 && (73)\end{aligned}$$

It then follows that

$$\begin{aligned}\mathbb{E}_t\left[\delta_s\right] &= e^{\alpha_0(s-t)+\alpha_1(s-t)\ln\delta_t+\alpha_2(s-t)\widehat{\beta}_t} && (74)\\ &= \exp\left[\ln\delta_t + \left(\bar{f} + x_0\widehat{\beta}_t\right)(s-t) + \frac{x_0^2\nu_t}{2}(s-t)^2\right] && (75)\\ &= \delta_t\exp\left[\left(\bar{f} + x_0\widehat{\beta}_t\right)(s-t) + \frac{x_0^2\nu_t}{2}(s-t)^2\right], && (76)\end{aligned}$$

which is increasing in the level of experimentation x_0. Write now

$$\mathbb{E}_t\left[\delta_s\right] = \mathbb{E}_t\left[e^{\ln\delta_s}\right] = \exp\left[\mathbb{E}_t[\ln\delta_s] + \frac{1}{2}\mathrm{Var}_t[\ln\delta_s]\right], \tag{77}$$

which, using (54) and (76) yields

$$\exp\left[\frac{\sigma^2(1+kx_0)^2}{2}(s-t)+\frac{x_0^2\nu_t}{2}(s-t)^2\right]=\exp\left[\frac{1}{2}\mathrm{Var}_t[\ln\delta_s]\right], \tag{78}$$

and thus

$$\mathrm{Var}_t\left[\ln\delta_s\right]=\sigma^2(1+kx_0)^2(s-t)+x_0^2\nu_t(s-t)^2, \tag{79}$$

which is increasing in the level of experimentation x_0. □

A.1.4 Value Function of the Active Investor

Proposition 11 *The value function of the active investor unambiguously increases with expected future dividends $\mathbb{E}_t[\delta_s]$ and unambiguously decreases with the future variance $\mathrm{Var}_t[\ln\delta_s]$, for any $s \geq t$. Furthermore, the value function can be written*

$$J_a(\delta_t,\widehat{\beta}_t,\nu_t,t)=\frac{e^{-\rho t}(\theta\delta_t)^{1-\gamma}}{1-\gamma}F(\widehat{\beta}_t,\nu_t,t), \tag{80}$$

where

$$F(\widehat{\beta}_t,\nu_t,t)\equiv\int_t^T\exp\left[\kappa(x_0,\widehat{\beta}_t)(s-t)+\frac{(1-\gamma)^2}{2}x_0^2\nu_t(s-t)^2\right]ds \tag{81}$$

and

$$\kappa(x_0,\widehat{\beta}_t)\equiv(1-\gamma)\left(\bar{f}+x_0\widehat{\beta}_t-\gamma\frac{\sigma^2(1+kx_0)^2}{2}\right)-\rho. \tag{82}$$

Proof In equilibrium, the active investor consumes a fraction θ of the entire output δ_t and thus her lifetime expected utility of consumption can be computed as

$$J_a(\delta_t,\widehat{\beta}_t,\nu_t,t)=\mathbb{E}_t\left[\int_t^T e^{-\rho s}\frac{(\theta\delta_s)^{1-\gamma}}{1-\gamma}ds\right]=\frac{\theta^{1-\gamma}}{1-\gamma}\int_t^T e^{-\rho s}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]ds, \tag{83}$$

where the second equality results from application of Fubini's theorem. The expectation in Eq. (10) can be further expanded by using the property that for a normally distributed random variable $y=\ln(x)$, $\mathbb{E}[x^\alpha]=\exp(\alpha\mathbb{E}[y]+\alpha^2\mathrm{Var}[y]/2)$:

$$J_a(\delta_t,\widehat{\beta}_t,\nu_t,t)=\frac{\theta^{1-\gamma}}{1-\gamma}\int_t^T e^{-\rho s}\exp\left[(1-\gamma)\mathbb{E}_t[\ln\delta_s]+\frac{(1-\gamma)^2}{2}\mathrm{Var}_t[\ln(\delta_s)]\right]ds. \tag{84}$$

