Schumpeterian Competition in a Lucas Economy

Daniel Andrei* Bruce I. Carlin†

January 17, 2023

Abstract

We model a rent-seeking game where agents experiment with a new technology and compete for claims to a consumption stream. We characterize how creative destruction affects risk, wealth, and asset prices. Competition not only imposes excessive disruption risk on existing assets and higher technological uncertainty, it also increases the wealth duration (the weighted-average maturity of wealth). Because of hedging motives, a complementarity between wealth duration and technological uncertainty decreases systematic risk. If competition is sufficiently intense, a negative risk premium may arise. The model generates price paths consistent with boom-bust patterns and transient episodes of negative expected excess returns. We show that Schumpeterian competition may worsen income inequality.

Keywords: Schumpeterian Competition; Experimentation; Tullock contests; Creative Destruction; Return Predictability; Income Inequality.

JEL classification: G11; G12; G14; L13; O33.

We are grateful to Kerry Back, Tony Bernardo, Hank Bessembinder, Michael Brennan, Philip Bond, Murray Carlson, Ian Dew-Becker, Winston Dou (Red Rock discussant), Lorenzo Garlappi, Chris Hrdlicka, Mark Huson, Leonid Kogan, Francis Longstaff, Alexander Ober, Stavros Panageas, Lubos Pastor, Jeff Pontiff, Ed Rice, Andrei Shleifer, Tyler Shumway, Avanidhar Subrahmanyam, Johan Walden, Mark Westerfield, and participants in research seminars at UCLA, the 2016 HEC-McGill Winter Finance Conference, the University of Washington, UC Irvine, the 2017 UBC Summer Finance Meetings, the 2017 Red Rock Finance Conference, McGill, and Collegio Carlo Alberto for their comments and input. We thank Nan Ma, Alexander Ober, Wenyun Qin, and Yujin Yang for outstanding research assistance.

*Corresponding author. McGill University, Desautels Faculty of Management, 1001 Sherbrooke Street West, Office 549, Montréal, Québec H3A 1G5, danielandrei.info, daniel.andrei@mcgill.ca.

†Jones School of Business, Rice University, 1900 Rice Blvd., Houston, TX 77005, carlin@rice.edu.

1 Introduction

We characterize the effect of Schumpeterian competition on the aggregate consumption stream in a standard Lucas endowment economy. When paradigm shifts occur, agents fight for market share, which is a claim to both current and future cash flows, and includes real options such as the option to expand, flexibility, and synergies. We investigate the consequences that this type of competition has for the structure of the consumption stream, risk in the economy, and asset pricing dynamics.

Surely, aggregate consumption changes when innovation arises and agents compete for monopoly rents. The relationship between creative destruction and rent-seeking has been studied in a variety of contexts (e.g., Reinganum, 1983; Aghion and Howitt, 1992; Baye and Hoppe, 2003).[1] Though, it still remains unclear whether creative destruction is socially beneficial (e.g., Witt, 1996; Aghion, Akcigit, Bergeaud, Blundell, and Hémous, 2015; Komlos, 2014), and whether it worsens income inequality (e.g., Jones and Kim, 2018; Gabaix, Lasry, Lions, and Moll, 2016).

Beyond the potential gains that innovation promises, Schumpeter also stressed that creative destruction "takes considerable time in revealing its true features and ultimate effects" (Schumpeter, 1942). Some of this risk is due to the eventual fate of the innovation (*technological uncertainty*). But, creative destruction may also involve spillovers and irreversible changes in systematic risk (Pastor and Veronesi, 2009) because the fate of existing assets becomes uncertain: some assets face *disruption risk*, while others enjoy new growth options and complementarities. So, a paradigm shift may indeed change both the expected value and variance of future consumption.

To fix ideas, consider the recent widespread adoption of Zoom Video Communications. This not only displaced technologies like Skype and Face Time, it impacted many other dimensions of the economy. Remote meetings have made firms more efficient in connecting global employees and they have improved the ability of researchers from different academic institutions to conduct research together. At the same time, the future of corporate real estate has become more uncertain as businesses are re-assessing their needs for office space. Residential real estate has become more risky as well: people are migrating away

[1] See Kamien and Schwartz (1975) for an excellent survey of the literature, and Kamien and Schwartz (1982) and Nelson and Winter (2009) for textbook treatments. See also Swan (1970), Loury (1979), Futia (1980), Fudenberg and Tirole (1984), Reinganum (1985), D'Aspremont and Jacquemin (1988),Reinganum (1989), Klepper (1996), Witt (1996), Chung (1996), Aghion, Bloom, Blundell, Griffith, and Howitt (2005), Carree and Thurik (2005), Boldrin and Levine (2008), Aghion and Griffith (2008), Komlos (2014),Bampoky, Prieger, Blanco, and Liu (2016), Acemoglu and Restrepo (2020), Akcigit, Grigsby, and Nicholas (2017), and Jones and Kim (2018).

from city centers because they are able to work remotely. The evolution of this paradigm shift has also likely affected travel, hotel usage, the census of state populations (which have political implications), and labor markets themselves. So, while the direct impact of this creative destruction is on obvious substitutes, spillovers make many other existing assets more uncertain as well.[2]

We study a Tullock contest (Tullock, 1980) in which agents make simultaneous choices about how much to experiment with a new technology. Agents compete for market share, which is a proportion of an aggregate consumption stream that endogenously results from innovation (Hirshleifer, 1989). In our model, however, creative destruction by the agents has two effects on the aggregate consumption stream: it changes the growth rate of consumption—hopefully for the better—and it amplifies the magnitude of its diffusion. The former characterizes technological uncertainty and the latter captures disruption risk.

We solve for a first-best benchmark. The socially optimal amount of experimentation results from the classic tradeoff between maximizing the expected value of future consumption and minimizing its future volatility. Naturally, experimentation is higher with less technological uncertainty, a higher expected benefit, lower risk aversion, and less spillover to existing assets through disruption risk.

Rent-seeking competition fundamentally alters this tradeoff and catalyzes inefficient investments. When agents fight for a slice of the pie, aggregate experimentation is higher than the social optimum and grows monotonically with competition. Moreover, if the experimental good offers some degree of technological diversification (Koren and Tenreyro, 2013), then the agents have higher incentives to over-invest, which again increases experimentation. But, the agents do not internalize their effect on risk. As such, even though excessive experimentation allows for faster learning about the new technology, it amplifies the impact of technological uncertainty and causes more disruption of existing assets.

Disruption risk permanently increases both the equilibrium stock market volatility and the risk premium, and competition amplifies this. However, technological uncertainty has a negative and dynamic effect on volatility and the risk premium. This effect arises in general equilibrium because of a hedging demand from agents, who insulate themselves against the adverse effect of technological uncertainty.

[2]Other examples include: the rise of ride-sharing companies such as Uber and Lyft, which have disrupted the traditional taxi industry; the growth of *e*-commerce and online retail companies such as Amazon, which have led to the decline of brick-and-mortar retail stores; the emergence of renewable energy sources such as solar and wind power, which have disrupted the traditional fossil fuel industry; the advent of streaming services such as Netflix and Spotify, which have disrupted the traditional cable and music industries; and the rise of social media platforms such as Facebook and Instagram, which have disrupted the traditional advertising industry.

From the model, we learn that experimentation affects systematic risk by changing the term structure of wealth. The *wealth duration*, which measures the weighted average maturity of the price of a claim to aggregate consumption, describes the risks that agents bear based on the timing of consumption from the dividend stream. Experimentation tends to increase wealth duration, and this amplifies the effect of technological uncertainty and increases agents' hedging motives. This complementarity may even cause the risk premium to be negative. But this effect is transient: as learning progresses and technological uncertainty drops, both systematic risk and the risk premium rise, which is a result of persistent disruption risk.

The price-dividend ratio increases with experimentation but decreases as uncertainty about the new technology is resolved through learning. Higher rent-seeking and increased competition make the occurrence of boom-bust episodes more likely, which does not occur in the social optimum. To test this empirically, we discretize our continuous-time setup, run simulations of the model, and conduct predictive regressions. Our results show that high competition causes the price-dividend ratio to negatively predict future returns. This not only lowers predicted expected excess returns, but it may also result in negative expected excess returns when competition is particularly intense. This pattern aligns with the empirical characterization of booms and busts in Hoberg and Phillips (2020) and the theoretical implications of DeMarzo, Kaniel, and Kremer (2007).

We investigate extensions of the model, several of which amplify our baseline results. For instance, when the new technology is subject to obsolescence (Aghion and Howitt, 1992), Schumpeterian competition leads to a higher degree of over-experimentation than in our baseline setup. Additionally, we expand the representative agent model to include the option for agents to change their experimentation over time. This allows agents to abandon a technology if it is determined to be non-value enhancing. This option for abandonment increases experimentation, which in turn increases risk as agents may experiment more knowing that they have the option to back off.

To determine the effect of Schumpeterian competition on residual claimants in the economy, we incorporate passive agents into our model. These agents do not participate in competition for market share, but instead are residual claimants of the consumption stream. Our findings indicate that the welfare of passive agents decreases as competition increases. Not only do they receive a smaller portion of the consumption stream due to over-experimentation, but the aggregate value of the consumption stream also decreases. Thus, our model suggests that creative destruction may exacerbate income inequality, particularly with higher levels of competition (e.g., Jones and Kim, 2018). Finally, our analysis of a setting with recursive preferences reveals that learning by doing endogenously creates long-run risk (Bansal and Yaron, 2004), aligning with a growing literature that provides equilibrium foundations of long-run risk through innovation and R&D (Kung and Schmid, 2015; Kung, 2015).

Related Literature Our paper is closest to Pastor and Veronesi (2009). In their model, learning about the new technology takes place with an infinitesimally small consumption stream, so the new technology does not perturb the rest of the economy. Once the new technology becomes sufficiently promising, there is an instantaneous transition in which the new consumption stream displaces the status quo. Prices rise and then fall as risk changes from idiosyncratic to systematic. Our work complements theirs. First, we consider that learning by doing (Arrow, 1962) does disrupt the status quo and there is an observer effect. Second, because of our game-theoretic framework, we can consider the effect of competition on risk, wealth, and prices. Third, both technological uncertainty and systematic risk evolve continuously, so that we can evaluate their effects over the long term. Last, because we model game-theoretic claims to a consumption stream, we can show how the term structure of those claims changes and affects volatility and risk premia.

In a related paper, DeMarzo et al. (2007) develop a model of relative wealth concerns where risky technologies attract excessive investment that can be largely unprofitable, which they argue is consistent with a real investment bubble. In our setting with Schumpeterian competition, excess investment arises in a rent-seeking contest. So, we extend the analysis in DeMarzo et al. (2007) on two dimensions. First, we establish a novel link between excessive real investment and asset price bubbles. As over-investment magnifies technological uncertainty, it leads to price paths consistent with boom-bust patterns and transient episodes of negative expected excess returns. Second, our model shows how over-investment may worsen wealth inequality and expose agents who do not actively invest in the new technology to excessive risks.

Our study builds on the seminal contribution of Tullock (1980) and a number of papers that have explored various aspects of rent-seeking competitions.[3] We contribute to this literature by considering a game in which the aggregate prize changes non-linearly with effort. This is a key aspect of our analysis, whose outcome is that agents over-experiment despite the fact that the aggregate prize decreases, because their incentives are mainly

[3]See, among many others, D'Aspremont and Jacquemin (1988); Hirshleifer (1989); Alexeev and Leitzel (1996); Chung (1996); Baye and Hoppe (2003); Chowdhury and Sheremeta (2011).

focused on their own share of the prize. Furthermore, because we place the Tullock contest into a general equilibrium asset pricing setting with heterogeneous agents, we can consider the effect of competition on systematic risk, social welfare, and wealth inequality.

Our model relates to stochastic variants of Schumpeterian growth models (Barlevy, 2007; Kung, 2015; Kung and Schmid, 2015), which provide microfoundations for the way innovation impacts economic growth and its fluctuations. Our framework builds on these microfoundations and analyzes the effect of competition and technological uncertainty on risk and asset prices. The long term effects of the model provide a new perspective on how competition may drive secular trends related to innovation and risk premia. Recent studies have attributed slumping investment and innovation (Gutiérrez and Philippon, 2018) and rising risk premia in equity markets (Corhay, Kung, and Schmid, 2020) to declining competition, an observation that aligns well with the predictions of our model. Finally, through its effect on duration and risk, experimentation changes the term structure of risk and implies lower returns for long-duration assets, consistent with a recent literature that studies the risk premia of equity claims with different maturities (Lettau and Wachter, 2007, 2011; Weber, 2018; Van Binsbergen, 2020; Gonçalves, 2021).

The rest of the paper proceeds as follows. Section 2 poses the model, characterizes learning and socially optimal experimentation, and contrasts the equilibrium behavior of competitive agents to the social optimum. Section 3 characterizes the risks of Schumpeterian competition and its implications for asset prices and return predictability. Section 4 discusses various extensions of the model. Section 5 concludes. All proofs and additional calculations are in the Appendix.

2 Experimentation and Competition

The continuous-time model of this section builds on a pure exchange economy (Lucas, 1978) whose growth is disrupted by the advent of a new technology. We study agents' decisions to deploy the new technology and reap the associated rewards. For expositional purposes, we keep the model as simple as possible. We discuss the model's assumptions in Section 4, where we analyze several extensions and alternative specifications.

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

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

where $\bar{f}$ and σ are known constants and W_t^S is a one-dimensional standard Brownian motion. For now, we assume that the economy is populated by one representative agent. Later, we study an economy with multiple agents who compete for shares of the aggregate consumption stream. We also discuss the finite horizon assumption in Section 4.1, where we consider an extension of the model to the infinite horizon case.

The agent has isoelastic preferences over lifetime consumption,

$$\mathbb{E}_0 \left[\int_0^T e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} dt \right], \tag{2}$$

where γ is the coefficient of relative risk aversion, ρ is the time discount parameter, and $\mathbb{E}_0[\cdot]$ denotes the expectation under the filtration of the agent, to be defined below.

At time $t = 0$, the agent has the opportunity to allocate resources $X \geq 0$ to an *experimental asset*. This new asset is sufficiently important to affect the entire economy and to change aggregate consumption. The asset employs a new technology that has never been implemented on a large scale before. Its output follows the dynamic process

$$\frac{d\delta_t^E}{\delta_t^E} = \left(g - \frac{k^2\sigma^2}{2} X \right) dt + k\sigma\varphi dW_t^S + k\sigma\sqrt{1-\varphi^2} dW_t^E. \tag{3}$$

The parameter g in the drift of the process (3) is unknown at time $t = 0$ (all the other parameters are known). We refer to the uncertainty about g as *technological uncertainty.* The drift of δ_t^E increases in g and decreases linearly with X; that is, the expected output growth of the experimental asset has decreasing returns to scale, an assumption that will simplify our analysis, and that is likely to dampen our main results.

According to the diffusion terms in the output process of the experimental good, the new technology experiences the same shocks as the status quo, dW_t^S, but also new shocks generated by the independent Brownian motion W_t^E. The parameter $\varphi \in [-1, 1]$ regulates the exposure of the new technology to dW_t^E. If $\varphi < 1$, the shock dW_t^E provides technological diversification benefits (Koren and Tenreyro, 2013). Finally, the parameter $k > 0$ controls the output volatility of the new technology.

The agent commits to an experimentation level X that remains constant thereafter. As such, the agent chooses how far to open Pandora's box and lives with the consequences.[4] Allocating resources to the new asset has an opportunity cost: if $X > 0$, the consumption

[4] We discuss the case in which the agent can dynamically change experimentation in Section 4.3.

stream that results from the old asset grows at a slower pace:

$$\frac{d\delta_t^S}{\delta_t^S} = \left(\bar{f} - \frac{c}{2}X^2\right) dt + \sigma dW_t^S. \tag{4}$$

where $c > 0$. The cost parameter c in our model relates to the concept of creative destruction. For instance, in the endogenous growth model pioneered by Aghion and Howitt (1992), innovation is the driving force of growth, but it also poses a threat to the rents generated by current research, similar to our model, where investment in new technology redirects resources away from the existing asset, leading to the destruction of its future rents. Additionally, c can also be understood as the displacement cost that arises from transferring resources away from existing assets (Kogan, Papanikolaou, and Stoffman, 2020). Finally, this opportunity cost specification is both mathematically convenient and economically relevant, as argued by Lucas (1987) and Barlevy (2004), as reducing the growth rate of consumption entails large welfare costs.[5]

If $X = 0$, the economy remains in the status quo and only δ_t^S is consumed. When $X > 0$, the experimental asset produces a new consumer good. Thus, the agent's decision to allocate resources to a new asset and create a new consumer good embodies Schumpeter's *"fundamental impulse that sets and keeps the capitalist engine in motion"* (1942, p. 83): it yields a new variety in the aggregate consumption basket, which we define as

$$\delta_t \equiv \delta_t^S \left(\delta_t^E\right)^X. \tag{5}$$

In the new economy, the agent consumes a composite of the status-quo and experimental goods. Replacing the aggregate consumption (5) in (2) yields a Cobb-Douglas/isoelastic specification commonly encountered in the international economics and finance literature[6], where X is interpreted as a taste parameter. A similar interpretation can be made here. Referring to the aggregate consumption basket δ_t as the *numéraire* and normalizing its price to unity, optimal intratemporal choice (within period and across goods) implies that $X/(1+X)$ is the agent's expenditure share on the experimental good (see Appendix A.1). In other words, maximizing utility by choosing the amount of experimentation X is equiv-

[5]Our results go through even in the case $c = 0$. As we will elaborate below, experimentation itself generates costs by raising the future variance of aggregate consumption, which guarantees a solution to the agent's problem even when $c = 0$.