Replacing (54) and (79) yields

$$J_a(\delta_t,\widehat{\beta}_t,\nu_t,t)=\frac{e^{-\rho t}(\theta\delta_t)^{1-\gamma}}{1-\gamma}\int_t^T\exp\left[\kappa(x_0,\widehat{\beta}_t)(s-t)+\frac{(1-\gamma)^2}{2}x_0^2\nu_t(s-t)^2\right]ds, \tag{85}$$

with $\kappa(x_0,\widehat{\beta}_t)$ defined in (82). In order to show that $J_a(\delta_t,\widehat{\beta}_t,\nu_t,t)$ increases in $\mathbb{E}_t[\delta_s]$ and decreases in $\mathrm{Var}_t[\ln \delta_s]$, replace the relation

$$\mathbb{E}_t[\ln \delta_s] = \ln(\mathbb{E}_t[\delta_s]) - \frac{1}{2}\mathrm{Var}_t[\ln \delta_s] \tag{86}$$

in (84) to obtain

$$J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) = \frac{\theta^{1-\gamma}}{1-\gamma}\int_t^T e^{-\rho s}\mathbb{E}_t[\delta_s]^{1-\gamma}\exp\left[-\frac{\gamma(1-\gamma)}{2}\mathrm{Var}_t[\ln \delta_s]\right] ds. \tag{87}$$

From Eq. (87), it is a matter of algebra to show that $J(\delta_t, \widehat{\beta}_t, t)$ increases in $\mathbb{E}_t[\delta_s]$ and decreases in $\mathrm{Var}_t[\ln \delta_s]$, for any value of the risk aversion parameter γ. □

A.1.5 Proof of Proposition 2 (Socially Optimal Experimentation)

The expected lifetime utility of the passive investor equals:

$$J_p(\delta_t, \widehat{\beta}_t, \nu_t, t) = \frac{e^{-\rho t}((1-\theta)\delta_t)^{1-\gamma}}{1-\gamma}F(\widehat{\beta}_t, \nu_t, t). \tag{88}$$

Because the consumption share θ is constant, the choice of the active investor also maximizes the expected lifetime utility of the passive investor and is therefore socially optimal (in other words, different values of the parameter $\theta \in [0,1]$ will always yield the same amount of optimal experimentation). The first order condition for the active investor writes

$$\begin{aligned} 0 = \frac{\partial J_a(\delta_0, \widehat{\beta}_0, \nu_0, 0)}{\partial x_0} &= \frac{(\theta\delta_0)^{1-\gamma}}{1-\gamma}\int_0^T \left[(1-\gamma)\left(\widehat{\beta}_0 - \gamma k\sigma^2(1+kx_0)\right)t + (1-\gamma)^2 x_0\nu_0 t^2\right] \\ &\times \exp\left[\kappa(x_0, \widehat{\beta}_0)t + \frac{(1-\gamma)^2}{2}x_0^2\nu_0 t^2\right] dt. \end{aligned} \tag{89}$$

Define the functions G and H as

$$G(\widehat{\beta}_t, \nu_t, t) \equiv \int_t^T (s-t)\exp\left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2}x_0^2\nu_t(s-t)^2\right] ds \tag{90}$$

$$H(\widehat{\beta}_t, \nu_t, t) \equiv \int_t^T (s-t)^2\exp\left[\kappa(x_0, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2}x_0^2\nu_t(s-t)^2\right] ds. \tag{91}$$

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

$$\widehat{\beta}_0\mathbb{D}_0 - \gamma k\sigma^2(1+kx_0)\mathbb{D}_0 + (1-\gamma)x_0\nu_0\mathbb{C}_0 = 0, \tag{92}$$

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

$$\mathbb{D}_t \equiv \frac{G(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)} \tag{93}$$

$$\mathbb{C}_t \equiv \frac{H(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)}. \tag{94}$$

The value $\mathbb{D}_t$ represents the weighted average maturity (i.e., the *equity duration*), whereas the value $\mathbb{C}_t$ represents the weighted average squared maturity (i.e., the *equity convexity*). Both $\mathbb{D}_t$ and $\mathbb{C}_t$ ar positive. Solving for x_0 in (92) yields (12). □

A.1.6 Proof of Proposition 3 (Experimentation under Competition)

Agent i faces the following choice in terms of consumption share and output:

$$\theta_i(x_i, \tilde{x}_C) = \frac{1}{N} e^{-\varphi + x_i(w - \tilde{x}_C)} \tag{95}$$

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \widehat{\beta}_t \tilde{x}_C)dt + \sigma(1 + k\tilde{x}_C)d\widehat{W}_t. \tag{96}$$

The expected lifetime utility of consumption for agent i at time t is

$$J_i(\delta_t, \widehat{\beta}_t, \nu_t, t) = \mathbb{E}_t\left[\int_t^T e^{-\rho s} \frac{\left(\frac{1}{N} e^{-\varphi + x_i(w - \tilde{x}_C)}\delta_s\right)^{1-\gamma}}{1-\gamma} ds\right] \tag{97}$$

$$= \frac{\left(\frac{1}{N} e^{-\varphi + x_i(w - \tilde{x}_C)}\right)^{1-\gamma}}{1-\gamma} \int_t^T e^{-\rho s} \mathbb{E}_t\left[\delta_s^{1-\gamma}\right] ds \tag{98}$$

$$= \underbrace{\frac{e^{-\rho t}\left(\frac{1}{N} e^{-\varphi + x_i(w - \tilde{x}_C)}\delta_t\right)^{1-\gamma}}{1-\gamma}}_{\equiv f(x_i)} F(\widehat{\beta}_t, \nu_t, t), \tag{99}$$

where $F(\cdot)$ is defined as before, except that now it is a function of $\tilde{x}_C$ instead of x_0. We prove first the following result.

Proposition 12 *At $t = 0$, the function $J_i(\delta_0, \widehat{\beta}_0, \nu_0, 0)$ admits a unique maximum in x_i.*

Proof *Write the function $J_i(\delta_0, \widehat{\beta}_0, \nu_0, 0)$ as the product of a negative constant and another function $z(x_i)$:*

$$J_i(\delta_t, \widehat{\beta}_t, \nu_t, t) = \underbrace{\frac{e^{-\rho t}\left(\frac{1}{N}\delta_t\right)^{1-\gamma} e^{(1-\gamma)(-\varphi)}}{1-\gamma}}_{<0} z(x_i), \tag{100}$$

where the function $z(x_i)$ is defined as

$$z(x_i) \equiv e^{(1-\gamma)x_i(w - \tilde{x}_C)} F(\widehat{\beta}_0, \nu_0, 0). \tag{101}$$

The function $z(x_i)$ is a product of two strictly positive functions (an exponential function and the price-dividend ratio).[22] *Consider a strictly monotonic transformation of $z(x_i)$:*

$$
\begin{aligned}
\ln z(x_i) &= (1-\gamma)x_i(w - x_i - \tilde{x}_C^-) + \ln F(\widehat{\beta}_0, \nu_0, 0) && (102)\\
&= \underbrace{(w - \tilde{x}_C^-)(1-\gamma)x_i + (\gamma-1)x_i^2}_{\equiv A(x_i)} + \underbrace{\ln F(\widehat{\beta}_0, \nu_0, 0)}_{\equiv B(x_i)}, && (103)
\end{aligned}
$$

where $\tilde{x}_C^- \equiv \tilde{x}_C - x_i$. The term $A(x_i)$ in (103) *is strictly convex when $\gamma > 1$. The term $B(x_i)$ represents the log price-dividend ratio. Close inspection of* (81)-(82) *reveals that the price-dividend ratio at time $t = 0$ has this particular form:*

$$
F(\widehat{\beta}_0, \nu_0, 0) = \int_0^T y(\tilde{x}_C, s)ds. \tag{104}
$$

It can be quickly verified that $y(\tilde{x}_C, s)$ is log-convex in x_i. This implies that the price-dividend ratio $F(\widehat{\beta}_0, \nu_0, 0)$ is log-convex in x_i (Boyd and Vandenberghe, 2004, p. 105). It then follows that the term $B(x_i)$ in (103) *is convex. Thus, $\ln z(x_i)$ is strictly convex and therefore the function $z(x_i)$ admits a unique minimum. This implies that $J_i(\delta_t, \widehat{\beta}_t, \nu_t, t)$ admits a unique maximum.* □

Starting from (99), write the first order condition for agent i at time $t = 0$:

$$
\begin{aligned}
0 &= \frac{\partial f(x_i)}{\partial x_i} F(\widehat{\beta}_0, \nu_0, 0) + f(x_i)\frac{\partial F(\widehat{\beta}_0, \nu_0, 0)}{\partial x_i} && (105)\\
&= (1-\gamma)(w - x_i - \tilde{x}_C) f(x_i) F(\widehat{\beta}_0, \nu_0, 0) + f(x_i)\frac{\partial F(\widehat{\beta}_0, \nu_0, 0)}{\partial x_i}, && (106)
\end{aligned}
$$

and replace

$$
\begin{aligned}
\frac{\partial F(\widehat{\beta}_0, \nu_0, 0)}{\partial x_i} &= \frac{\partial \tilde{x}_C}{\partial x_i}\frac{\partial F(\widehat{\beta}_0, \nu_0, 0)}{\partial \tilde{x}_C} && (107)\\
&= (1-\gamma)\left[\widehat{\beta}_0 - \gamma k \sigma^2 (1 + k\tilde{x}_C)\right] G(\widehat{\beta}_0, \nu_0, 0) + \tilde{x}_C (1-\gamma)^2 \nu_0 H(\widehat{\beta}_0, \nu_0, 0). && (108)
\end{aligned}
$$