[6]See Helpman and Razin (1978, p. 101), Cole and Obstfeld (1991, p. 7), Zapatero (1995, p. 790), Pavlova and Rigobon (2007, p. 1144). For analytical convenience, the sum of the exponents in (5) is $1+X$ rather than being normalized to unity. Our results are preserved if we specify instead $\delta_t \equiv (\delta_t^S)^{1-X}(\delta_t^E)^X$. However, this alternative would impose the constraint $X \leq 1$ and would further insert a non linearity in (6) arising from the cost $cX^2/2$ that experimentation imposes on the existing asset in (4).

alent with maximizing utility by choosing the expenditure share $X/(1+X)$.[7]

There are numerous real-life examples of new technologies capturing a significant share of consumer expenditure. One notable example is the advent of smartphones, particularly those produced by Apple, which replaced more traditional mobile phones such as Nokia and BlackBerry. Another example is Tesla, which has captured a significant share of the car market, eroding the rents generated by sales of gas-powered cars. In our model, these effects on consumption shares are driven by the agent's choice of X.

The Cobb-Douglas specification (5) considerably simplifies the dynamics of *aggregate consumption* δ_t. Formally, Itô's lemma implies:

$$\frac{d\delta_t}{\delta_t} = \left(\bar{f} + \beta X - \frac{c}{2}X^2\right) dt + \sigma\sqrt{1 + k^2X^2 + 2k\varphi X} dW_t, \tag{6}$$

where the parameter β is defined as

$$\beta \equiv g + \frac{1}{2}(2\varphi - k)k\sigma^2, \tag{7}$$

and W_t is a Brownian motion correlated with both W_t^S and W_t^E (it is possible to keep track of W_t^S and W_t^E separately, but not necessary in what follows).

Investment in the new technology simultaneously affects the drift and diffusion coefficients of aggregate consumption growth. In the new drift coefficient, which becomes $(\bar{f} + \beta X - cX^2/2)$, the parameter $\beta \in \mathbb{R}$ is unknown and encompasses that innovation can both fuel growth when $\beta > cX/2$ and also be destructive when $\beta < cX/2$. Our new economy is thus one in which growth is endogenous and driven by technological change (Aghion and Howitt, 1992; Garleanu, Panageas, and Yu, 2012; Kung and Schmid, 2015; Kung, 2015). Importantly, in our model the true value of β is unknown. Thus, the agent does not exclude the possibility that the technology may have negative aggregate effects (Acemoglu and Restrepo, 2020).

Investment in the new technology yields the diffusion coefficient $\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}$. This captures the *disruption risk* that the new asset imposes on the economy. Unlike in Pastor and Veronesi (2009), here the new technology is sufficiently developed so that its adoption contributes to aggregate fluctuations, although it also provides diversification benefits when $\varphi < 1$. The literature recognizes that innovation creates systematic risk

[7] A small quantity δ_E would imply a large initial price for the experimental asset (Appendix A.1). However, good prices do not generally play a role with a Cobb-Douglas specification. Considering a constant elasticity of substitution (CES) function would have the potential benefit of giving good prices a role in the game, a matter which we leave for future research.

(Gârleanu, Kogan, and Panageas, 2012; Kung and Schmid, 2015). Our focus will be on the effect of competition on the risks created by innovation.

In summary, the agent's selection of the level of experimentation, represented by the variable X, has a multifaceted impact on the future trajectory of the economy. While it may facilitate faster economic growth due to the unknown parameter β, it also engenders increased output volatility, as denoted by the parameter k, and diverts resources from existing assets, as represented by the parameter c. These various roles of X align our model with stochastic Schumpeterian growth models (e.g., Kung, 2015; Kung and Schmid, 2015), which provide microeconomic foundations for how innovation affects both the drift and diffusion components of aggregate output. Our model captures these effects in a simplified form and concentrates on the agent's experimentation trade-off. The agent maximizes the welfare gains of accelerated growth, taking into account the welfare losses brought about by more significant aggregate fluctuations, a trade-off that has also been discussed in previous literature such as in Barlevy (2004, 2007). However, it is essential to note that our model also features an additional source of risk, which is not immediately apparent from (6). This source of risk is the technological uncertainty surrounding the actual value of β. Characterizing the agent's learning process is crucial in comprehending how technological uncertainty influences this trade-off.

Learning In this economy learning occurs *by doing*: the agent can evaluate the benefits of the new technology only if it is deployed in the real economy and $X > 0$.[8] This implies that the agent learns via a *field experiment*, rather than through an isolated, laboratory experiment. What we have in mind is that a controlled laboratory setting where economic agents learn about a new technology by implementing it at an infinitesimally small scale (Pastor and Veronesi, 2009) is not sufficient to assess the efficacy of innovation (e.g., a vaccine; ride-sharing technology).

Before proceeding, it should be noted that the unknown aspect of the new technology in our model is its drift. This assumption seems somewhat restrictive, since the riskiness of the new technology, or its *diffusion*, may also be unknown (e.g., crypto assets). However, the agent is able to observe without error the unknown diffusion parameter of any continuous-time process from its quadratic variation by increasing the sampling frequency

[8] This is distinct from what is typically modeled in the literature with incomplete information where agents update their beliefs for free or pay a financial cost to obtain information (for surveys of this literature, see Ziegler, 2003; Veldkamp, 2011, and the references therein). Instead, here the agent acquires knowledge through effort (Arrow, 1962; Grossman, Kihlstrom, and Mirman, 1977; Rob, 1991; Johnson, 2007) and the cost of learning is its observer effect: it disturbs the process of aggregate consumption by inserting uncertainty about its drift and by increasing the magnitude of its diffusion.

(as shown by Williams, 1977, p. 222-224), hence the focus of our paper on the unknown drift. It is a potentially interesting extension to think about modeling a situation in which the riskiness of the new technology is also unknown.[9]

The agent learns by observing the history of aggregate consumption (6).[10] We define the information filtration of the agent as $\{\mathcal{F}_t^\delta\}$, where $\mathcal{F}_t^\delta = \sigma(\delta_u : u \leq t)$. The agent has initial beliefs

$$\beta \sim N(\widehat{\beta}_0, \nu_0), \tag{8}$$

where $\widehat{\beta}_0 > 0$, so that the agent initially believes that investing in the experimental asset is a good idea. The parameter ν_0 captures *technological uncertainty.* The agent's learning from observing δ_t implies that both $\widehat{\beta}$ and ν evolve dynamically over time.

Proposition 1 *This partially observed economy is equivalent to a perfectly observed economy with aggregate consumption process*

$$\frac{d\delta_t}{\delta_t} = \left(\bar{f} + \widehat{\beta}_t X - \frac{c}{2}X^2\right) dt + \sigma\sqrt{1 + k^2X^2 + 2k\varphi X}\, d\widehat{W}_t, \tag{9}$$

where

$$d\widehat{\beta}_t = \frac{X\nu_t}{\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}} d\widehat{W}_t, \tag{10}$$

$$\frac{d\nu_t}{dt} = -\frac{X^2\nu_t^2}{\sigma^2(1 + k^2X^2 + 2k\varphi X)}, \tag{11}$$

and $\widehat{W}_t$ is a standard Brownian motion with respect to the agent's filtration $\{\mathcal{F}_t^\delta\}$ (the shock $d\widehat{W}_t$ can be interpreted as the "surprise" change in aggregate consumption):

$$d\widehat{W}_t \equiv \frac{d\delta_t/\delta_t - (\bar{f} + \widehat{\beta}_t X - cX^2/2)dt}{\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}} = dW_t + \frac{X(\beta - \widehat{\beta}_t)}{\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}} dt. \tag{12}$$

Proposition 1 is an application of the Kalman-Bucy filter (Theorem 12.1, p. 22, Liptser and Shiryaev, 2001).[11] According to (10), the agent revises the estimate $\widehat{\beta}$ in the direction of the consumption surprises she observes (Brennan, 1998). The expression in (11) and

[9] We thank an anonymous referee for this point.

[10] An alternative would be to allow the agent to learn by observing (3) and (4) directly. We discuss this alternative in Appendix A.11 and show that it does not qualitatively affect our results.

[11] The seminal applications of the Kalman-Bucy filter in continuous-time finance are Detemple (1986), Dothan and Feldman (1986), and Gennotte (1986).

the initial value ν_0 imply a deterministic, decreasing path for the Bayesian (technological) uncertainty ν_t:

$$\nu_t = \left(\frac{1}{\nu_0} + \frac{X^2}{\sigma^2(1+k^2X^2+2k\varphi X)} t \right)^{-1}. \tag{13}$$

Technological uncertainty starts at ν_0, but decays to zero as t goes to infinity. One benefit of experimentation is that the agent learns about the new technology and lowers Bayesian uncertainty. However, when $k > 0$, experimentation also has a negative effect on learning. Although the *speed of learning*—the term multiplying time t in (13)—increases with experimentation, it is always lower than $1/(k^2\sigma^2)$. Thus, the agent cannot perfectly learn β in any finite time.

2.1 Socially Optimal Experimentation

The agent's choice of X affects her expected lifetime utility of consumption conditional on information at time 0. We denote the agent's expected lifetime utility at time 0, as defined in (2), by $\mathcal{U}_0(\delta_0, \widehat{\beta}_0, \nu_0, X)$. For notational convenience, we will suppress the dependence of $\mathcal{U}_0$ on the variables δ_0, $\widehat{\beta}_0$, and ν_0. The agent trades a riskless asset in zero net supply and a risky asset claim to the aggregate consumption (dividend) stream δ_t. The risky asset is in positive supply of one share. Proposition 2 characterizes $\mathcal{U}_0(X)$ together with the equilibrium price-dividend ratio in this economy, $\mathcal{P}_0(X)$.

Proposition 2 *Define the function* $\kappa(X, \widehat{\beta}_0)$ *as follows:*

$$\kappa(X, \widehat{\beta}_0) \equiv (1-\gamma)\left(\bar{f} + X\widehat{\beta}_0 - \gamma \frac{\sigma^2(1+k^2X^2+2k\varphi X)}{2} - \frac{c}{2}X^2 \right) - \rho. \tag{14}$$

(a) For any $X \geq 0$*, the representative agent's lifetime expected utility is*

$$\mathcal{U}_0(X) = \mathbb{E}_0\left[\int_0^T e^{-\rho t} \frac{\delta_t^{1-\gamma}}{1-\gamma} dt \right] = \frac{\delta_0^{1-\gamma}}{1-\gamma} \int_0^T \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma-1)^2 X^2 \nu_0}{2} t^2 \right] dt. \tag{15}$$

(b) The stochastic discount factor in this economy is $\xi_t \equiv e^{-\rho t}(\delta_t/\delta_0)^{-\gamma}$*, with* $\xi_0 = 1$*. Thus, the equilibrium price-dividend ratio at time* 0 *equals*

$$\mathcal{P}_0(X) = \frac{1}{\delta_0}\mathbb{E}_0\left[\int_0^T \xi_t \delta_t dt \right] = \int_0^T \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma-1)^2}{2} X^2 \nu_0 t^2 \right] dt. \tag{16}$$

Part (a) of Proposition 2 requires computing the expectation $\mathbb{E}_0[\delta_t^{1-\gamma}]$. We compute this expectation in Appendix A.3 using the theory of affine processes (Duffie, Filipović, and Schachermayer, 2003). The stochastic discount factor in part (b) of Proposition 2 follows from standard results in asset pricing (Duffie 2010, Dumas and Luciano 2017, Ch. 12, Munk 2013, Ch. 8). Proposition 2 then shows that both $\mathcal{U}_0(X)$ and $\mathcal{P}_0(X)$ depend on the level of experimentation, and we also notice that $\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X)$. The following Corollary guarantees the existence and uniqueness of an optimal level of experimentation when $\gamma > 1$.

Corollary 2.1 *If $\gamma > 1$, the expected lifetime utility $\mathcal{U}_0(X)$ is strictly concave in X.*

Corollary 2.1 implies that the first order condition with respect to X is necessary and sufficient for maximizing lifetime utility. This is guaranteed when $\gamma > 1$, an assumption that we maintain throughout. (We discuss the case $\gamma < 1$ in Appendix A.15 and we consider an extension to recursive preferences (Epstein and Zin, 1989) in Section 4.4.)

Before solving for the optimal level of X, we define two important quantities that enter into the experimentation tradeoff, *wealth duration* and *wealth convexity.*

Definition 1 *The wealth duration and wealth convexity, $\mathcal{D}_0(X)$ and $\mathcal{C}_0(X)$, are defined respectively as the weighted average maturity and the weighted average squared maturity of the price of a claim to aggregate consumption:*

$$\mathcal{D}_0(X) = \frac{1}{\mathcal{P}_0(X)} \int_0^T t \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma-1)^2}{2} X^2 \nu_0 t^2\right] dt, \tag{17}$$

$$\mathcal{C}_0(X) = \frac{1}{\mathcal{P}_0(X)} \int_0^T t^2 \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma-1)^2}{2} X^2 \nu_0 t^2\right] dt. \tag{18}$$

The denominators in (17)-(18) dictate the weights that are used to compute the weighted averages. The *wealth duration* $\mathcal{D}_0$ measures the sensitivity of aggregate wealth to changes in expected growth. The longer the wealth duration, the greater the impact of experimentation on the agent's wealth. The *wealth convexity* $\mathcal{C}_0$ measures the sensitivity of wealth to changes in uncertainty about expected growth. The higher the wealth convexity, the greater the impact of the uncertainty generated by experimentation on wealth. Thus, $\mathcal{D}_0$ and $\mathcal{C}_0$ characterize the effect of experimentation on the term structure of wealth. Both $\mathcal{D}_0$ and $\mathcal{C}_0$ are strictly positive quantities.

The first-order condition with respect to X, $0 = \partial\mathcal{U}_0(X)/\partial X$, implies

$$0 = \int_0^T \left(\frac{\partial\kappa(X, \widehat{\beta}_0)}{\partial X} t + (\gamma-1)^2 X \nu_0 t^2\right) \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma-1)^2 X^2 \nu_0}{2} t^2\right] dt. \tag{19}$$

Dividing by $\mathcal{P}_0(X)$ allows us to interpret the first-order condition using Definition 1:

$$0 = \frac{\partial \kappa(X, \widehat{\beta}_0)}{\partial X} \mathcal{D}_0(X) + (\gamma - 1)^2 X \nu_0 \mathcal{C}_0(X). \tag{20}$$

This equation determines the socially optimal level of experimentation.

Proposition 3 *At time $t = 0$, the representative agent chooses a socially optimal level of experimentation that solves the following implicit equation*

$$X^* = \begin{cases} \frac{(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X^*)}{(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X^*) + (\gamma - 1)\nu_0 \mathcal{C}_0(X^*)}, & \text{if } \widehat{\beta}_0 - \gamma k \varphi \sigma^2 > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{21}$$

The first-best level of experimentation is strictly positive when $\widehat{\beta}_0 - \gamma k \varphi \sigma^2 > 0$, and resembles a mean-variance portfolio.[12] A more favorable belief about the promise of the new technology (higher $\widehat{\beta}_0$) increases experimentation, whereas more disruption or a higher cost of transferring resources (higher k or c) decrease experimentation. Moreover, the agent's incentive to experiment with the new technology increases with diversification (lower φ). Experimentation has a stronger impact when the wealth duration is longer. When $(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)$ is positive, the agent's incentive to experiment increases further with the wealth duration.

These tradeoffs are best understood by writing the first-order condition (20) as follows:

$$\underbrace{\widehat{\beta}_0 \mathcal{D}_0(X^*) + X^* \nu_0 \mathcal{C}_0(X^*)}_{\text{marginal benefit}} = \underbrace{c X^* \mathcal{D}_0(X^*) + \gamma k \sigma^2 (\varphi + k X^*) \mathcal{D}_0(X^*) + \gamma X^* \nu_0 \mathcal{C}_0(X^*)}_{\text{marginal cost}}, \tag{22}$$

from where we can see that c is not the only cost that the agent faces when experimenting with the new technology. Two additional terms on the right-hand side are the disruption risk and technological uncertainty imposed by the experimental asset.

The effect of technological uncertainty on the optimal level of experimentation depends on risk aversion. When $\gamma > 1$, uncertainty about the new technology dampens experimentation, to a greater extent when the wealth convexity is higher. This can be understood by observing that technological uncertainty ν_0 affects the tradeoff on both sides of (22). On the left-hand side, uncertainty increases the expected value of future consumption. This is similar to the positive effect of risk on the value of real options

[12] In Appendix A.6, we show that the agent's optimal choice is indeed consistent with a classic mean-variance tradeoff (Markowitz, 1952).

with convex payoffs. But on the right-hand side, uncertainty increases the future variance of consumption. If the agent is sufficiently risk averse, the marginal cost of higher uncertainty exceeds its benefit. With power utility this condition is met when $\gamma > 1$.[13]

One caveat is that both the wealth duration and convexity are themselves functions of X^*. In Section 3, we further characterize this dependence and show that wealth duration increases in experimentation. Nevertheless, the overall message from Proposition 3 is that a risk-averse agent prefers technologies that have more potential (higher $\widehat{\beta}_0$); prefers technological diversification (lower φ); dislikes more disruptive technologies (higher k); and, if sufficiently risk averse, dislikes the uncertainty that new technologies introduce into the economy (higher ν_0).

2.2 Experimentation Under Competition

Now that we have characterized the socially optimal level of experimentation, we proceed to describe how competition causes the level of aggregate experimentation to deviate from the social optimum. We depart from the previous setup by assuming that the economy consists of $n \geq 1$ agents. All agents derive utility from consumption as in (2), with the same coefficient of risk aversion. Together they compete to consume the aggregate consumption basket δ_t.

This section assumes that all agents can actively participate in experimenting with the new technology. In reality, some agents are passive and merely benefit from the discovery of new technologies without actively investing in them. We consider this distinction between active and passive agents in Section 4.2, where we show how our results are amplified when active agents compete for the value of the rents produced exclusively by the new technology δ_t^E.

We model each agent's share of aggregate consumption, θ_i, with a symmetric, logit Tullock contest success function (Tullock, 1980). Here, θ_i represents the proportion of the prize won based on each individual experimentation level x_i (Hirshleifer, 1989). Our modeling choice is motivated by Baye and Hoppe (2003), who show that rent-seeking competitions, patent-race games, and innovation tournaments are often strategically equivalent to a Tullock contest, and by the fact that contest success functions are commonly used to characterize R&D races (D'Aspremont and Jacquemin, 1988; Chung, 1996).