After dividing by $(1-\gamma)f(x_i)F(\widehat{\beta}_0, \nu_0, 0)$, the first order condition for agent i at time $t = 0$ becomes:

$$
0 = (w - x_i - \tilde{x}_C) + \left[\widehat{\beta}_0 - \gamma k \sigma^2 (1 + k\tilde{x}_C)\right] \mathbb{D}_0 + \tilde{x}_C(1-\gamma)\mathbb{C}_0. \tag{109}
$$

The first term on the right hand side is the marginal benefit of experimentation arising from an increased consumption share. In the case of a monopolist, this term becomes $(w - 2x_i)$. Therefore, if $w > 2x_S^*$, then we are guaranteed that a monpolist will invest more than in the socially optimal case. Furthermore, it can be verified from Proposition 2 that the highest level of socially optimal experimentation is obtained in the log-utility case and equals

[22] Although both functions are convex, this does not guarantee that their product is convex.

$(\widehat{\beta}_0 - k\sigma^2)/(k^2\sigma^2)$. Thus, a sufficient condition necessary to obtain more experimentation in the monopolistic case is Eq. (14) in the text:

$$w > \frac{2(\widehat{\beta}_0 - k\sigma^2)}{k^2\sigma^2}. \tag{110}$$

Solving for x_i in Eq. (109) yields

$$x_i = \left[\widehat{\beta}_0 - \gamma k\sigma^2(1 + k\tilde{x}_C)\right]\mathbb{D}_0 + (w - \tilde{x}_C) - \tilde{x}_C(\gamma - 1)\nu_0\mathbb{C}_0, \quad \forall i = 1, ..., N. \tag{111}$$

This shows that agents' best response functions are symmetric. Aggregating across agents and dividing by N yields an equation in the total experimentation in the economy, $\tilde{x}_C$:

$$\frac{1}{N}\tilde{x}_C = \left[\widehat{\beta}_0 - \gamma k\sigma^2(1 + k\tilde{x}_C)\right]\mathbb{D}_0 + (w - \tilde{x}_C) - \tilde{x}_C(\gamma - 1)\nu_0\mathbb{C}_0, \tag{112}$$

with solution

$$\tilde{x}_C^* = \frac{\widehat{\beta}_0\mathbb{D}_0 - \gamma k\sigma^2\mathbb{D}_0 + w}{\gamma k^2\sigma^2\mathbb{D}_0 + (\gamma - 1)\nu_0\mathbb{C}_0 + \frac{N+1}{N}}. \tag{113}$$

This is a fixed point problem ($\mathbb{D}_0$ and $\mathbb{C}_0$ depend on $\tilde{x}_C^*$). The function on the right hand side is continuous and starts from a strictly positive value at $\tilde{x}_C = 0$. The solution is obtained when this function crosses the 45-degree line. The fixed point is unique, since we have proven in Proposition 12 that the function $J_i(\delta_t, \widehat{\beta}_t, \nu_t, t)$ admits a unique maximum and x_i is unique $\forall i = 1, ..., N$. To prove that $\tilde{x}_C^*$ is strictly increasing in N, consider two consecutive cases, N and $N + 1$:

$$\tilde{x}_C^*(N) = \frac{\widehat{\beta}_0\mathbb{D}_0 - \gamma k\sigma^2\mathbb{D}_0 + w}{\gamma k^2\sigma^2\mathbb{D}_0 + (\gamma - 1)\nu_0\mathbb{C}_0 + \frac{N+1}{N}} \tag{114}$$

$$\tilde{x}_C^*(N + 1) = \frac{\widehat{\beta}_0\mathbb{D}_0 - \gamma k\sigma^2\mathbb{D}_0 + w}{\gamma k^2\sigma^2\mathbb{D}_0 + (\gamma - 1)\nu_0\mathbb{C}_0 + \frac{N+1}{N} - \frac{1}{N(N+1)}}. \tag{115}$$

The two expressions on the right hand side in (114)-(115) are both functions of the experimentation level x. Naming these functions $g(x)$ and and $h(x)$ respectively, it is clear that, for any given experimentation level $x \geq 0$:

$$g(x) < h(x). \tag{116}$$

This inequality is easily verified, since for a given experimentation level x the duration $\mathbb{D}_0$ and the convexity $\mathbb{C}_0$ are the same in both functions. It then follows that the point at which the function $h(\cdot)$ crosses the 45-degree line is strictly higher than the point at which the function $g(\cdot)$ crosses the 45-degree line (see Figure 1). Since the equilibrium is unique, the quantity $\tilde{x}_C^*$ is strictly increasing in N. This completes the proof of Proposition 3. □

A.2 Implications for Asset Prices

A.2.1 Proof of Proposition 4 (Stochastic Discount Factor)

The proof follows standard results in asset pricing (Duffie, 2010). Assuming time-additive expected utility, we can define a stochastic discount factor from the optimal consumption plan of any individual as

$$\xi_t = e^{-\rho t} \frac{u'(\delta_t)}{u'(\delta_0)}. \tag{117}$$

Note that in our case agents consume fixed shares of the aggregate output and observe the economy under the same probability measure. Given the CRRA assumption, the dynamics of the stochastic discount factor can then be expressed as

$$\frac{d\xi_t}{\xi_t} = -\left[\rho + \gamma(\bar{f} + x_0\widehat{\beta}_t) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kx_0)^2\right]dt - \gamma\sigma(1+kx_0)d\widehat{W}_t. \tag{118}$$

The continuously compounded risk-free rate is the negative of the drift of the stochastic discount factor, whereas the market price of risk process is the negative of the diffusion of the stochastic discount factor. This yields (19) and (20) in Proposition 4. □

A.2.2 Proof of Proposition 5 (Asset Prices)

Recall that the state variables in this economy evolve according to (5)-(7). The equilibrium price of the risky asset is

$$P_t = \frac{1}{\xi_t}\mathbb{E}_t\left[\int_t^T \xi_s \delta_s ds\right] = \delta_t^\gamma \int_t^T e^{-\rho(s-t)}\mathbb{E}_t[\delta_s^{1-\gamma}]ds. \tag{119}$$

Using Proposition 11, we obtain

$$P_t = \frac{1-\gamma}{\theta^{1-\gamma}} e^{\rho t}\delta_t^\gamma J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) = \delta_t F(\widehat{\beta}_t, \nu_t, t), \tag{120}$$

which proves that indeed $F(\widehat{\beta}_t, \nu_t, t)$ is the price-dividend ratio. We then obtain F_t, F_β, F_ν, and $F_{\beta\beta}$, where $F(\widehat{\beta}_t, \nu_t, t)$ is defined in (81) and $\kappa(x_0, \widehat{\beta}_t)$ is defined in (82):

$$F_t \equiv \frac{\partial F(\widehat{\beta}_t, \nu_t, t)}{\partial t} = -1 - \kappa(x_0, \widehat{\beta}_t)F(\widehat{\beta}_t, \nu_t, t) - (1-\gamma)^2 x_0^2 \nu_t G(\widehat{\beta}_t, \nu_t, t) \tag{121}$$

$$F_\beta \equiv \frac{\partial F(\widehat{\beta}_t, \nu_t, t)}{\partial \widehat{\beta}_t} = (1-\gamma)x_0 G(\widehat{\beta}_t, \nu_t, t) \tag{122}$$

$$F_{\beta\beta} \equiv \frac{\partial^2 F(\widehat{\beta}_t, \nu_t, t)}{\partial \widehat{\beta}_t^2} = (1-\gamma)^2 x_0^2 H(\widehat{\beta}_t, \nu_t, t) \tag{123}$$

$$F_\nu \equiv \frac{\partial F(\widehat{\beta}_t, \nu_t, t)}{\partial \nu_t} = \frac{(1-\gamma)^2}{2} x_0^2 H(\widehat{\beta}_t, \nu_t, t), \tag{124}$$

where $G(\widehat{\beta}_t, \nu_t, t)$ and $H(\widehat{\beta}_t, \nu_t, t)$ are defined in (90)-(91). Apply Ito's formula to P_t:

$$dP_t = \delta_t F \frac{d\delta_t}{\delta_t} + \delta_t F_\beta d\widehat{\beta}_t + \delta_t F_\nu d\nu_t + \delta_t F_t dt + \frac{1}{2}\left[\delta_t F_{\beta\beta}(d\widehat{\beta}_t)^2 + 2F_\beta (d\delta_t)(d\widehat{\beta}_t)\right], \quad (125)$$