[13]More precisely, for any time-additive utility function that is increasing and concave in current consumption, uncertainty about consumption growth will decrease 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 Ziegler (2003, p. 52).

We define

$$\theta_i = \begin{cases} \frac{1}{n} & \text{if } x_j = 0 \text{ for all } j \in \{1, ..., n\}, \\ \frac{x_i}{\sum_{j=1}^n x_j} & \text{otherwise.} \end{cases} \tag{23}$$

The contest success function (23) is a member of the class of "power" success functions, for which a comprehensive axiomatic characterization in the case of n-player contests has been provided by seminal works of Skaperdas (1996) and Clark and Riis (1998). This justifies the utilization of equation (23) in our setup. In Section 2.3, we analyze extensions to this function and their effect on the equilibrium level of experimentation.

Agents' experimentation choices are perfectly observable and all agents have the same information filtration $\{\mathcal{F}_t^\delta\}$. Any agent i's lifetime expected utility, $\mathcal{U}_0^i$, can be written as

$$\mathcal{U}_0^i(x_1, ..., x_n) = \mathbb{E}_0 \left[\int_0^T e^{-\rho t} \frac{(\theta_i \delta_t)^{1-\gamma}}{1-\gamma} dt \right] = \theta_i^{1-\gamma} \mathcal{U}_0 \left(\sum_{j=1}^n x_j \right), \tag{24}$$

where $\mathcal{U}_0(\cdot)$ is the function defined and characterized in Proposition 2.

Effort by the agents in our Tullock contest affects the aggregate value of the rents available from the consumption stream. When agents experiment, they not only fight for consumption share, they affect the technological uncertainty and disruption risk that everyone faces. Prior work has considered rent-seeking with spillovers like this where the size of the pie in contests increases with effort (D'Aspremont and Jacquemin, 1988; Chung, 1996) or shrinks (Alexeev and Leitzel, 1996). Chowdhury and Sheremeta (2011) generalize this to consider linear combinations of effort complementarities in duopoly contests. Thus, our work contributes in two ways to this literature: by considering a non-linear combination and extending the analysis beyond a duopoly.

After learning (which is identical across agents and continues to hold as in Proposition 1), the aggregate consumption stream that evolves based on the total level of experimentation, $X_n \equiv \sum_{j=1}^n x_j$, is

$$\frac{d\delta_t}{\delta_t} = \left(\bar{f} + \widehat{\beta}_t X_n - \frac{c}{2} X_n^2 \right) dt + \sigma \sqrt{1 + k^2 X_n^2 + 2k\varphi X_n} d\widehat{W}_t. \tag{25}$$

This is the counterpart of Equation (9) in the representative agent version of the model, with the aggregate level of experimentation X_n replacing X.

Any Nash equilibrium involves an optimal choice of $x_i \geq 0$, taking into account the simultaneous choices of the other players. The first-order condition for agent i's maximization problem is

$$0 = \frac{\partial \mathcal{U}_0^i(x_1, ..., x_n)}{\partial x_i} = \frac{(1-\gamma)(X_n - x_i)}{x_i X_n} \theta_i^{1-\gamma} \mathcal{U}_0(X_n) + \theta_i^{1-\gamma} \frac{\partial \mathcal{U}_0(X_n)}{\partial X_n}, \tag{26}$$

which, after dividing by $\theta_i^{1-\gamma}\mathcal{U}_0(X_n)$ and replacing $\mathcal{U}_0(X_n)$ by $\frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X_n)$ in accordance with Proposition 2, yields

$$\frac{(\gamma-1)(X_n - x_i)}{x_i X_n} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{27}$$

Equation (27) can further be written as

$$\frac{X_n}{x_i} = 1 + \frac{X_n}{\gamma - 1} \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}, \quad \text{for all } i \in \{1, ..., n\}. \tag{28}$$

By inspection, it is easy to appreciate that any candidate equilibrium in this setting must be symmetric. This is because the right-hand side of (28) is the same for all agents. Any agent i that changes her experimentation has the same marginal impact on X_n and therefore on the aggregate consumption stream in (25). It follows then that X_n/x_i is identical across agents.

Denote the symmetric equilibrium by $(x^*, ..., x^*)$, where $x^* = X_n^*/n$. By substituting x^* for x_i in (27), the first-order condition can be rewritten as an equation in X_n^*:

$$\frac{(\gamma-1)(n-1)}{X_n^*} = \frac{\partial \ln \mathcal{P}_0(X_n^*)}{\partial X_n^*}. \tag{29}$$

A key result that follows from Corollary 2.1 is that the equilibrium price-dividend ratio $\mathcal{P}_0(X)$ is log-convex in the aggregate level of experimentation X (see Appendix A.4).[14] This result implies that the right-hand side of (29) strictly increases in $X_n^* \in [0, \infty)$. Appendix A.4 further shows that $\partial \ln \mathcal{P}_0(X_n^*)/\partial X_n^*$ increases from $(1-\gamma)(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(0)$, which is finite, to ∞. When $n \geq 2$, the left-hand side strictly decreases in $X_n^* \in [0, \infty)$, taking values from ∞ to 0. Thus, any equilibrium $X_n^* > 0$ that satisfies (29) is unique.

Proposition 4 *There exists a unique symmetric Nash equilibrium in which the aggregate level of experimentation under competition among n agents is strictly positive and solves*

[14] A function $f : \mathbb{R} \to \mathbb{R}$ is *logarithmically convex* or *log-convex* if $f(x) > 0$ for all $x \in \mathbb{R}$ and $\ln f$ is convex (Boyd and Vandenberghe, 2004, p. 104).

the implicit equation

$$X_n^* = \frac{(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*}}{(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma - 1)\nu_0 \mathcal{C}_0(X_n^*)}, \tag{30}$$

and $x_i^* = X_n^*/n \ \forall i$*. The quantity* X_n^* *is strictly increasing in* n *and* $X_n^* > X^*$*,* $\forall n \geq 2$*.*

To provide an intuition for the result that $X_n^* > X^*$, let $x_j = X^*/n$ for all j, where X^* is the social optimum of Proposition 3, and write each agent's utility function:

$$\mathcal{U}_0^i(x_1, ..., x_n) = \left(\frac{x_i}{\sum_{j=1}^n x_j} \right)^{1-\gamma} \mathcal{U}_0(X^*). \tag{31}$$

Starting at this social optimum in (31), a marginal change in x_i does not change $\mathcal{U}_0(X^*)$ by the envelope theorem (at the social optimum, $\partial \mathcal{U}_0(X)/\partial X = 0$). However, it does increase agent i's consumption share because $x_i / \sum_{j=1}^n x_j$ is increasing in x_i (note that $\mathcal{U}_0(X_n^*)$ takes strictly negative values when $\gamma > 1$, and thus $\mathcal{U}_0^i(x_1, ..., x_n)$ *increases* in x_i). Therefore, every agent has an incentive to increase her own experimentation from the socially optimal level, so that the equilibrium generates excessive experimentation.

A major difference between our setup and a standard Tullock contest is that the prize, $\mathcal{U}_0(X_n)$, depends on the aggregate level of experimentation. While (31) clearly shows that in a multiple-agent economy every agent has an incentive to increase experimentation from the socially optimal level, an additional result of Proposition 4 is that X_n^* *strictly increases* in n. This is an endogenous outcome of our model and can be understood as follows.

Consider an economy with $n + 1$ agents, where $n \geq 2$. In this economy, start at the aggregate level X_n^*, with $x_j = X_n^*/(n + 1)$ for all j. That is, start from the premise that the $n+1$ agents experiment in aggregate at the same level as in an economy with n agents and write each agent's utility function as

$$\mathcal{U}_0^i(x_1, ..., x_{n+1}) = \left(\frac{x_i}{\sum_{j=1}^{n+1} x_j} \right)^{1-\gamma} \mathcal{U}_0(X_n^*). \tag{32}$$

Equation (32) bears similarity with (31), with the key difference that now $\mathcal{U}_0(\cdot)$ is evaluated at X_n^* instead of X^*. Since $X_n^* > X^*$, it follows that the value X_n^* is suboptimal and $\partial \mathcal{U}_0(X_n^*)/\partial x_i < 0$. In words, further increasing experimentation when the aggregate level is X_n^* leads to a smaller aggregate "pie." Considering this, every agent has a clear incentive to *decrease* experimentation. But although a marginal change in x_i shrinks the pie, it also increases the consumption share $x_i / \sum_{j=1}^{n+1} x_j$, or agent i's share of the pie.

Appendix A.7 shows that this second force always dominates in our model and therefore every agent has an incentive to increase her own experimentation from the level $X_n^*/(n+1)$, so that in equilibrium the aggregate level of experimentation strictly increases in n.

The principal outcome is that creative destruction and Schumpeterian competition ultimately result in over-experimentation and a reduction in aggregate welfare from the socially optimal level. As competition intensifies, the aggregate level of experimentation deviates increasingly from the socially optimal level. This finding aligns with existing models of technological change (e.g., Acemoglu, 2009, Ch. 12.3), where *business stealing* generates a divergence between the private and social value of innovation, thereby enabling the possibility of excessive innovation.

Moreover, it can be shown that competition can promote inefficient technologies. Consider the case $\widehat{\beta}_0 - \gamma k \varphi \sigma^2 < 0$ from Proposition 3. From the perspective of a welfare-maximizing agent, such a technology is inefficient and $X^* = 0$. But when $n \geq 2$ agents compete for consumption share, Proposition 4 implies $X_n^* > 0$. Thus, innovation accompanied by Schumpeterian competition can lead to over-investment in inefficient technologies.

Illustration We illustrate the effect of Schumpeterian competition on experimentation with a numerical example. The parameters that we choose are economically plausible and will be the same for the rest of the paper. We fix the risk aversion to $\gamma = 2$, below the level of ten deemed reasonable by Mehra and Prescott (1985, p. 154).[15] We calibrate the long-term growth and the volatility of consumption in the status-quo economy to $\bar{f} = 0.03$ (Andrei and Hasler, 2015, Table 1) and $\sigma = 0.05$ (this parameter allows us to obtain reasonable levels for the risk premium in Section 3). The agents' initial beliefs about the new technology are set to $\widehat{\beta}_0 = 0.03$ (the same value as $\bar{f}$) and $\nu_0 = 0.03^2$ (in line with the variance of expected growth of consumption, estimated at 0.029^2 in Andrei and Hasler, 2015). We fix the disruption parameter to $k = 3$ and the cost of transferring resources to $c = 0.02$. These values imply that at the social optimum ($X^* = 0.16$), the price-dividend ratio is $\mathcal{P}_0(X^*) = 17$, which is consistent with Robert Shiller's Cyclically Adjusted PE Ratio that averaged 16.91 from 1870 to 2021. Given this, the representative agent is willing to forgo 0.03% expected growth in the status-quo to experiment with the new technology, as shown in Equation (4). For the rest of our illustrations, we fix $\varphi = 1$ (i.e., no technological diversification), unless indicated otherwise. Finally, we fix $T = 100$ and the subjective discount rate to $\rho = 0.03$. These values yield a duration

[15] Friend and Blume (1975) estimate an average coefficient well in excess of one and perhaps in excess of two. Dreze (1981) finds values between 0.6 and 10 using an analysis of deductibles in insurance contracts.

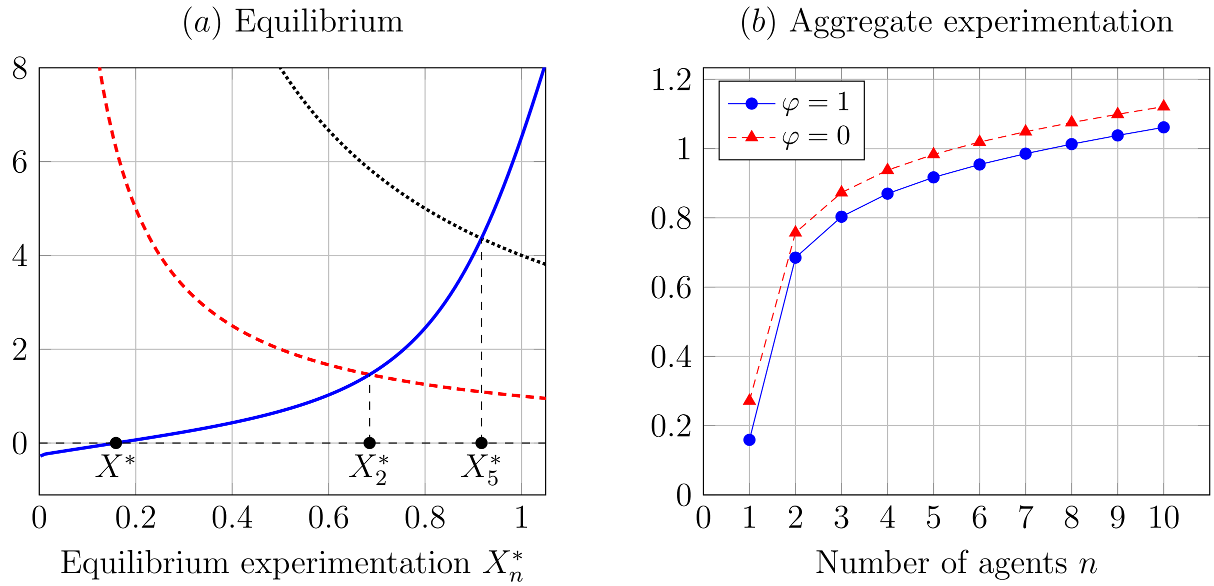

Figure 1: **Equilibrium Experimentation with Competition.** Panel (a) depicts the solution of Equation (29), where the solid line represents the right-hand side and the dashed (dotted) line represent the left-hand side for $n = 2$ ($n = 5$). The social optimum, X^*, is obtained when the solid line crosses the zero axis. Panel (b) depicts the aggregate level of experimentation as a function of the number of agents in the economy, for two levels of technological diversification, $\varphi \in \{0, 1\}$. The calibration used is: $\gamma = 2$, $\bar{f} = 0.03$, $\widehat{\beta}_0 = 0.03$, $\nu_0 = 0.03^2$, $\sigma = 0.05$, $k = 3$, $\rho = 0.03$, $T = 100$, $c = 0.02$, and $\delta_0 = 1$.

$\mathcal{D}_0(X_n^*)$ between 17 years (for $n = 1$) and 65 years (for $n = 10$), consistent with numbers estimated by Van Binsbergen (2020).

Panel (a) of Figure 1 depicts the solution of Equation (29), where the solid line represents the right-hand side and the dashed (dotted) line represents the left-hand side for $n = 2$ ($n = 5$). The social optimum X^* is obtained when the solid line crosses the zero axis.[16] The plot shows that when $n \geq 2$ the equilibrium generates excessive levels of experimentation. Panel (b) plots the aggregate level of experimentation X_n^* when the number of agents varies from 1 to 10, for two levels of technological diversification. The lines start at the social optimum X^* and experimentation increases as competition rises in the market, further confirming the results of Proposition 4. Diversification ($\varphi < 1$) increases the aggregate level of experimentation, as shown with the dashed line. The amplification arises because the term $\widehat{\beta}_0 - \gamma k \varphi \sigma^2$ now takes into account the diversification

[16] This example is only intended to be illustrative, and in reality, X^* should increase gradually. However, in the context of our theoretical model, the decision to experiment is made only once. A more realistic model with dynamic experimentation—as the one we develop in Section 4.3—would take into account the gradual adoption of the new technology over time.

benefit of the new technology.

2.3 Economic forces that may dampen experimentation

Based on the results so far, it is natural to consider what might dampen the over-experimentation result obtained in Proposition 4 and illustrated in Figure 1. One possibility is that there are decreasing returns to experimentation. This can be analyzed by an extension of the contest success function (23) that preserves the axiomatic characterization provided by Skaperdas (1996) and Clark and Riis (1998).

Consider the following extension of the Tullock contest:

$$\theta_i = \begin{cases} \frac{1}{n} & \text{if } x_j = 0 \text{ for all } j \in \{1, ..., n\}, \\ \frac{x_i^r}{\sum_{j=1}^n x_j^r} & \text{otherwise.} \end{cases} \tag{33}$$

The parameter $r > 0$ measures returns to scale: the case $0 < r < 1$ represents decreasing returns, while $r > 1$ represents increasing returns to experimentation (Baye, Kovenock, and De Vries, 1994; Chung, 1996). We will focus here on the case of decreasing returns.

As in Section 2.2, any Nash equilibrium involves an optimal choice of $x_i \geq 0$, taking into account the simultaneous choices of the other players. The first-order condition for agent i's maximization problem is

$$0 = \frac{\partial \mathcal{U}_0^i(x_1, ..., x_n)}{\partial x_i} = \frac{(1-\gamma) r (x_i^{-1} \sum_{j=1}^n x_j^r - x_i^{r-1})}{\sum_{j=1}^n x_j^r} \theta_i^{1-\gamma} \mathcal{U}_0(X_n) + \theta_i^{1-\gamma} \frac{\partial \mathcal{U}_0(X_n)}{\partial X_n}, \tag{34}$$

which, after dividing by $\theta_i^{1-\gamma}\mathcal{U}_0(X_n)$ and replacing $\mathcal{U}_0(X_n)$ by $\frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X_n)$ in accordance with Proposition 2, yields

$$\frac{(\gamma - 1) r (\sum_{j=1}^n x_j^r - x_i^r)}{x_i \sum_{j=1}^n x_j^r} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{35}$$

In a symmetric equilibrium, this leads to

$$\frac{(\gamma - 1)(n-1) r}{X_n^*} = \frac{\partial \ln \mathcal{P}_0(X_n^*)}{\partial X_n^*}, \tag{36}$$

which yields an equilibrium aggregate level of experimentation that solves

$$X_n^* = \frac{(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X_n^*) + \frac{(n-1)r}{X_n^*}}{(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma - 1)\nu_0 \mathcal{C}_0(X_n^*)}. \tag{37}$$

Comparing this with the result of Proposition 4, the aggregate level of experimentation is lower when there are decreasing returns ($0 < r < 1$). The dampening effect is stronger for smaller values of r. The only situation in which the social optimum of Proposition 3 is reached is in the limit $r \to 0$. Otherwise, agents will over-experiment as in Section 2.2.

A second factor that can dampen aggregate experimentation is the presence of negative externalities imposed by competition on the utility gained through experimentation (Chowdhury and Sheremeta, 2011). Our framework already accounts for a negative externality in the form of deviation from the social optimum, which decreases aggregate welfare and shrinks the resources available to all agents. However, it is not uncommon for competition for new technologies to lead to patent warfare, which can harm innovation by diverting resources towards litigation or discouraging experimentation, as highlighted by the works of Shaver (2012), Karakashian (2015), and Trappey, Trappey, and Wang (2016).

To examine this, assume that instead of (24) each agent gets

$$\mathcal{U}_0^i(x_1, ..., x_n) = f(n, x_i)\theta_i^{1-\gamma}\mathcal{U}_0\left(\sum_{j=1}^n x_j\right), \tag{38}$$

where $f(n, x_i)$ is a function that satisfies three assumptions: (i) $f(1, x_i) = 1$, (ii) $f(n, x_i) \geq 1$, and (iii) $\partial f(n, x_i)/\partial x_i > 0$ for $n > 1$.

Assumption (i) implies that if $n = 1$ we obtain the social optimum of Proposition 3. Assumptions (ii) and (iii) imply $\partial \ln f(n, x_i)/\partial x_i > 0$, a property that will be useful below. Moreover, since $f(n, x_i) \geq 1$, the multiplication with $f(n, x_i)$ in (38) magnifies $\mathcal{U}_0(\cdot)$. Keeping in mind that $\mathcal{U}_0(\cdot)$ is negative when $\gamma > 1$, this magnification generates an utility loss for agent i. This utility loss occurs only when competition is present, i.e., when $n > 1$. One can think of this effect as a negative externality that arises from competition (e.g., patent warfare).

The first-order condition for agent i's maximization problem is

$$0 = \frac{(1-\gamma)(X_n - x_i)}{x_i X_n}\theta_i^{1-\gamma}\mathcal{U}_0(X_n)f(n, x_i) + \theta_i^{1-\gamma}\frac{\partial \mathcal{U}_0(X_n)}{\partial X_n}f(n, x_i) + \theta_i^{1-\gamma}\mathcal{U}_0(X_n)\frac{\partial f(n, x_i)}{\partial x_i}, \tag{39}$$

which, after dividing by $\theta_i^{1-\gamma}\mathcal{U}_0(X_n)f(n, x_i)$ and replacing $\mathcal{U}_0(X_n)$ by $\frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X_n)$ in accordance with Proposition 2, yields

$$\frac{(\gamma-1)(X_n - x_i)}{x_i X_n} - \frac{\partial \ln f(n, x_i)}{\partial x_i} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{40}$$

Comparing this condition with (29), we notice the term $\partial \ln f(n, x_i)/\partial x_i$, which is positive according to assumptions (ii) and (iii) above. Thus, this term decreases the left-hand side and according to panel (a) of Figure 1 dampens the equilibrium level of experimentation. The aggregate level of experimentation now solves

$$X_n^* = \frac{(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*} - \frac{1}{\gamma-1}\frac{\partial \ln f(n,x_i)}{\partial x_i}}{(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma - 1)\nu_0 \mathcal{C}_0(X_n^*)}. \tag{41}$$

A sufficiently strong negative externality imposed on agent i by competition can generate under-investment in the new technology. Consider the externality function $f(n, x_i) = e^{\alpha(n-1)x_i}$, with $\alpha > 0$. This function satisfies assumptions (i)-(iii) above. It can be shown that $\alpha > \gamma - 1$ leads to an aggregate level of experimentation that is below the social optimum.

3 The Implications of Schumpeterian Competition

In this section, we evaluate the effects of Schumpeterian competition and excessive levels of experimentation on technological uncertainty, systematic risk, and asset prices.

3.1 Risk and Asset Prices

Both disruption risk ($k > 0$) and technological uncertainty ($\nu_0 > 0$) affect the path of the future volatility of consumption, $\{\mathrm{Vol}_0[\delta_t]\}_{t=0}^T$. Figure 2 illustrates this for one maturity ($t = 30$). When competition increases experimentation in the economy, it magnifies the disturbance to the status quo, which is the hashed area in the plot. Experimentation also initially introduces technological uncertainty, which further magnifies risk and is captured by the gray area in the plot. So, experimentation creates macroeconomic risk through two channels and competition makes this worse.

To explore how these two sources of risk affect financial markets, we solve for the equilibrium asset prices when the n agents from the previous section share the aggregate output in proportions given by (23) and determined in Proposition 4. Agents trade a risky asset, which is a claim to the aggregate output stream δ_t, and a risk-free asset that is in zero net supply. The equilibrium price-dividend ratio at any time $t \in [0, T]$ follows from Proposition 2:

$$\mathcal{P}_t(X_n^*) = \int_t^T \exp\left[\kappa(X_n^*, \widehat{\beta}_t)(s - t) + \frac{(1-\gamma)^2}{2}(X_n^*)^2 \nu_t (s-t)^2\right] ds. \tag{42}$$

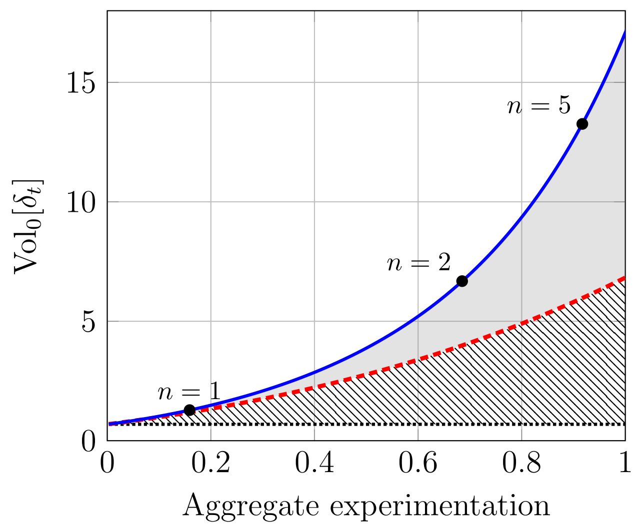

Figure 2: **Equilibrium Experimentation with Competition and its Effect on Macroeconomic Risk.** The figure depicts the increase in the volatility of future consumption, $\text{Vol}_0[\delta_t]$, for $t = 30$ years. The increase is due to disruption risk ($k > 0$, hashed area) and technological uncertainty ($\nu_0 > 0$, gray area).

The price-dividend ratio increases with the uncertainty about β. To see this, suppose that β were known (i.e., $\widehat{\beta}_t = \beta$ and $\nu_t = 0$ for all t). Then, an application of the Leibniz integral rule would confirm that $\partial^2 \mathcal{P}_t(X_n^*)/\partial\beta^2 > 0$ and thus the price-dividend ratio is *convex* in β. Since β is in fact a random variable, Jensen's inequality implies that the price-dividend ratio under incomplete information ($\nu_t > 0$) must be greater than the one under complete information ($\nu_t = 0$). This is a general result and does not depend on the value of γ (e.g., Pástor and Veronesi, 2003, 2006).

Our analysis implies, then, that over-valuation is magnified with Schumpeterian competition because rent seeking amplifies technological uncertainty. Panel (a) in Figure 3 depicts the price-dividend ratio as a function of aggregate experimentation (solid line). The dashed line depicts the price-dividend ratio while setting technological uncertainty to zero. The gap between the two lines, illustrated with the gray area, provides a measure of over-valuation due to technological uncertainty. We also plot the point that corresponds to the socially optimal level of experimentation. Schumpeterian competition leads to over-experimentation, which increases the wealth duration and the gap between the two lines. The gap grows from being almost negligible at the social optimum to roughly 100% when five agents compete for consumption share.

The over-valuation of the asset caused by competition and technological uncertainty has consequences for the future path of the price-dividend ratio. In panel (b), we compare

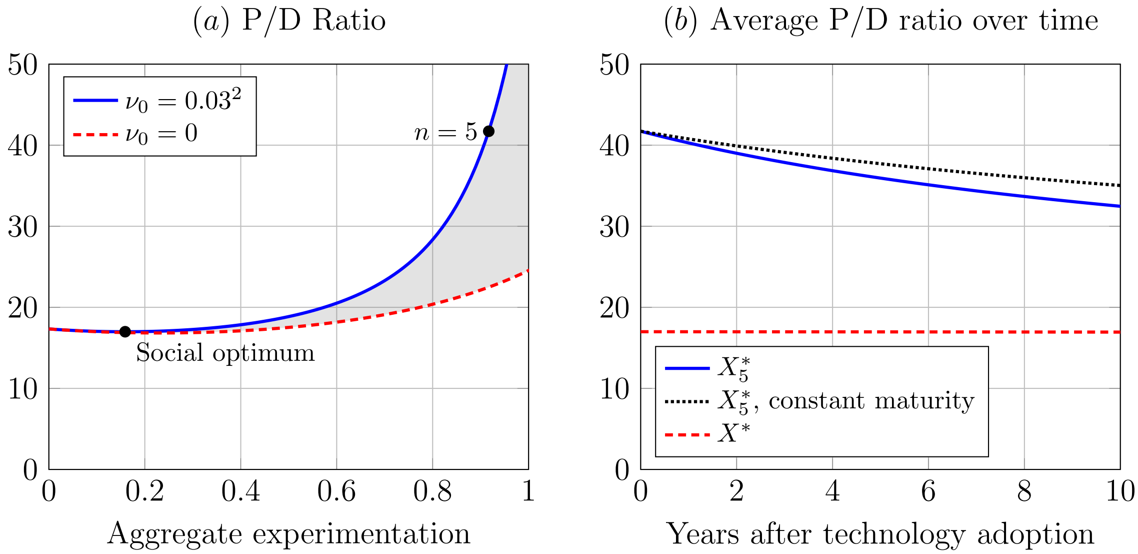

Figure 3: **Experimentation and Asset Prices.** The solid line in panel (a) depicts the price-dividend ratio as a function of the level of aggregate experimentation in the economy. The dashed line depicts the price-dividend ratio with uncertainty set to $\nu_0 = 0$. The dot labeled "$n = 5$" corresponds to the equilibrium aggregate level of experimentation in an economy with five agents. Panel (b) depicts average price paths over time starting from "$n = 5$" and $\widehat{\beta}_0 = \beta$ (solid and dotted lines) and from the social optimum (dashed line). The dotted line maintains the maturity of the asset constant at $T = 100$ years.

the average paths of the price-dividend ratio once the points "$n = 5$" and the social optimum have been reached. We start from a prior $\widehat{\beta}_0$ that is exactly equal to the true value of β, so that agents are initially right, but are still uncertain about the expected growth of the new technology. As such, we are not imposing any ex-ante bias on future price paths. This assumption ensures that simulations of the economy over time will result in an average $\widehat{\beta}_t$ that is equal to the true value of β.[17] Because the only remaining state variable that enters into the price-dividend ratio, ν_t, decreases deterministically over time, we can plot the average price-dividend ratio over time by simply using (13) for the value of ν_t, without resorting to simulations.

The solid line in panel (b) depicts the average path of the price-dividend ratio starting from a competitive equilibrium with five agents. The dashed line shows the average path of the price-dividend ratio starting from the social optimum. Because the remaining life of the asset diminishes, we also plot the dotted line, which isolates the effect of

[17] 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 $t = 0 + dt$.

the decrease in uncertainty by holding the remaining asset life constant at $T = 100$ years.[18] With competition ($n = 5$), the average asset price has the tendency to decrease as uncertainty about the new technology resolves. The pace of the decline depends on the rate at which uncertainty about the new technology is resolved. Equation (13) shows that uncertainty decays faster with over-experimentation, which accelerates the decline. In contrast, after experimentation at the social optimum, the average asset price decrease due to resolution of uncertainty is negligible. Thus, markets characterized by intense Schumpeterian competition are prone to asset price over-valuation and to subsequent declines through learning and resolution of uncertainty.

Now we explore the effects of experimentation and rent-seeking on how the agents trade the risky asset. We rely on results from the theory of dynamic portfolio choice (Merton, 1973) under incomplete information (Brennan, 1998). Once experimentation begins, all agents face the same time-varying investment opportunity set and the expected growth rate of the economy is driven by $\widehat{\beta}_t$. Thus, the agents need to hedge and their willingness to hold the risky asset evolves over time. Denote the risk premium by $\pi_t(X_n^*)$ and the instantaneous diffusion of the risky asset by $\sigma_{P,t}(X_n^*)$. Both quantities will be determined in equilibrium once we impose market clearing.

Proposition 5 *At $t > 0$, agents invest the same proportion of their wealth in the risky asset,*

$$\phi_t = \frac{1}{\gamma}\frac{\pi_t(X_n^*)}{\sigma_{P,t}(X_n^*)^2} + (1-\gamma)\frac{(X_n^*)^2\nu_t}{\sigma\sqrt{1+k^2(X_n^*)^2+2k\varphi X_n^*}\sigma_{P,t}(X_n^*)}\mathcal{D}_t(X_n^*), \tag{43}$$

where $\mathcal{D}_t(X_n^*)$ is the wealth duration at time $t \in [0,T]$, which follows from Definition 1.

The first term in (43) is the myopic demand that results from a classic risk-return tradeoff (Markowitz, 1952). The second is a hedging demand, $\mathcal{H}_t(X_n^*, \nu_t)$, which results from a desire to manage the uncertainty of future changes in expected growth (Merton, 1973). As time evolves, agents trade in order to hedge the technological uncertainty generated by experimentation. For example, if $\gamma > 1$ and $\sigma_{P,t}(X_n^*) > 0$, the hedging demand $\mathcal{H}_t(X_n^*, \nu_t)$ is negative: agents dislike uncertainty and therefore decrease their demand for the risky asset, consistent with the analysis in Brennan (1998). In this case, an increase in $\widehat{\beta}_t$ is "good news" for consumption, but because the return of the risky asset is positively correlated with consumption growth ($\sigma_{P,t} > 0$), the hedging weight is negative in order to stabilize agents' utility.

[18] This adjustment is unnecessary in the social optimum, where the decrease in the price-dividend ratio is almost imperceptible with our calibration.

Before deriving explicit expressions for $\pi_t(X_n^*)$ and $\sigma_{P,t}(X_n^*)$, we clarify the impact that technological uncertainty has on both of these quantities. As in Breeden (1979), the market price of risk in this economy is equal to the diffusion of aggregate consumption times aggregate risk aversion, $\zeta(X_n^*) = \gamma\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}$. Proposition 5 and market clearing together imply:[19]

$$\pi_t(X_n^*) = \frac{\zeta^2(X_n^*)/\gamma}{1 - \mathcal{H}_t(X_n^*, \nu_t)} \quad \text{and} \quad \sigma_{P,t}(X_n^*) = \frac{\zeta(X_n^*)/\gamma}{1 - \mathcal{H}_t(X_n^*, \nu_t)}. \tag{44}$$

Since the market price of risk does not depend on ν_t,[20] technological uncertainty affects $\pi_t(X_n^*)$ and $\sigma_{P,t}(X_n^*)$ only through the agents' hedging motives: a negative hedging demand lowers both $\pi_t(X_n^*)$ and $\sigma_{P,t}(X_n^*)$. *Ceteris paribus*, the risk premium and stock return diffusion are lower in an economy with technological uncertainty ($\nu_t > 0$) than in an economy without technological uncertainty ($\nu_t = 0$ and $\mathcal{H}_t(X_n^*, 0) = 0$).

Proposition 6 fully characterizes the equilibrium stock diffusion and the risk premium in this economy, based on primitives in the model and on the wealth duration.[21]

Proposition 6 *The equilibrium stock diffusion with experimentation is given by*

$$\sigma_{P,t}(X_n^*) = \sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*} - (\gamma - 1)\frac{(X_n^*)^2\nu_t}{\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}}\mathcal{D}_t(X_n^*), \tag{45}$$

and the equilibrium risk premium is

$$\pi_t(X_n^*) = \zeta(X_n^*)\sigma_{P,t}(X_n^*) = \gamma\sigma^2(1 + k^2(X_n^*)^2 + 2k\varphi X_n^*) - \gamma(\gamma - 1)(X_n^*)^2\nu_t\mathcal{D}_t(X_n^*). \tag{46}$$

Because $k > 0$, the first term in (45) shows that experimentation increases $\sigma_{P,t}(X_n^*)$ by disturbing the economy and amplifying macroeconomic fluctuations. However, the second term implies that when $\gamma > 1$, technological uncertainty lowers $\sigma_{P,t}(X_n^*)$, which is due to a hedging motive and is consistent with the previous discussion of (44). Thus, the two types of risk imposed by creative destruction—disruption risk and technological uncertainty—may have opposite effects on systematic risk. The balance between these

[19] To obtain (44), impose market clearing ($\phi_t = 1$) in (43) and use $\zeta(X_n^*) = \pi_t(X_n^*)/\sigma_{P,t}(X_n^*)$ to solve for $\pi_t(X_n^*)$ and $\sigma_{P,t}(X_n^*)$.

[20] This statement is true in any setting with time-additive utility. It is only in settings with recursive utility that the market price of risk depends on the level of uncertainty about the growth rate. See Section 4.4 for an extension of our setup to recursive utility.

[21] An analogous of this proposition in a standard general equilibrium model with learning can be found in Brennan and Xia (2001), Theorem 5.