to obtain

$$\frac{dP_t}{P_t} = \mu_{P,t}dt + \sigma_{P,t}d\widehat{W}_t, \quad (126)$$

with

$$\mu_{P,t} \equiv \bar{f} + \widehat{\beta}_t x_0 - \kappa(x_0, \widehat{\beta}_t) - \frac{1}{F(\widehat{\beta}_t, \nu_t, t)} + \gamma(1-\gamma)x_0^2 \nu_t \frac{G(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)} \quad (127)$$

$$\sigma_{P,t} \equiv \sigma(1 + kx_0)\left(1 + (1-\gamma)\frac{x_0^2 \nu_t}{\sigma^2(1+kx_0)^2}\frac{G(\widehat{\beta}_t, \nu_t, t)}{F(\widehat{\beta}_t, \nu_t, t)}\right). \quad (128)$$

To obtain the risk premium as in (24), multiply the market price of risk, $\theta_t = \gamma\sigma(1 + kx_0)$ (Proposition 4, Eq. 20), with the diffusion of stock returns. □

A.3 Dynamic Experimentation

A.3.1 Proof of Proposition 6 (Optimal Dynamic Experimentation)

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

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \widehat{\beta}_t x_t)dt + \sigma(1 + kx_t)d\widehat{W}_t, \quad (129)$$

with

$$d\widehat{\beta}_t = \frac{x_t}{\sigma(1+kx_t)}\nu_t d\widehat{W}_t \quad (130)$$

$$d\nu_t = -\frac{x_t^2}{\sigma^2(1+kx_t)^2}\nu_t^2 dt. \quad (131)$$

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

$$0 = \max_x \left[\mathcal{D}J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) + e^{-\rho t}\frac{(\theta c_t)^{1-\gamma}}{1-\gamma}\right], \quad (132)$$

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

$$x_t \geq 0, \forall t. \quad (133)$$

In equilibrium consumption equals total output and therefore

$$
\begin{aligned}
0 = \max_x \frac{e^{-\rho t}\delta_t^{1-\gamma}}{1-\gamma} + J_{a,t} + \delta_t(\bar{f} + \widehat{\beta}_t x)J_{a,\delta} \\
+ \frac{x^2\nu_t^2}{2\sigma^2(1+kx)^2}(J_{a,\beta\beta} - 2J_{a,\nu}) + \frac{1}{2}\delta_t^2\sigma^2(1+kx)^2 J_{a,\delta\delta} + \delta_t\nu_t x J_{a,\delta\beta}.
\end{aligned} \tag{134}
$$

With CRRA utility, we make the customary assumption

$$
J_a(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t}\frac{(\theta\delta_t)^{1-\gamma}}{1-\gamma}F(\widehat{\beta}_t, \nu_t, t), \tag{135}
$$

and thus the PDE (134) becomes

$$
0 = \max_x F_t + \kappa(x, \widehat{\beta}_t)F + \frac{x^2\nu_t^2}{2\sigma^2(1+kx)^2}(F_{\beta\beta} - 2F_\nu) + (1-\gamma)x\nu_t F_\beta + 1, \tag{136}
$$

with boundary condition $F(\widehat{\beta}_T, \nu_T, T) = 0$ and $\kappa(x, \widehat{\beta}_t)$ defined as (same as in Eq. 22):

$$
\kappa(x, \widehat{\beta}_t) \equiv -\rho - (\gamma - 1)(\bar{f} + x\widehat{\beta}_t) + \frac{\gamma(\gamma-1)}{2}\sigma^2(1+kx)^2. \tag{137}
$$

The first order condition for x is

$$
\begin{aligned}
0 &= \kappa' F + \frac{x\nu_t^2}{\sigma^2(1+kx)^3}(F_{\beta\beta} - 2F_\nu) + (1-\gamma)\nu_t F_\beta && (138) \\
&= (\gamma - 1)\left[k(1+kx)\gamma\sigma^2 - \widehat{\beta}_t\right]F + \frac{x\nu_t^2}{\sigma^2(1+kx)^3}(F_{\beta\beta} - 2F_\nu) + (1-\gamma)\nu_t F_\beta. && (139)
\end{aligned}
$$

This is a quartic equation in x. Re-arranging yields Eq. (34) in Proposition 6. □

A.3.2 Proof of Proposition 7 (AP with Dynamic Experimentation)