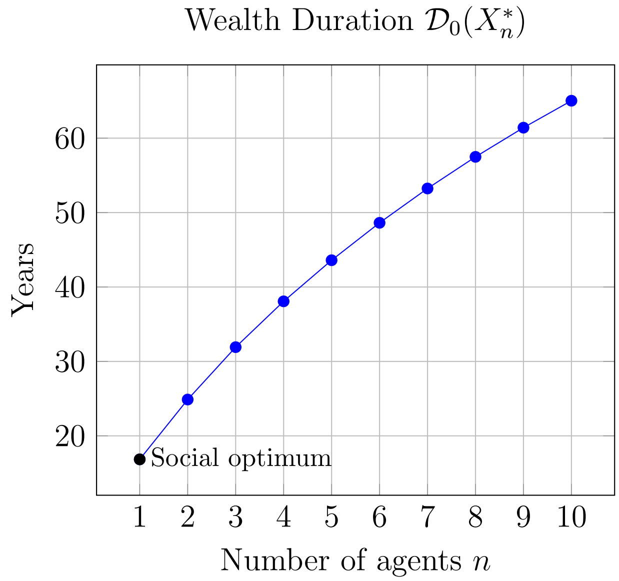

Figure 4: **Experimentation with Competition and its Effect on the Wealth Duration.** The line depicts the wealth duration as a function of the number of agents in the economy.

two forces depends on the level of experimentation.[22]

The expression for the equity risk premium in (46) is also made up of two terms. Agents require a premium for holding the risky asset, but at the same time their hedging motives yield a negative (insurance) premium. The balance between these two forces is affected by the level of experimentation for two reasons. First, higher experimentation increases the uncertainty in the economy and the impact of hedging. Second, higher experimentation tends to increase the wealth duration $\mathcal{D}_t(X_n^*)$. This can be seen from (17), where the denominator dictates the weights that are used in the numerator to compute the weighted-average maturity. Technological uncertainty increases the long-maturity weights, amplifying the long-term impact of experimentation. More competition increases experimentation, which puts more weight on long-term payoffs and increases $\mathcal{D}_t(X_n^*)$. Figure 4 plots the wealth duration for different values of n. It confirms that the wealth duration increases with the degree of competition.

Figure 5 illustrates the effect of experimentation and competition on the risk premium. The risk premium is hump-shaped in experimentation. The initial increase is driven

[22]Another way to understand the opposite impact of the two terms in (45) is to start from the asset price, $P_t = \delta_t \mathcal{P}_t(X_n^*)$, where $\mathcal{P}_t(X_n^*)$ is the equilibrium price-dividend ratio defined in (42). A positive shock to δ_t increases P_t, hence the first term in (45). But the same positive shock increases $\widehat{\beta}_t$ through agents' learning (Proposition 1). A higher $\widehat{\beta}_t$ implies high future consumption and low future marginal utility. When $\gamma > 1$ the latter effect dominates, hence the second term in (45).

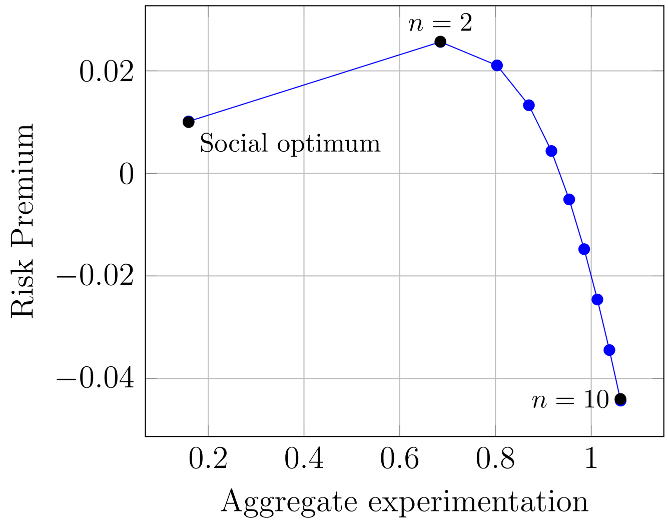

Figure 5: **Experimentation and the Risk Premium.** This figure plots the effect of experimentation on the risk premium in the economy.

by disruption risk, and the subsequent decrease is driven by technological uncertainty. The effect of competition on the risk premium depends on the number of agents in the market. When n is low, the impact of disruption risk dominates and competition increases the risk premium. However, when n is large, the complementarity between duration and technological uncertainty creates a significant hedging demand and lowers the risk premium. Intense competition may lead to a negative risk premium, when $\sigma_{P,t}(X_n^*) < 0$. In this case, the risky asset becomes a good hedge against fluctuations in expected growth, $\mathcal{H}_t(X_n^*, \nu_t) > 0$, and agents increase their overall demand of the risky asset beyond the myopic position dictated by the standard tradeoff between risk and return.

Proposition 6 has additional implications for the term structure of risk. A recent literature studies the risk premia of equity claims with different maturities and documents that long-duration assets earn lower returns than short-duration assets (Lettau and Wachter, 2007, 2011; Weber, 2018; Van Binsbergen, 2020; Gonçalves, 2021). Over the past century, long-duration dividend risk has received little to no compensation (Van Binsbergen, 2020). The negative impact of duration on the risk premium in our model (Equation 46) is consistent with these findings and proposes that long-duration assets may provide hedging against technological uncertainty. The fact that in the model assets that are more exposed to technological uncertainty have longer duration is consistent with the idea that growth firms are long-duration assets (Lettau and Wachter, 2007). Understanding the impact of Schumpeterian competition on the equity duration and the cross-section of returns may prove to be a fruitful avenue for future research.

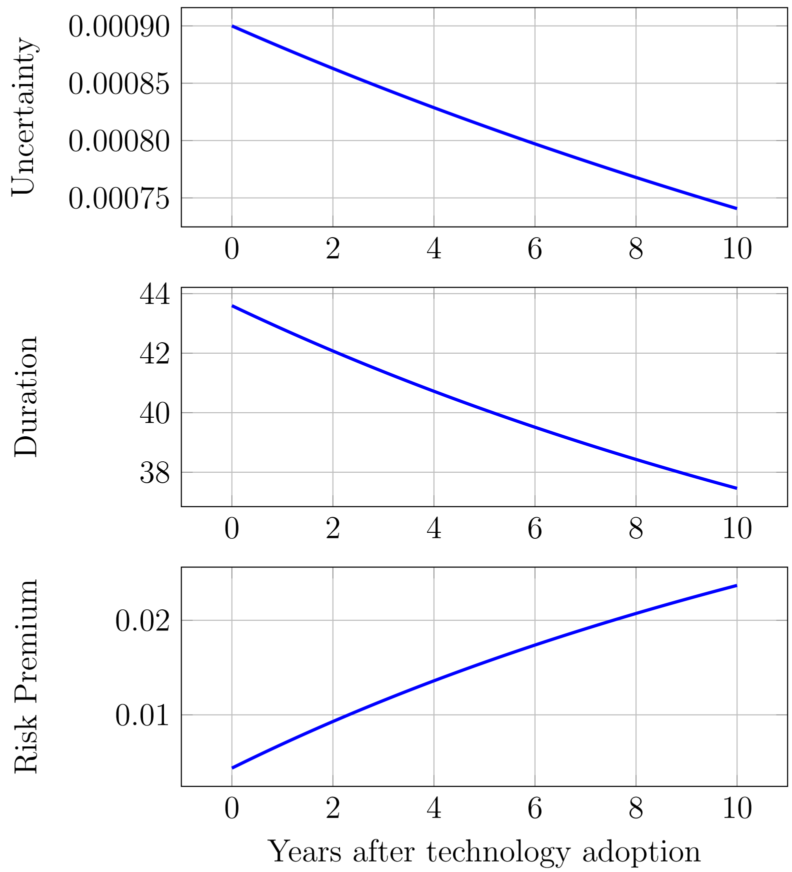

Figure 6: **Dynamic Patterns of Risk.** The top panel shows the evolution of ν_t after the time of technology adoption ($t = 0$), in an economy with $n = 5$ agents. The middle panel shows the evolution of wealth duration (while keeping the maturity T constant, to avoid a mechanical effect arising from the reduction of maturity). The bottom panel depicts the dynamics of the risk premium over the same period of time.

3.2 Return Predictability and Secular Trends

While the impact of disruption risk on the risk premium is permanent, the effect of technological uncertainty is time-varying, depending on the speed of learning and the amount of wealth duration. Proposition 6 shows that the risk premium is time-varying and depends on both ν_t and $\mathcal{D}_t(X_n^*)$. When a new technology is initially adopted, technological uncertainty and wealth duration are both high. Agents build hedging positions to manage risk, which decreases the risk premium. However, as learning takes place, both technological uncertainty and wealth duration decline, so that the risk premium rises over time. Figure 6 illustrates these patterns in an economy with $n = 5$ competitors.

Since the price-dividend ratio in (42) is partially driven by movements in technological uncertainty, when agents over-experiment, uncertainty increases the price-dividend ratio,

Panel I: $n = 5$, $X_5^* = 0.917$ (medians of 10,000 simulations)

	4Q	12Q	20Q	40Q
Coefficient of Log (P/D)	0.043	0.128	0.207	0.411
t-stat	1.263	1.257	1.302	1.555
R-squared	0.010	0.030	0.051	0.109
Expected Excess Returns (annualized)	1.1%	1.3%	1.5%	1.9%

Panel II: $n = 20$, $X_{20}^* = 1.271$ (medians of 10,000 simulations)

	4Q	12Q	20Q	40Q
Coefficient of Log (P/D)	-0.033	-0.093	-0.144	-0.243
t-stat	-2.087	-2.315	-2.469	-2.614
R-squared	0.054	0.137	0.190	0.274
Expected Excess Returns (annualized)	-7.5%	-6.8%	-6.1%	-4.5%

Table 1: **Return Predictability with the Price-Dividend Ratio (Simulations).** This table reports the predictability of excess stock returns with the log price-dividend ratio, estimated from (47). There are two panels, each one corresponding to a different value of n. Each panel reports medians from 10,000 simulations for the regression coefficients, t-stats computed with Newey and West (1987) standard errors with $2(K-1)$ lags, and the R^2 coefficients. The panels show results for different horizons: 4, 12, 20, and 40 quarters. The last row of each panel computes the median annualized excess return between the initial date and the horizon indicated in each column.

but after $t = 0$ the downward drift of uncertainty caused by learning (Figure 6, top panel) deflates the price-dividend ratio (Figure 3, panel (b)). These dynamics of risk and pricing suggest that competition impacts the time-series predictability of returns. Using a discretization of our continuous-time setup (Appendix A.9), we perform simulations of the model at a quarterly frequency for 100 years and run the following predictive regressions (Beeler and Campbell, 2012):

$$\sum_{k=1}^{K}(r_{t+k} - r_{f,t+k}) = a_K + b_K pd_t + \epsilon_{t+K}, \tag{47}$$

where the dependent variable is the K-period time-aggregated excess return, the independent variable is the log price-dividend ratio, and $K = \{4, 12, 20, 40\}$ quarters. We keep the remaining life of the asset constant at $T = 100$ years in order to eliminate effects arising from the diminishing time to maturity.

The results in Table 1 show that return predictability is positive for $n = 5$, but negative for $n = 20$. When $n = 20$, the last line of the table shows that agents expect negative excess returns: they are willing to hold the risky asset at a high price, for which they demand a negative risk premium. Thus, high prices may negatively predict future excess returns, particularly when Schumpeterian competition is intense.

Finally, our framework seems to provide an economic interpretation of how changes in competition may drive secular trends relating to innovation and risk premia. In recent decades, there has been a secular decline in competition that accelerated in the early 2000s (e.g., Grullon, Larkin, and Michaely, 2019; De Loecker, Eeckhout, and Unger, 2020). Recent studies have attributed slumping investment and innovation (Gutiérrez and Philippon, 2018) and rising risk premia in equity markets (Corhay et al., 2020) to declining competition. Our model's predictions appear to be consistent with these stylized facts: a decline in the number of competitors would indeed result in depressed investment in new technologies (Figure 1) and higher risk premia (Figure 5). The model can therefore provide an economic explanation for why declining competition leads to slumping innovation and a rising equity premium since the early 2000s.

4 Extensions and Discussion of Assumptions

Our baseline model relied on several simplifying assumptions, which served our purpose of isolating the main results. In this section, we relax some of these assumptions and explore various extensions of the model. For the rest of the paper, we will assume $\varphi = 1$ to keep the analysis of these extensions as simple as possible.

4.1 Infinite Horizon and Obsolescence

An important feature of our model is that time is a state variable: as can be seen from (13), when agents experiment with the new technology, the technological uncertainty ν_t decays deterministically with time. This gradual resolution of uncertainty, arguably a common feature of many new technologies, generates our dynamic results in Section 3.2. But it comes with a caveat: it is a well-known observation that learning about a constant, as in Proposition 1, implies that the posterior belief $\widehat{\beta}_t$ is a martingale and thus has infinite persistence (Collin-Dufresne, Johannes, and Lochstoer, 2016). The immediate implication of this observation is that all of the integrals in Proposition 2 would not converge when $T \to \infty$ (the terms that multiply t^2 in (15)-(16) are strictly positive). Thus, the main purpose of our finite horizon assumption is to keep the model stationary.

The model can be extended to an infinite horizon case without losing the realistic feature of a gradual resolution of uncertainty, by supposing that the new technology will eventually become obsolete (Aghion and Howitt, 1992). Assume that the unknown parameter β decays deterministically to zero

$$d\beta_t = -\lambda\beta_t dt, \tag{48}$$

and that the *rate of obsolescence* ($\lambda \geq 0$) for the new technology is commonly known. (In the baseline model, $\lambda = 0$.) Continue to assume that agents have a common initial prior,

$$\beta_0 \sim N(\widehat{\beta}_0, \nu_0), \tag{49}$$

with $\widehat{\beta}_0 > 0$. They learn about β as follows.

Proposition 7 *This partially observed economy is equivalent to a perfectly observed economy with aggregate consumption process*

$$\frac{d\delta_t}{\delta_t} = \left(\bar{f} + \widehat{\beta}_t X - \frac{c}{2}X^2\right)dt + \sigma(1+kX)d\widehat{W}_t, \tag{50}$$

where

$$d\widehat{\beta}_t = -\lambda\widehat{\beta}dt + \frac{X\nu_t}{\sigma(1+kX)}d\widehat{W}_t, \tag{51}$$

$$\frac{d\nu_t}{dt} = -2\lambda\nu_t - \frac{X^2\nu_t^2}{\sigma^2(1+kX)^2}, \tag{52}$$

and $\widehat{W}_t$ *is a standard Brownian motion with respect to the agents' filtration* $\{\mathcal{F}_t^\delta\}$.

Equation (52) implies

$$\nu_t = \left(\frac{e^{2\lambda t}}{\nu_0} + \frac{X^2}{\sigma^2(1+kX)^2}\frac{e^{2\lambda t}-1}{2\lambda}\right)^{-1}, \tag{53}$$

which has the same interpretation as (13).

The following is the counterpart of Proposition 2 in the infinite horizon case.

Proposition 8 *Define the function* $\kappa(X, \widehat{\beta}_0, t)$ *as follows:*

$$\kappa(X, \widehat{\beta}_0, t) \equiv \left[(1-\gamma)\left(\bar{f} - \frac{c}{2}X^2 - \gamma\frac{\sigma^2(1+kX)^2}{2}\right) - \rho\right]t + (1-\gamma)\widehat{\beta}_0 X\frac{1-e^{-\lambda t}}{\lambda}. \tag{54}$$

(a) For any $X \geq 0$, the representative agent's lifetime expected utility is

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma} \int_0^{\infty} \exp\left[\kappa(X, \widehat{\beta}_0, t) + \frac{(\gamma-1)^2}{2} X^2 \nu_0 \left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2\right] dt. \tag{55}$$

(b) The equilibrium price-dividend ratio equals

$$\mathcal{P}_0(X) \equiv \int_0^{\infty} \exp\left[\kappa(X, \widehat{\beta}_0, t) + \frac{(\gamma-1)^2}{2} X^2 \nu_0 \left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2\right] dt. \tag{56}$$

When $\lambda > 0$, the last terms in (54)-(56) are bounded as $t \to \infty$. Thus, we can derive a necessary and sufficient condition for the expected lifetime utility $\mathcal{U}_0(X)$ and the price-dividend ratio $\mathcal{P}_0(X)$ to be finite (the *transversality condition*, see Appendix A.10):

$$(1-\gamma)\left(\bar{f} - \frac{c}{2}X^2 - \gamma\frac{\sigma^2(1+kX)^2}{2}\right) - \rho < 0. \tag{57}$$

Obsolescence keeps the infinite horizon model stationary, while at the same time preserves concavity of the expected lifetime utility, as in Corollary 2.1. We now write a modified version of Proposition 4.

Proposition 9 *There exists a unique symmetric Nash equilibrium in which the aggregate level of experimentation under competition satisfies* (57) *and solves*

$$X_n^* = \frac{\widehat{\beta}_0 \mathcal{D}_0^E(X_n^*) - \gamma k \sigma^2 \mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*}}{(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma - 1)\nu_0 \mathcal{C}_0^E(X_n^*)}, \tag{58}$$

and $x_i^* = X_n^*/n \; \forall i$. The quantity X_n^* is strictly increasing in n and $X_n^* > X^*$, $\forall n \geq 2$.

The main difference with Proposition 4 is that the rate of obsolescence modifies the wealth duration and wealth convexity that pertain to the economic value generated by the new technology. Hence the new quantities $\mathcal{D}_0^E(X_n^*)$ and $\mathcal{C}_0^E(X_n^*)$ in (58) are distinct from $\mathcal{D}_0(X_n^*)$, and $\mathcal{C}_0(X_n^*)$, as we show in Appendix A.10. With this in mind, all of our previous results go through in the infinite horizon case.

Higher technological obsolescence impacts both $\mathcal{D}_0^E(X_n^*)$ and $\mathcal{C}_0^E(X_n^*)$. Depending on which effect dominates in (58), higher obsolescence may lead to stronger or weaker aggregate experimentation. Figure 7 depicts the aggregate experimentation X_n^* as a function of the number of agents in the economy. We augment our baseline calibration with $\lambda \in [0.01, 0.15]$, which implies depreciation rates between 1% and 14% and a half-life for

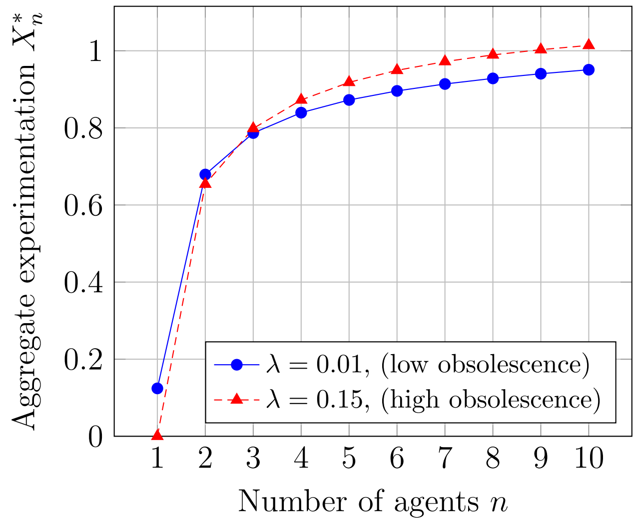

Figure 7: **The Impact of Obsolescence on Experimentation.** The two lines depict the aggregate amount of experimentation X_n^* as a function of the number of agents in the economy, for two different values of the rate of obsolescence λ.

β_t between 5 and 65 years.[23] The plot has two lines, one for $\lambda = 0.01$ (low obsolescence) and one for $\lambda = 0.15$ (high obsolescence). With our calibration, high obsolescence generates weaker experimentation when competition is low but stronger experimentation when competition is high. While we have not been able to prove the generality of this result, the illustration suggests that intense competition for technologies that are under stronger threat of becoming obsolete may exacerbate the over-experimentation result obtained in our baseline setup.

4.2 Competition and Inequality

We assumed in Section 2.2 that all agents can actively experiment with the new technology. We now depart from this and consider that there are two groups of agents. The first are $n \geq 1$ *active agents*, who are similar to those in Section 2. The second are *passive agents*, who are residual claimants in the economy and, since they are passive, we aggregate them into one agent p.

All agents derive utility from consumption with the same coefficient of risk aversion as in (2). Together they consume the aggregate consumption basket δ_t. Thus, the passive agents benefit from the introduction of the new technology by enjoying the same

[23] These are relatively wide intervals that encompass measurements of depreciation rates of technological knowledge by Park, Shin, and Park (2006) and of technology cycle times (the median age of the patents cited on the front page of a patent document) by Kayal and Waters (1999).

consumption basket (5) as the active agents. But the distinguishing feature between active and passive agents is that only the active agents are able to allocate capital to the production of the experimental good: as in our baseline model with competition, the n active agents in aggregate choose a level of experimentation $X_n = \sum_{i=1}^{n} x_i$. By doing so, they are rewarded for introducing innovation into the economy and collectively receive the rents associated with the experimental good. We show in Appendix A.1 that the size of these rents is measured as the aggregate expenditure share of the experimental good,

$$p_t^E \delta_t^E = \frac{X_n}{1 + X_n} \delta_t, \tag{59}$$

where p_t^E is the equilibrium price of the experimental good, expressed in terms of the numéraire (the aggregate good δ_t). Thus, the main difference with the baseline model of Section 2.2 is that the active agents compete for a share $X_n/(1 + X_n)$ of the total consumption stream, which itself is *increasing* in aggregate experimentation.

On the other hand, the passive agents are left with the remaining share of the aggregate consumption stream, $p_t^S \delta_t^S = \delta_t/(1 + X_n)$. Thus, while the passive agents may still enjoy the benefits of the new technology by adding its experimental good to their own consumption bundle (5), their aggregate expenditure share is driven by the aggregate experimentation X_n. One example of a passive agent might be someone who enjoys using Zoom technology but is not a Zoom shareholder.

In keeping with Section 2.2, we assume that each individual active agent's experimentation choice, x_i, dictates her own share of the pie, where the active agents' pie is now given by (59):

$$\theta_i = \frac{x_i}{\sum_{j=1}^{n} x_j} \frac{X_n}{1 + X_n} = \frac{x_i}{1 + X_n}. \tag{60}$$

This yields a modified Tullock contest, in which the success functions are $x_i/(1 + X_n)$ instead of x_i/X_n. In sum, the consumption shares of all agents in the economy are

$$\begin{cases} \frac{x_i}{1+X_n} \delta_t & \text{for any active agent } i \in \{1, ..., n\}, \\ \frac{1}{1+X_n} \delta_t & \text{for the passive agent.} \end{cases} \tag{61}$$

Note that it is unnecessary to separately define the case when no active agents experiments, i.e., $x_i = 0 \ \forall i \in \{1, ..., n\}$. In that case, one can consider all agents to be one passive agent who consumes the dividend stream provided by the status-quo economy (1).

The following is the counterpart of Proposition 4 in this modified version of the game.

Proposition 10 *There exists a unique symmetric Nash equilibrium in which the aggregate level of experimentation under competition among n agents solves*

$$X_n^* = \frac{(\widehat{\beta}_0 - \gamma k \sigma^2)\mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*} + \frac{1}{X_n^*(1+X_n^*)}}{(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma - 1)\nu_0 \mathcal{C}_0(X_n^*)}, \tag{62}$$

and $x_i^* = X_n^*/n \ \forall i$. *The quantity* X_n^* *is strictly increasing in* n *and* $X_n^* > X^*$, $\forall n \geq 1$.

For any value $n \geq 1$, the equilibrium level of experimentation that results from Proposition 10 is always higher than the level of experimentation in the baseline model. This arises because the active agents together increase their total share of consumption, $X_n^*/(1+X_n^*)$, which provides them with an additional incentive to experiment. It creates a new term in the numerator of (62), $1/[X_n^*/(1+X_n^*)]$. Since this term is positive, the left-hand side must be higher than in Proposition 4 in order to restore the equality.

In turn, when active agents experiment more here than in the baseline model, this imposes a displacement effect and added risk on the passive agent (Kogan et al., 2020). So, rent seeking and competition by active agents not only reduces the consumption share for passive agents, it worsens the quality of the dividend stream. This latter effect arises because rent seeking exposes passive agents to too much risk relative to the expected benefits, and lowers their expected utility from future consumption. In turn, then, innovation and competition may worsen income inequality (Jones and Kim, 2018).

These results are consistent with existing empirical observations. Innovation and top income inequality in developed countries tend to follow a parallel evolution. According to Aghion et al. (2015), 11 out of the 50 wealthiest individuals across US in 2015 *"are listed as inventors in a US patent and many more manage or own firms that patent."*

4.3 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 (e.g., Gort and Klepper, 1982). We return to the socially optimal case of Section 2.1 and consider that the agent can choose the experimentation level X_t at any time t, in order to maximize her expected lifetime utility.[24] Thus, she can alter the level of experimentation

[24] We were unable to solve a model in which multiple agents dynamically experiment and compete in the market. We leave this as an interesting open question for future research and analyze here the implications of dynamic experimentation in the representative agent setting.

dynamically and retains the option to expand or abandon her investment at every instant.

The agent's expected lifetime utility of consumption J satisfies the following partial differential equation at any time t:

$$0 = \max_X \left[e^{-\rho t} \frac{\delta_t^{1-\gamma}}{1-\gamma} + \mathcal{L}J(\delta_t, \widehat{\beta}_t, \nu_t, t) \right], \quad \text{where} \quad \mathcal{L}J = \frac{\mathbb{E}[dJ]}{dt}, \tag{63}$$

with boundary condition $J_a(\delta_T, \widehat{\beta}_T, \nu_T, T) = 0$ and subject to $X_t \geq 0$, $\forall t$. We assume that $c = 0$ for simplicity. In equilibrium, the CRRA utility conjecture $J(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t}\delta_t^{1-\gamma}\mathcal{P}(\widehat{\beta}_t, \nu_t, t)/(1-\gamma)$ results in a partial differential equation for the price-dividend ratio $\mathcal{P}(\cdot)$, which we relegate to Appendix A.13 for the sake of brevity. The optimal level of experimentation then follows from the first order condition on X_t.

Proposition 11 *If the problem (63) has an interior maximum, then the optimal level of experimentation at time t solves*

$$X_t^* = \frac{\widehat{\beta}_t - \gamma k \sigma^2}{\gamma k^2 \sigma^2} + \frac{\nu_t}{\gamma k^2 \sigma^2} \left(\frac{\mathcal{P}_\beta}{\mathcal{P}} - \frac{X_t^* \nu_t}{(\gamma - 1)\sigma^2 (1 + k X_t^*)^3} \frac{\mathcal{P}_{\beta\beta} - 2\mathcal{P}_\nu}{\mathcal{P}} \right). \tag{64}$$

The solution (64) 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 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 $\widehat{\beta}_t$ and ν_t.

Impact of dynamic experimentation on the real economy Dynamic experimentation provides the agent with an additional option to abandon the new technology at any point in the future, which results in more aggressive experimentation than in the static case. Although this additional option is an improvement in agent's set of choices, it can have adverse consequences on the economy. A primary consequence pertains to the volatility of aggregate consumption, which in the dynamic case evolves stochastically with X_t^*. Experimentation, thus, may amplify the volatility of consumption through the observer effect that it imposes on the economy.

Furthermore, the abandonment option embedded in dynamic experimentation has the power to decide the future of a new technology. This is best illustrated by the example

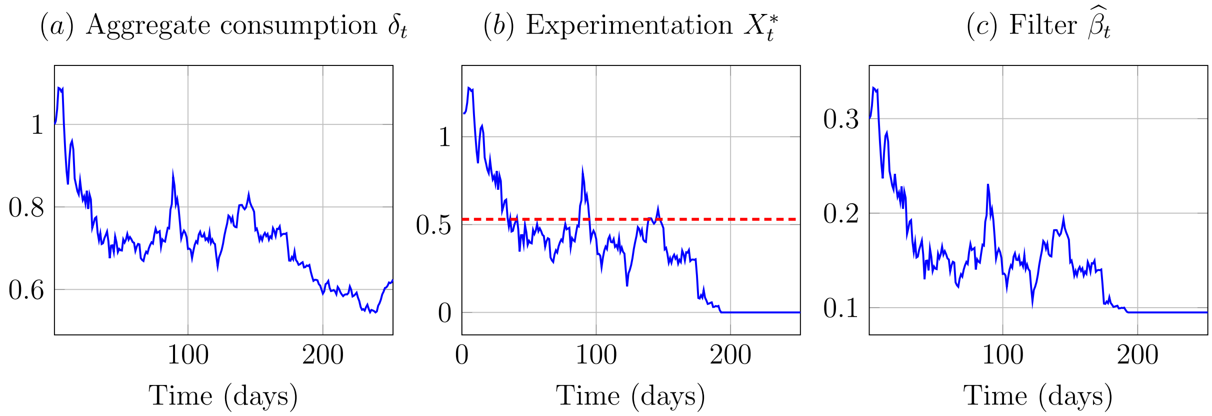

Figure 8: **Experimentation and the Abandonment of a New Technology.** Panel (a) shows a simulated path of consumption over 252 days. Panel (b) shows agent's optimal experimentation choice, X_t^*, given the observed history of consumption (the red dashed line shows the optimal experimentation level in the static case of Section 2). Panel (c) shows the filter $\widehat{\beta}_t$ which results from the agent's updating. The calibration used for this figure is only illustrative and is different than for the rest of the paper to improve numerical accuracy: $\gamma = 1.2$, $\bar{f} = 0.1$, $\widehat{\beta}_0 = 0.3$, $\nu_0 = 0.15$, $\sigma = 0.2$, $k = 2$, $\rho = 0.05$, $\delta_0 = 1$, and $T = 1$.

in Figure 8, where we consider a situation in which the agent stops experimenting just because an unusually bad stream of consumption has occurred. In this example, the agent starts with a positive prior $\widehat{\beta}_0 > 0$, which also happens to be equal to the true β. As such, this is a perfectly viable technology that can improve agent's welfare if adopted. However, lack of perfect knowledge about its productivity and learning by doing leads the agent to conclude after the unusually strong downward trend observed in the left panel that the technology is not productive. In the middle panel, the agent stops experimenting after about 200 days. This also stops the learning process: in the right panel, the estimated $\widehat{\beta}_t$ remains constant at a relatively low value once experimentation stops.

Asset Prices with Dynamic Experimentation Proposition 11 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. We relegate asset pricing details to Appendix A.13 and directly discuss here the implications of dynamic experimentation on asset pricing quantities.

Figure 9 compares the optimal level of experimentation, the risk premium and the volatility in the static case (solid lines) versus the dynamic case (dashed lines), as functions of the filter $\widehat{\beta}_t$. When the filter $\widehat{\beta}_0$ is sufficiently low, neither the "static" or the "dynamic"

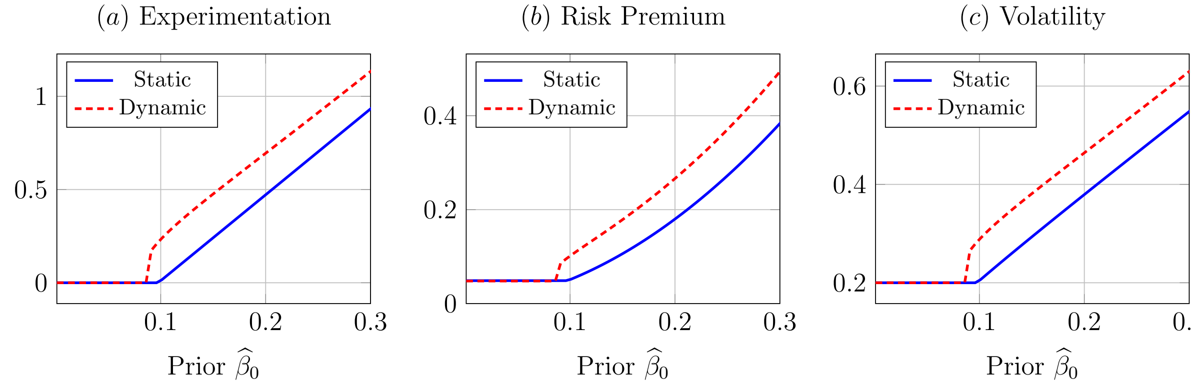

Figure 9: **Static Versus Dynamic Experimentation.** The three panels depict the optimal level of experimentation, the risk premium and the volatility of asset returns with static experimentation (solid lines) versus dynamic experimentation (dashed lines). The panels plot functions of the prior $\widehat{\beta}_0$. The parameters are the same as in Figure 8.

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 always has the option to stop or decrease later on. In contrast, the static experimenter is more cautious when fixing an initial experimentation level.

The risk premium and the volatility are generally higher with dynamic experimentation than with static experimentation, as shown in panels (b) and (c). The option to abandon allows the dynamic experimenter to experiment more aggressively, which raises the risk premium and volatility.

4.4 Competition in a Model with Recursive Preferences

In a final extension, we attempt to relax the assumption of isoelastic preferences employed in Section 2. An unfortunate feature of standard isoelastic preferences is that, by imposing equality between risk aversion and aversion to intertemporal substitution, they hide the determinant role played by the elasticity of intertemporal substitution. In Appendix A.14, we develop an alternative model with stochastic differential utility (Epstein and Zin, 1989), which separates risk aversion from aversion to intertemporal substitution.

With stochastic differential utility, a log-linear approximation of the price-dividend ratio is necessary to reach a solution (Restoy and Weil, 2010; Bansal and Yaron, 2004; Beeler and Campbell, 2012; Benzoni, Collin-Dufresne, and Goldstein, 2011). We are able to solve the alternative model using this approximation, and we show that it delivers qualitatively similar results: 1) rent-seeking competition yields over-experimentation, which in turn increases with the number of competitors in the economy; 2) technological uncertainty generates over-valuation of the risky asset; and 3) this latter effect is further exacerbated by competition.

When agents have recursive preferences, they are sensitive to long-run uncertainty about consumption growth. Section 2 shows that our setup with experimentation creates such uncertainty (which in our model is represented by technological uncertainty) and inserts a small but persistent component $\widehat{\beta}_t$ in the growth rate of consumption. In short, our model endogenously creates long-run risks (Bansal and Yaron, 2004). A growing literature initiated by Kung and Schmid (2015) and Kung (2015) provides equilibrium foundations of long-run risk through innovation and R&D. In line with this literature, Appendix A.14 shows that the market price of risk in our version of the model with recursive utility increases with technological uncertainty. Intense levels of competition and experimentation further exacerbate this. While our main observations (over-experimentation and over-valuation caused by Schumpeterian competition) remain valid in this model extension, competition and its effect on experimentation increase long-run risks. An important avenue for future research is the characterization of the interaction between Schumpeterian competition, experimentation through learning by doing, and long-run risk.

The Epstein-Zin framework further highlights the role of the elasticity of intertemporal substitution (EIS) in our results. Specifically, the value of EIS dictates how technological uncertainty affects the quantity of risk (the volatility of stock returns). For a parameter EIS higher than one (consistent with the calibration of the long-run risk literature, e.g., Bansal and Yaron, 2004), technological uncertainty magnifies stock return volatility; the opposite occurs when the EIS is lower than one.[25] These results point to a potentially important role played by the EIS (or the variation of it) during technological revolutions.

5 Conclusion

This paper proposes a risk-return perspective on Schumpeter (1934)'s evolutionary economics ideas. Using a rent-seeking game in which agents compete for a share of the consumption stream, we study how competition affects risk, wealth, and prices.

Risk arises from various sources. First, the new technology disrupts existing assets, making their future use uncertain. The second source of risk is technological uncertainty,

[25] Empirical studies disagree about reasonable values for the EIS: EIS is greater than one in Vissing-Jørgensen and Attanasio (2003), or smaller than one in Campbell (1999); Vissing-Jørgensen (2002).

which naturally arises when agents adopt and experiment with new ways of doing things. Agents' choice to experiment with the new technology creates both sources of risks, and competition for consumption share magnifies them.

What we learn from the model is that the agents' actions also affect the duration of the consumption stream, that is, the weighted-average maturity of wealth. Higher wealth duration in our model exposes agents to more technological uncertainty. As experimentation with the new technology grows, wealth duration increases. Because of agents' hedging motives, a complementarity between wealth duration and technological uncertainty decreases systematic risk and the risk premium. Sufficiently intense competition can lead to a negative risk premium, but this effect is transient due to learning. Eventually, as agents learn more about the new technology, the risk premium becomes positive and reflects mainly the persistent disruption risk that agents bear when they hold the risky asset.