The stochastic discount factor follows

$$
\frac{d\xi_t}{\xi_t} = -\left[\rho + \gamma(\bar{f} + x_t^*\widehat{\beta}_t) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kx_t^*)^2\right]dt - \gamma\sigma(1+kx_t^*)d\widehat{W}_t^\delta, \tag{140}
$$

and thus the risk-free rate and the market price of risk from (35)-(36) follow. The stock price at time t is $P_t = \delta_t F(\widehat{\beta}_t, \nu_t, t)$. The main 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

$$
\frac{dP_t}{P_t} = \left[\bar{f} + x_t^*\widehat{\beta}_t - \kappa(x_t^*, \widehat{\beta}_t) - \frac{1}{F} + \gamma x_t^*\nu_t\frac{F_\beta}{F}\right]dt + \sigma(1+kx_t^*)\left[1 + \frac{x_t^*\nu_t}{\sigma^2(1+kx_t^*)^2}\frac{F_\beta}{F}\right]d\widehat{W}_t, \tag{141}
$$

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

$$RP_t = \gamma\sigma^2(1+kx_t^*)^2\left[1+\frac{x_t^*\nu_t}{\sigma^2(1+kx_t^*)^2}\frac{F_\beta}{F}\right]. \tag{142}$$

The volatility and the risk premium depend on F_β/F. Using (135), at the optimal experimentation level x_t^*,

$$F(\widehat{\beta}_t,\nu_t,t) = \mathbb{E}_t\left[\int_t^T \frac{e^{-\rho(s-t)}\delta_s^{1-\gamma}}{\delta_t^{1-\gamma}}ds\right] = \int_t^T \frac{e^{-\rho(s-t)}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]}{\delta_t^{1-\gamma}}ds, \tag{143}$$

and

$$\mathbb{E}_t\left[\delta_s^{1-\gamma}\right] = e^{(1-\gamma)\mathbb{E}_t[\ln\delta_s]+\frac{(1-\gamma)^2}{2}\mathrm{Var}_t[\ln\delta_s]}. \tag{144}$$

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

$$\mathbb{E}_t\left[\ln\delta_s\right] = \ln\delta_t + \mathbb{E}_t\left[\int_t^s\left(\bar{f}+\widehat{\beta}_\tau x_\tau^* - \frac{1}{2}\sigma^2(1+kx_\tau^*)^2\right)d\tau\right], \tag{145}$$

and

$$\frac{\partial\mathbb{E}_t\left[\ln\delta_s\right]}{\partial\widehat{\beta}_t} = \mathbb{E}_t\left[\int_t^s x_\tau^* d\tau\right]. \tag{146}$$

Therefore

$$\begin{aligned}
\frac{F_\beta}{F} &= \frac{(1-\gamma)\int_t^T \mathbb{E}_t\left[\int_t^s x_\tau^* d\tau\right]\frac{e^{-\rho(s-t)}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]}{\delta_t^{1-\gamma}}ds}{\int_t^T\frac{e^{-\rho(s-t)}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]}{\delta_t^{1-\gamma}}ds} && (147)\\
&= (1-\gamma)x_t^*\frac{\int_t^T \mathbb{E}_t\left[\int_t^s \frac{x_\tau^*}{x_t^*}d\tau\right]\frac{e^{-\rho(s-t)}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]}{\delta_t^{1-\gamma}}ds}{\int_t^T\frac{e^{-\rho(s-t)}\mathbb{E}_t\left[\delta_s^{1-\gamma}\right]}{\delta_t^{1-\gamma}}ds} && (148)\\
&= (1-\gamma)x_t^*\widetilde{\mathbb{D}}_t, && (149)
\end{aligned}$$

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 (38) and the risk premium (37) follow by replacing (149) in (141)-(142). □

References

Abreu, D. and M. K. Brunnermeier (2003). Bubbles and crashes. *Econometrica 71*(1), 173–204.

Acemoglu, D. and P. Restrepo (2017). Robots and jobs: Evidence from us labor markets. *Working Paper*.

Aghion, P., U. Akcigit, A. Bergeaud, R. Blundell, and D. Hémous (2015). Innovation and top income inequality. Technical report, National Bureau of Economic Research.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted-u relationship. *The Quarterly Journal of Economics 120*(2), 701–728.

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991). Optimal learning by experimentation. *The review of economic studies 58*(4), 621–654.

Aghion, P. and R. Griffith (2008). *Competition and growth: reconciling theory and evidence.* MIT press.

Aghion, P. and P. Howitt (1992). A model of growth through creative destruction. *Econometrica 60*(2), 323–351.

Akcigit, U., J. Grigsby, and T. Nicholas (2017). The rise of american ingenuity: Innovation and inventors of the golden age.

Arrow, K. (1962). The economic implications of learning by doing. *Review of Economic Studies*, 131.