Finally, not all agents in the economy benefit from the new technology: in the model, passive agents are missing the rewards of the new technology and their welfare decreases monotonically when an increasing number of active agents fight for a share of the pie. Thus, Schumpeterian competition can worsen income inequality.

References

Acemoglu, D. (2009). *Introduction to modern economic growth.* Princeton, NJ: Princeton University Press.

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy 128*(6), 2188–2244.

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. 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.

Alexeev, M. and J. Leitzel (1996). Rent shrinking. *Southern Economic Journal 62*, 90–113.

Andrei, D. and M. Hasler (2015). Investor attention and stock market volatility. *The review of financial studies 28*(1), 33–72.

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.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–1509.

Barlevy, G. (2004). The cost of business cycles under endogenous growth. *American Economic Review 94*(4), 964–990.

Barlevy, G. (2007). On the cyclicality of research and development. *American Economic Review 97*(4), 1131–1164.

Baye, M. and H. Hoppe (2003). The strategic equivalence of rent-seeking, innovation, and patent-race games. *Games and Economic Behavior 44*, 217–226.

Baye, M. R., D. Kovenock, and C. G. De Vries (1994). The solution to the tullock rent-seeking game when r¿ 2: Mixed-strategy equilibria and mean dissipation rates. *Public Choice 81*(3), 363–380.

Beeler, J. and J. Y. Campbell (2012). The long-run risks model and aggregate asset prices: An empirical assessment. *Critical Finance Review 1*(1), 141–182.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2011). Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics 101*(3), 552–573.

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.

Breeden, D. T. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics 7*, 265–296.

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

Brennan, M. J. and Y. Xia (2001). Stock price volatility and equity premium. *Journal of monetary Economics 47*(2), 249–283.

Campbell, J. Y. (1999). Asset prices, consumption, and the business cycle. *Handbook of macroeconomics 1*, 1231–1303.

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

Chowdhury, S. M. and R. M. Sheremeta (2011). A generalized tullock contest. *Public Choice 147*(3), 413–420.

Chung, T.-Y. (1996). Rent-seeking contest when the prize increases with aggregate efforts. *Public Choice 87*(1), 55–66.

Clark, D. J. and C. Riis (1998). Contest success functions: an extension. *Economic Theory 11*(1), 201–204.

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.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review 106*(3), 664–98.

Corhay, A., H. Kung, and L. Schmid (2020). Competition, markups, and predictable returns. *The Review of Financial Studies 33*(12), 5906–5939.

D'Aspremont, C. and A. Jacquemin (1988). Cooperative and noncooperative r&d in duopoly with spillovers. *American Economic Review 78*(5), 1133–1137.

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics 135*(2), 561–644.

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.

Devereux, M. B. and A. Sutherland (2011). Country portfolios in open economy macro-models. *Journal of the european economic Association 9*(2), 337–369.

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. and L. G. Epstein (1992). Asset pricing with stochastic differential utility. *Review of Financial Studies 5*(3), 411–436.

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

Dumas, B., A. Kurshev, and R. Uppal (2009). Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility. *The Journal of Finance 64*(2), 579–629.

Dumas, B. and E. Luciano (2017, December). *The Economics of Continuous-Time Finance*, Volume 1 of *MIT Press Books*. The MIT Press.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–969.

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.

Fudenberg, D. and J. Tirole (1984). The fat-cat effect, the puppy-dog ploy, and the lean and hungry look. *American Economic Review 74*(2), 361–366.

Futia, C. A. (1980). Schumpeterian competition. *The Quarterly Journal of Economics 94*(4), 675–695.

Gabaix, X., J.-M. Lasry, P.-L. Lions, and B. Moll (2016). The dynamics of inequality. *Econometrica 84*(6), 2071–2111.

Gârleanu, N., L. Kogan, and S. Panageas (2012). Displacement risk and asset returns. *Journal of Financial Economics 105*(3), 491–510.

Garleanu, N., S. Panageas, and J. Yu (2012). Technological growth and asset pricing. *The Journal of Finance 67*(4), 1265–1292.

Gennotte, G. (1986, July). Optimal portfolio choice under incomplete information. *Journal of Finance 41*(3), 733–46.

Gonçalves, A. S. (2021). The short duration premium. *Journal of Financial Economics*.

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

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.

Grullon, G., Y. Larkin, and R. Michaely (2019). Are us industries becoming more concentrated? *Review of Finance 23*(4), 697–743.

Gutiérrez, G. and T. Philippon (2018). Ownership, concentration, and investment. In *AEA Papers and Proceedings*, Volume 108, pp. 432–37.

Helpman, E. and A. Razin (1978). *A Theory of International Trade under Uncertainty*. Academic Press.

Hirshleifer, J. (1989). Conflict and rent-seeking success functions: Ratio vs. difference models of relative success. *Public Choice 41*, 101–112.

Hoberg, G. and G. Phillips (2020). Real and financial industry booms and busts. *Forthcoming, Journal of Finance*.

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

Jones, C. I. and J. Kim (2018). A schumpeterian model of top income inequality. *Journal of Political Economy 126*(5), 1785–1826.

Kamien, M. I. and N. L. Schwartz (1975). Market structure and innovation: A survey. *Journal of economic literature 13*(1), 1–37.

Kamien, M. I. and N. L. Schwartz (1982). *Market structure and innovation*. Cambridge University Press.

Karakashian, S. (2015). A software patent war: The effects of patent trolls on startup companies, innovation, and entrepreneurship. *Hastings Bus. LJ 11*, 119.

Kayal, A. A. and R. C. Waters (1999). An empirical evaluation of the technology cycle time indicator as a measure of the pace of technological progress in superconductor technology. *IEEE Transactions on Engineering Management 46*(2), 127–131.

Klepper, S. (1996). Entry, exit, growth, and innovation over the product life cycle. *American Economic Review 86*(3), 562–583.

Kogan, L., D. Papanikolaou, and N. Stoffman (2020). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy 128*(3), 855–906.

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

Koren, M. and S. Tenreyro (2013). Technological diversification. *American Economic Review 103*(1), 378–414.

Kung, H. (2015). Macroeconomic linkages between monetary policy and the term structure of interest rates. *Journal of Financial Economics 115*(1), 42–57.

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

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? a duration-based explanation of the value premium. *The Journal of Finance 62*(1), 55–92.

Lettau, M. and J. A. Wachter (2011). The term structures of equity and interest rates. *Journal of Financial Economics 101*(1), 90–113.

Liptser, R. S. and A. N. Shiryaev (2001). *Statistics of Random Processes* (Second ed.), Volume II. Applications. Springer-Verlag.

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

Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica: journal of the Econometric Society*, 1429–1445.

Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.

Markowitz, H. (1952). Portfolio selection. *The journal of finance 7*(1), 77–91.

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.

Munk, C. (2013). *Financial asset pricing theory.* OUP Oxford.

Nelson, R. R. and S. G. Winter (2009). *An evolutionary theory of economic change.* Harvard University Press.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–08.

Park, G., J. Shin, and Y. Park (2006). Measurement of depreciation rate of technological knowledge: Technology cycle time approach. *Journal of Scientific & Industrial Research 65*, 121–127.

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.

Pohl, W., K. Schmedders, and O. Wilms (2018). Higher order effects in asset pricing models with long-run risks. *The Journal of Finance 73*(3), 1061–1111.

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.

Reinganum, J. F. (1989). *The Timing of Innovation: Research, Development and Diffusion*, Chapter 14, pp. 849–908. The Handbook of Industrial Organizationl. Elsevier.

Restoy, F. and P. Weil (2010). Approximate equilibrium asset prices. *Review of Finance 15*(1), 1–28.

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

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.

Shaver, L. (2012). Illuminating innovation: From patent racing to patent war. *Wash. & Lee L. Rev. 69*, 1891.

Skaperdas, S. (1996). Contest success functions. *Economic theory 7*(2), 283–290.

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.

Trappey, C. V., A. J. Trappey, and Y.-H. Wang (2016). Are patent trade wars impeding innovation and development? *World Patent Information 46*, 64–72.

Tullock, G. (1980). Efficient rent seeking. In J. M. Buchanan, R. D. Tollison, and G. Tullock (Eds.), *Toward a theory of the rent-seeking society.* Texas A&M University Press: College Station.

Van Binsbergen, J. H. (2020). Duration-based stock valuation: Reassessing stock market performance and volatility. Technical report, National Bureau of Economic Research.

Veldkamp, L. L. (2011). *Information choice in macroeconomics and finance.* Princeton University Press.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy 110*(4), 825–853.

Vissing-Jørgensen, A. and O. P. Attanasio (2003). Stock-market participation, intertemporal substitution, and risk-aversion. *The American Economic Review 93*(2), 383–391.

Weber, M. (2018). Cash flow duration and the term structure of equity returns. *Journal of Financial Economics 128*(3), 486–503.

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

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

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.

A Appendix

A.1 Intratemporal Choice and Expenditure Shares

Consider the intratemporal (within period and across goods) choice of the agent at any time $t \in [0, T]$. Define the numéraire as δ_t and let p_t^S (p_t^E) denote the price of the status-quo (experimental) good. Thus

$$\delta_t = p_t^S \delta_t^S + p_t^E \delta_t^E. \tag{A.1}$$

Given an equilibrium level of experimentation X, the expenditure shares $p_t^S \delta_t^S$ and $p_t^E \delta_t^E$ are determined by the following maximization problem subject to (A.1):

$$\max_{\delta_t^S, \delta_t^E} \frac{\left[\delta_t^S \left(\delta_t^E\right)^X\right]^{1-\gamma}}{1-\gamma}. \tag{A.2}$$

Solving for δ_t^E in (A.1) yields $\delta_t^E = (\delta_t - p_t^S \delta_t^S)/p_t^E$. Replacing this into (A.2) and taking the first-order condition with respect to δ_t^S yields $p_t^S \delta_t^S = \delta_t/(1+X)$. Thus, the agent spends a fraction $1/(1+X)$ on the status-quo good and a fraction $X/(1+X)$ on the experimental good.

A.2 Proof of Proposition 1

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

Theorem A.1 ***(Liptser and Shiryaev, 2001, theorem 12.1, p. 22)*** *Let $(\theta, \xi) = (\theta_t, \xi_t)$, $0 \le t \le T$, be a continuous random diffusion-type, conditionally Gaussian process with:*

$$d\theta_t = [a_0(t,\xi) + a_1(t,\xi)\theta_t]dt + b_1(t,\xi)dW_1(t) + b_2(t,\xi)dW_2(t), \tag{A.3}$$

$$d\xi_t = [A_0(t,\xi) + A_1(t,\xi)\theta_t]dt + B(t,\xi)dW_2(t), \tag{A.4}$$

where W_1 and W_2 are mutually independent standard Brownian motions. If conditions (11.4)-(11.8) and (12.16)-(12.18) in Liptser and Shiryaev (2001) are satisfied and the conditional distribution $P(\theta_0 \le a|\xi_0)$ is Gaussian, $N(m_0, \nu_0)$, then the a posteriori mean $m_t = \mathbb{E}(\theta_t|\mathcal{F}_t^\xi)$ and the a posteriori variance $\nu_t = \mathbb{E}[(\theta_t - m_t)^2|\mathcal{F}_t^\xi]$ satisfy equations

$$dm_t = [a_0(t,\xi) + a_1(t,\xi)m_t]dt + \frac{b_2(t,\xi)B(t,\xi) + \nu_t A_1(t,\xi)}{B^2(t,\xi)}[d\xi_t - (A_0(t,\xi) + A_1(t,\xi)m_t)dt], \tag{A.5}$$

$$\frac{d\nu_t}{dt} = 2a_1(t,\xi)\nu_t + b_1^2(t,\xi) + b_2^2(t,\xi) - \left(\frac{b_2(t,\xi)B(t,\xi) + \nu_t A_1(t,\xi)}{B(t,\xi)}\right)^2, \tag{A.6}$$

subject to the conditions $m_0 = \mathbb{E}(\theta_0|\xi_0)$, $\nu_0 = \mathbb{E}[(\theta_0 - m_0)^2|\xi_0]$.

In the present setup, the unobservable variable is the constant β. Hence, $a_0 = a_1 = b_1 = b_2 = 0$. The observable process is δ_t. Applying Itô's lemma on $\ln \delta_t$ yields

$$d\ln\delta_t = \left[\bar{f} + \beta X - \frac{c}{2}X^2 - \frac{1}{2}\sigma^2(1 + k^2X^2 + 2k\varphi X)\right]dt + \sigma\sqrt{1 + k^2X^2 + 2k\varphi X}\,dW_t, \tag{A.7}$$

and thus

$$A_0 = \bar{f} - \frac{c}{2}X^2 - \frac{1}{2}\sigma^2(1 + k^2X^2 + 2k\varphi X), \quad A_1 = X, \quad B = \sigma\sqrt{1 + k^2X^2 + 2k\varphi X}. \tag{A.8}$$

Because a_0, a_1, b_1, b_2, A_0, A_1, and B are constants, all the conditions of Theorem A.1 are satisfied. Direct application of (A.5) then yields

$$d\widehat{\beta}_t = \frac{\nu_t X}{\sigma^2(1 + k^2X^2 + 2k\varphi X)} \left[d\ln\delta_t - \left(\bar{f} + \widehat{\beta}_t X - \frac{cX^2}{2} - \frac{\sigma^2(1 + k^2X^2 + 2k\varphi X)}{2} \right) dt \right] \tag{A.9}$$

$$= \frac{\nu_t X}{\sigma^2(1 + k^2X^2 + 2k\varphi X)} \left[X(\beta - \widehat{\beta}_t)dt + \sigma\sqrt{1 + k^2X^2 + 2k\varphi X}\, dW_t \right] \tag{A.10}$$

$$= \frac{\nu_t X}{\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}} \underbrace{\left[\frac{X(\beta - \widehat{\beta}_t)}{\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}} dt + dW_t \right]}_{\equiv d\widehat{W}_t}. \tag{A.11}$$

Eq. (A.11) defines $d\widehat{W}_t$ as a shock proportional to the "surprise" change in consumption that occurs in (A.9) in the square brackets. The division of this surprise by $\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}$ in (A.11) ensures that $\widehat{W}_t$ is a standard Brownian motion with respect to the agent's filtration $\{\mathcal{F}_t^\delta\}$. Eq. (A.6) further implies

$$\frac{d\nu_t}{dt} = -\frac{X^2\nu_t^2}{\sigma^2(1 + k^2X^2 + 2k\varphi X)}. \quad \square \tag{A.12}$$

A.3 Proof of Proposition 2

Part (a) of Proposition 2 Part (a) follows from the theory of affine processes (Duffie et al., 2003). An application of Fubini's theorem yields

$$\mathcal{U}_0(X) = \mathbb{E}_0 \left[\int_0^T e^{-\rho t} \frac{\delta_t^{1-\gamma}}{1-\gamma} dt \right] = \frac{1}{1-\gamma} \int_0^T e^{-\rho t} \mathbb{E}_0[\delta_t^{1-\gamma}] dt. \tag{A.13}$$

Note that the expectation $\mathbb{E}_0[\delta_t^{1-\gamma}]$ can be written as

$$\mathbb{E}_0[\delta_t^{1-\gamma}] = \mathbb{E}_0 \left[e^{(1-\gamma)\ln\delta_t} \right]. \tag{A.14}$$

The exponent in (A.14), $(1-\gamma)\ln\delta_t$, is an affine function of the vector $[\ln\delta_t \;\; \widehat{\beta}_t]'$, whose dynamics can be written under an affine form:

$$\begin{bmatrix} d\ln\delta_t \\ d\widehat{\beta}_t \end{bmatrix} = \left(\begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1 + k^2X^2 + 2k\varphi X) - \frac{c}{2}X^2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & X \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \ln\delta_t \\ \widehat{\beta}_t \end{bmatrix} \right) dt + \begin{bmatrix} \sigma\sqrt{1 + k^2X^2 + 2k\varphi X} \\ A(t) \end{bmatrix} d\widehat{W}_t, \tag{A.15}$$

where $A(t)$ is a function of time that results from Proposition 1 and Eq. (13):

$$A(t) = \frac{\nu_0 X\sigma\sqrt{1 + k^2X^2 + 2k\varphi X}}{\nu_0 X^2 t + \sigma^2(1 + k^2X^2 + 2k\varphi X)}. \tag{A.16}$$

Because the diffusion of $\widehat{\beta}$ depends on time, (A.15) is a *time-inhomogeneous* multi-factor affine process (Filipović, 2005). Define:

$$K_0 \equiv \begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1 + k^2X^2 + 2k\varphi X) - \frac{c}{2}X^2 \\ 0 \end{bmatrix}, \tag{A.17}$$

$$K_1 \equiv \begin{bmatrix} 0 & X \\ 0 & 0 \end{bmatrix}, \tag{A.18}$$

$$H_0(t) \equiv \begin{bmatrix} \sigma\sqrt{1 + k^2X^2 + 2k\varphi X} \\ A(t) \end{bmatrix} \begin{bmatrix} \sigma\sqrt{1 + k^2X^2 + 2k\varphi X} \\ A(t) \end{bmatrix}'. \tag{A.19}$$