Bampoky, C., J. E. Prieger, L. R. Blanco, and A. Liu (2016). Economic growth and the optimal level of entrepreneurship. *World Development 82*, 95–109.

Boldrin, M. and D. K. Levine (2008). Perfectly competitive innovation. *Journal of Monetary Economics 55*(3), 435–453.

Boyd, S. and L. Vandenberghe (2004). *Convex optimization.* Cambridge university press.

Brennan, M. J. (1998a). The role of learning in dynamic portfolio decisions. *European Finance Review 1*(3), 295–306.

Brennan, M. J. (1998b). Stripping the s&p 500 index. *Financial Analysts Journal*, 12–22.

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.

Bryson, A. E. and Y.-C. Ho (1975). *Applied optimal control: optimization, estimation and control.* Taylor and Francis.

Carree, M. A. and A. R. Thurik (2005). The impact of entrepreneurship on economic growth. *Handbook of entrepreneurship research*, 437–471.

Cole, H. L. and M. Obstfeld (1991). Commodity trade and international risk sharing: How much do financial markets matter? *Journal of monetary economics 28*(1), 3–24.

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.

Duffie, D., D. Filipović, and W. Schachermayer (2003). Affine processes and applications in finance. *Annals of applied probability*, 984–1053.

Fama, E. F. (2014). Two pillars of asset pricing. *The American Economic Review 104*(6), 1467–1485.

Filipović, D. (2005). Time-inhomogeneous affine processes. *Stochastic Processes and their Applications 115*(4), 639–659.

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.

Giglio, S., M. Maggiori, and J. Stroebel (2016). No-bubble condition: Model-free tests in housing markets. *Econometrica 84*(3), 1047–1091.

Gort, M. and S. Klepper (1982). Time paths in the diffusion of product innovations. *The economic journal 92*(367), 630–653.

Grammig, J. and S. Jank (2015). Creative destruction and asset prices. *Journal of Financial and Quantitative Analysis*.

Greenwood, R., A. Shleifer, and Y. You (2017). Bubbles for fama. Technical report, National Bureau of Economic Research.

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.

Helpman, E. and A. Razin (2014). *A theory of international trade under uncertainty.* Academic Press.

Johnson, T. C. (2007). Optimal learning and new technology bubbles. *Journal of Monetary Economics 54*(8), 2486–2511.

Komlos, J. (2014). Has creative destruction become more destructive? Technical report, National Bureau of Economic Research.

Kung, H. and L. Schmid (2015). Innovation, growth, and asset prices. *The Journal of Finance 70*(3), 1001–1037.

Liptser, R. and A. Shiryayev (1977). Statistics of random processes. Springer-Verlag, Berlin, New York.

Loury, G. C. (1979). Market structure and innovation. *The quarterly journal of economics*, 395–410.

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of monetary Economics 15*(2), 145–161.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, 867–887.

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.

Pavlova, A. and R. Rigobon (2007). Asset prices and exchange rates. *Review of Financial Studies 20*(4), 1139–1180.

Reinganum, J. F. (1983). Uncertain innovation and the persistence of monopoly. *The American Economic Review 73*(4), 741–748.

Reinganum, J. F. (1985). Innovation and industry evolution. *The Quarterly Journal of Economics 100*(1), 81–99.

Rob, R. (1991). Learning and capacity expansion under demand uncertainty. *The Review of Economic Studies 58*(4), 655–675.

Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy 98*(5, Part 2), S71–S102.

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.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy.* New York: Harper and Brothers.

Shiller, R. J. (2015). *Irrational exuberance.* Princeton university press.

Swan, P. L. (1970). Market structure and technological progress: The influence of monopoly on product innovation. *The Quarterly Journal of Economics 84*(4), 627–638.

Williams, J. T. (1977). Capital asset prices with heterogeneous beliefs. *Journal of Financial Economics 5*(2), 219–239.

Williamson, O. E. (1965). Innovation and market structure. *Journal of political economy 73*(1), 67–73.

Witt, U. (1996). Innovations, externalities and the problem of economic progress. *Public Choice 89*(1), 113–130.

Xiong, W. (2013). Bubbles, crises, and heterogeneous beliefs. *Handbook on Systemic Risk*, 663–713.

Zapatero, F. (1995). Equilibrium asset prices and exchange rates. *Journal of Economic Dynamics and Control 19*(4), 787–811.

Ziegler, A. C. (2003). *Incomplete Information and Heterogeneous Beliefs in Continuous-time Finance.* Springer-Verlag Berlin Heidelberg.