Let $t, s \in [0, T]$ such that $t \leq s \leq T$, and define $\tau = s - t$. In order to compute the expectation $\mathbb{E}_t[\delta_s^{1-\gamma}] = \mathbb{E}_t[e^{(1-\gamma)\ln\delta_s}]$, we conjecture an exponential-affine solution of the form

$$\mathbb{E}_t\left[\delta_s^{1-\gamma}\right] = e^{\alpha_0(\tau) + \alpha_1(\tau)\ln\delta_t + \alpha_2(\tau)\widehat{\beta}_t}, \tag{A.20}$$

for some coefficient functions $\alpha_j(\cdot)$, $j = 0, 1, 2$, which satisfy

$$\begin{bmatrix} \alpha_1'(\tau) \\ \alpha_2'(\tau) \end{bmatrix} = K_1^\top \begin{bmatrix} \alpha_1(\tau) \\ \alpha_2(\tau) \end{bmatrix} \tag{A.21}$$

$$\alpha_0'(\tau) = K_0^\top \begin{bmatrix} \alpha_1(\tau) \\ \alpha_2(\tau) \end{bmatrix} + \frac{1}{2}\begin{bmatrix} \alpha_1(\tau) & \alpha_2(\tau) \end{bmatrix} H_0(t) \begin{bmatrix} \alpha_1(\tau) \\ \alpha_2(\tau) \end{bmatrix}, \tag{A.22}$$

with boundary conditions $\alpha_0(0) = 0$, $\alpha_1(0) = 1 - \gamma$, and $\alpha_2(0) = 0$. This is a system of Riccati ordinary differential equations (Duffie et al., 2003). Eq. (A.21) has a straightforward solution:

$$\alpha_1(\tau) = 1 - \gamma \tag{A.23}$$

$$\alpha_2(\tau) = (1 - \gamma)X\tau, \tag{A.24}$$

which can be now inserted in the remaining Riccati equation (A.22):

$$\alpha_0'(\tau) = \begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1 + k^2X^2 + 2k\varphi X) - \frac{c}{2}X^2 & 0 \end{bmatrix} \begin{bmatrix} 1 - \gamma \\ (1 - \gamma)X\tau \end{bmatrix} + \frac{1}{2}\begin{bmatrix} 1 - \gamma & (1 - \gamma)X\tau \end{bmatrix} H_0(t) \begin{bmatrix} 1 - \gamma \\ (1 - \gamma)X\tau \end{bmatrix}, \tag{A.25}$$

leading to

$$\alpha_0'(\tau) = \frac{1}{2}(\gamma - 1)\{-2\bar{f} + cX^2 + \sigma^2(1 + k^2X^2 + 2k\varphi X) + (\gamma - 1)[\sigma\sqrt{1 + k^2X^2 + 2k\varphi X} + X\tau A(t)]^2\}. \tag{A.26}$$

Replacing $A(t)$ from (A.16) and t by $s - \tau$ yields

$$\begin{aligned} \alpha_0'(\tau) = {} & \frac{1}{2}(\gamma - 1)\left[-2\bar{f} + cX^2 + \sigma^2(1 + k^2X^2 + 2k\varphi X)\right] \\ & + \frac{1}{2}(\gamma - 1)^2\sigma^2(1 + k^2X^2 + 2k\varphi X)\left(\frac{\sigma^2(1 + k^2X^2 + 2k\varphi X) + X^2\nu_0 s}{\sigma^2(1 + k^2X^2 + 2k\varphi X) + X^2\nu_0(s - \tau)}\right)^2, \end{aligned} \tag{A.27}$$

with boundary condition $\alpha_0(0) = 0$. Solving this equation by integration is straightforward because only the last term depends on τ. Further using (13) to replace ν_0 with a function of ν_t yields a solution for α_0 for any $t, s \in [0, T]$ with $s \geq t$ and $\tau = s - t$:

$$\alpha_0(\tau) = \frac{1}{2}(\gamma - 1)\left[-2\bar{f} + cX^2 + \gamma\sigma^2(1 + k^2X^2 + 2k\varphi X)\right]\tau + \frac{1}{2}X^2\nu_t(\gamma - 1)^2\tau^2. \qquad \text{(A.28)}$$

Written at time 0 and with $\tau = t$, the solution of the Riccati system (A.21)-(A.22) is

$$\alpha_0(t) = \frac{1}{2}(\gamma - 1)\left[-2\bar{f} + cX^2 + \gamma\sigma^2(1 + k^2X^2 + 2k\varphi X)\right]t + \frac{1}{2}X^2\nu_0(\gamma - 1)^2t^2, \qquad \text{(A.29)}$$

$$\alpha_1(t) = 1 - \gamma, \qquad \text{(A.30)}$$

$$\alpha_2(t) = (1 - \gamma)Xt, \qquad \text{(A.31)}$$

which can now be replaced into the conjecture (A.20). After multiplication with $e^{-\rho t}$, we obtain:

$$e^{-\rho t}\mathbb{E}_0\left[\delta_t^{1-\gamma}\right] = \exp\left[-\rho t + \alpha_0(t) + \alpha_1(t)\ln\delta_0 + \alpha_2(t)\widehat{\beta}_0\right] \qquad \text{(A.32)}$$

$$= \exp\left[(1 - \gamma)\ln\delta_0 + \kappa(X, \widehat{\beta}_0)t + \frac{(\gamma - 1)^2X^2\nu_0}{2}t^2\right], \qquad \text{(A.33)}$$

where $\kappa(X, \widehat{\beta}_0)$ is defined as in (14):

$$\kappa(X, \widehat{\beta}_0) \equiv (1 - \gamma)\left(\bar{f} + X\widehat{\beta}_0 - \gamma\frac{\sigma^2(1 + k^2X^2 + 2k\varphi X)}{2} - \frac{c}{2}X^2\right) - \rho. \qquad \text{(A.34)}$$

Replacing (A.33) into (A.13) yields Eq. (15) of Proposition 2:

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1 - \gamma}\int_0^T \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma - 1)^2X^2\nu_0}{2}t^2\right]dt. \qquad \text{(A.35)}$$

Part (b) of Proposition 2 Part (b) follows from standard asset pricing theory (Duffie, 2010); see also Appendix A.8. In this setup with time-additive utility, we define the stochastic discount factor from the optimal consumption plan of the aggregate consumer as

$$\xi_t = e^{-\rho t}\frac{u'(\delta_t)}{u'(\delta_0)} = e^{-\rho t}\frac{\delta_t^{-\gamma}}{\delta_0^{-\gamma}}. \qquad \text{(A.36)}$$

The equilibrium price of the risky asset at time 0 is then

$$P_0(X) = \mathbb{E}_0\left[\int_0^T \xi_t\delta_t dt\right] = \delta_0^{\gamma}\int_0^T e^{-\rho t}\mathbb{E}_0[\delta_t^{1-\gamma}]dt, \qquad \text{(A.37)}$$

and we recognize the same integral as in (A.13). Using (A.33), the price-dividend ratio is then

$$\mathcal{P}_0(X) = \frac{P_0(X)}{\delta_0} = \int_0^T \exp\left[\kappa(X, \widehat{\beta}_0)t + \frac{(\gamma - 1)^2}{2}X^2\nu_0t^2\right]dt. \qquad \text{(A.38)}$$

The following relation holds between expected lifetime utility and the price-dividend ratio

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X). \quad \square \tag{A.39}$$

A.4 Proof of Corollary 2.1

To prove Corollary 2.1, we write the lifetime utility under the following form:

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\int_0^T e^{y(\widehat{\beta}_0,\nu_0,X,t)}dt, \tag{A.40}$$

and consider the second partial derivative of the function $y(\widehat{\beta}_0, \nu_0, X, t)$ with respect to X,

$$\frac{\partial^2 y(\widehat{\beta}_0, \nu_0, X, t)}{\partial X^2} = t(\gamma-1)\left[c + \gamma\sigma^2 k^2 + (\gamma-1)\nu_0 t\right], \tag{A.41}$$

which follows from taking the second derivative of the exponent in (A.38).

If $\gamma > 1$, the second derivative $\partial^2 y/\partial X^2$ is strictly positive $\forall t \in (0,T] \Rightarrow$ the function y is strictly convex in X. Thus, e^y is log-convex (Boyd and Vandenberghe, 2004, p. 104). Log-convexity is preserved under sums and integrals (Boyd and Vandenberghe, 2004, p. 105-106) $\Rightarrow$ if $\gamma > 1$ the integral $\int_0^T e^{y(\widehat{\beta}_0,\nu_0,X,t)}dt$, which represents the price-dividend ratio $\mathcal{P}_0(X)$ as defined in (A.38), is log-convex in X. Since a log-convex function is strictly convex, $\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X)$ and $\gamma > 1$ imply that $\mathcal{U}_0(X)$ is strictly concave in X. $\square$

Log-convexity of $\mathcal{P}_0(X)$ implies that the function $\partial \ln \mathcal{P}_0(X)/\partial X$ is increasing in X. This result is important for the existence and uniqueness of an equilibrium level of experimentation with competition (Proposition 4). More precisely, using Definition 1,

$$\frac{\partial \ln \mathcal{P}_0(X)}{\partial X} = \frac{1}{\mathcal{P}_0(X)}\frac{\partial \mathcal{P}_0(X)}{\partial X} = \frac{\partial \kappa(X,\widehat{\beta}_0)}{\partial X}\mathcal{D}_0(X) + (\gamma-1)^2 X\nu_0\mathcal{C}_0(X) \tag{A.42}$$

$$= (1-\gamma)\left[\widehat{\beta}_0 - \gamma k\sigma^2(\varphi + kX) - cX\right]\mathcal{D}_0(X) + (\gamma-1)^2 X\nu_0\mathcal{C}_0(X) \tag{A.43}$$

$$= (1-\gamma)(\widehat{\beta}_0 - \gamma k\varphi\sigma^2)\mathcal{D}_0(X) + X\left[(\gamma-1)(\gamma k^2\sigma^2 + c)\mathcal{D}_0(X) + (\gamma-1)^2\nu_0\mathcal{C}_0(X)\right]. \tag{A.44}$$

Both $\mathcal{D}_0(X)$ and $\mathcal{C}_0(X)$ are weighted averages and thus their values are finite: $\mathcal{D}_0(X) \in (0,T]$ and $\mathcal{C}_0(X) \in (0,T^2]$. We can therefore find the limits of $\partial \ln \mathcal{P}_0(X)/\partial X$ at $X = 0$ and $X \to \infty$. If $X = 0$ then $\partial \ln \mathcal{P}_0(X)/\partial X = (1-\gamma)(\widehat{\beta}_0 - \gamma k\varphi\sigma^2)\mathcal{D}_0(0)$. Furthermore, when $\gamma > 1$, $\lim_{X\to\infty}\frac{\partial \ln \mathcal{P}_0(X)}{\partial X} = \infty$, and thus $\partial \ln \mathcal{P}_0(X)/\partial X$ increases in X from $(1-\gamma)(\widehat{\beta}_0 - \gamma k\varphi\sigma^2)\mathcal{D}_0(0)$ to ∞ (a statement we make on page 16).

A.5 Proof of Proposition 3

In the first-order condition (20), compute

$$\frac{\partial \kappa(X,\widehat{\beta}_0)}{\partial X} = (1-\gamma)\left(\widehat{\beta}_0 - \gamma\sigma^2 k(\varphi + kX) - cX\right), \tag{A.45}$$

which, after replacement into (20) leads to

$$(\widehat{\beta}_0 - \gamma k\varphi\sigma^2)\mathcal{D}_0(X) = X\left[(\gamma k^2\sigma^2 + c)\mathcal{D}_0(X) + (\gamma-1)\nu_0\mathcal{C}_0(X)\right]. \tag{A.46}$$

This equation in X has a positive solution only if $\widehat{\beta}_0 > \gamma k\varphi\sigma^2$, in which case the optimal level of experimentation solves the following implicit equation:

$$X^* = \frac{(\widehat{\beta}_0 - \gamma k\varphi\sigma^2)\mathcal{D}_0(X^*)}{(\gamma k^2\sigma^2 + c)\mathcal{D}_0(X^*) + (\gamma-1)\nu_0\mathcal{C}_0(X^*)}. \quad \square \tag{A.47}$$

A.6 Experimentation as a Mean-Variance Tradeoff

We will show here that the optimal level of experimentation in (21) results from a mean-variance tradeoff, whereby the agent prefers higher expected future consumption at each t, $\{\mathbb{E}_0[\delta_t], \ t \in (0,T]\}$, and dislikes higher variance of future consumption at each t, $\{\mathrm{Var}_0[\delta_t], \ t \in (0,T]\}$.

The expectation $\mathbb{E}_0[\delta_t]$ is computed using the same method as in Appendix A.3, with the sole difference that the boundary conditions for the Riccati system (A.21)-(A.22) are $\alpha_0(0)=0$, $\alpha_1(0)=1$, and $\alpha_2(0)=0$. This yields

$$\mathbb{E}_0[\delta_t] = \delta_0 \exp\left[\left(\bar{f} + X\widehat{\beta}_0 - \frac{c}{2}X^2\right)t + \frac{X^2\nu_0}{2}t^2\right]. \tag{A.48}$$

The variance $\mathrm{Var}_0[\delta_t]$ is

$$\mathrm{Var}_0[\delta_t] = \mathbb{E}_0\left[\delta_t^2\right] - (\mathbb{E}_0[\delta_t])^2, \tag{A.49}$$

where (using the same method as in Appendix A.3, with the sole difference that the boundary conditions for the Riccati system (A.21)-(A.22) are $\alpha_0(0)=0$, $\alpha_1(0)=2$, and $\alpha_2(0)=0$):

$$\mathbb{E}_0\left[\delta_t^2\right] = \delta_0^2 \exp\left[2\left(\bar{f} + X\widehat{\beta}_0 - \frac{c}{2}X^2 + \frac{\sigma^2(1+k^2X^2+2k\varphi X)}{2}\right)t + 2X^2\nu_0 t^2\right], \tag{A.50}$$

and thus

$$\mathrm{Var}_0[\delta_t] = \mathbb{E}_0[\delta_t]^2\left(\exp\left[\sigma^2(1+k^2X^2+2k\varphi X)t + X^2\nu_0 t^2\right] - 1\right). \tag{A.51}$$

The expected lifetime utility, after replacing $\kappa(X,\widehat{\beta}_0)$ from (14), is

$$\mathcal{U}_0(X) = \int_0^T \frac{e^{-\rho t}\delta_0^{1-\gamma}}{1-\gamma}\exp\left[(1-\gamma)\left(\bar{f} + X\widehat{\beta}_0 - \gamma\frac{\sigma^2(1+k^2X^2+2k\varphi X)}{2} - \frac{c}{2}X^2\right)t + \frac{(\gamma-1)^2X^2\nu_0}{2}t^2\right]dt \tag{A.52}$$

$$= \int_0^T e^{-\rho t}\frac{\mathbb{E}_0[\delta_t]^{1-\gamma}}{1-\gamma}\left(\exp\left[\sigma^2(1+k^2X^2+2k\varphi X)t + X^2\nu_0 t^2\right]\right)^{\frac{1}{2}\gamma(\gamma-1)}dt \tag{A.53}$$

$$= \int_0^T e^{-\rho t}\frac{\mathbb{E}_0[\delta_t]^{1-\gamma}}{1-\gamma}\left(\frac{\mathrm{Var}_0[\delta_t]}{\mathbb{E}_0[\delta_t]^2} + 1\right)^{\frac{1}{2}\gamma(\gamma-1)}dt, \tag{A.54}$$

where the last equality follows from (A.51). It is straightforward to verify that, for any value of $\gamma > 0$, (A.54) strictly increases with $\mathbb{E}_0[\delta_t]$ and strictly decreases with $\mathrm{Var}_0[\delta_t]$.

The agent's experimentation choice X impacts both $\mathbb{E}_0[\delta_t]$ and $\mathrm{Var}_0[\delta_t]$. Eqs. (A.48) and (A.51) imply that, at $X = 0$, both $\mathbb{E}_0[\delta_t]$ and $\mathrm{Var}_0[\delta_t]$ strictly increase with X, consistent with the idea that adopting a new technology increases expected growth, but also makes the future more risky. Thus, when choosing the amount of experimentation, the risk-averse agent trades off a higher path of expected consumption against a higher path of variance of future consumption, consistent with a classic mean-variance tradeoff (Markowitz, 1952).

A.7 Proof of Proposition 4

We begin by showing that the equilibrium cannot involve all agents $i \in \{1, \ldots, n\}$ choosing $x_i = 0$ and sharing the dividend stream equally as in (23). Suppose all other agents besides agent j choose $x_i = 0$. Based on (23), if agent j chooses $x_j = \epsilon > 0$, $\theta_j = 1$. Even for arbitrarily small ϵ (i.e., $\epsilon \to 0$), agent j gains the entire consumption stream δ_t. Since this represents a discrete jump in agent j's payoff, there exists a sufficiently small ϵ for which the agent's gain from the discrete jump dominates any continuous change in $\mathcal{U}_0\left(\sum_{i=1}^n x_i\right)$, and thus it is a profitable deviation. So, the equilibrium cannot involve $x_i = 0$ for all $i \in \{1, \ldots, n\}$.

The rest of the proof starts from (29):

$$\frac{(\gamma-1)(n-1)}{X_n} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{A.55}$$

As discussed in the text and in Appendix A.4, when $n \geq 2$ this equation in X_n has a unique positive solution. On the right-hand side of (A.55), we have

$$\frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n} = \frac{1}{\mathcal{P}_0(X_n)} \frac{\partial \mathcal{P}_0(X_n)}{\partial X_n} = \frac{\partial \kappa(X_n, \widehat{\beta}_0)}{\partial X_n} \mathcal{D}_0(X_n) + (\gamma-1)^2 X_n \nu_0 \mathcal{C}_0(X_n) \tag{A.56}$$

$$= (1-\gamma)\left[\widehat{\beta}_0 - \gamma k \sigma^2(\varphi + kX_n) - cX_n\right] \mathcal{D}_0(X_n) + (\gamma-1)^2 X_n \nu_0 \mathcal{C}_0(X_n) \tag{A.57}$$

$$= (1-\gamma)(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X_n) + X_n\left[(\gamma-1)(\gamma k^2\sigma^2 + c)\mathcal{D}_0(X_n) + (\gamma-1)^2 \nu_0 \mathcal{C}_0(X_n)\right], \tag{A.58}$$

and thus (A.55) yields

$$\frac{n-1}{X_n} + (\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X_n) = X_n\left[(\gamma k^2 \sigma^2 + c)\mathcal{D}_0(X_n) + (\gamma-1)\nu_0 \mathcal{C}_0(X_n)\right], \tag{A.59}$$

which leads to Eq. (30):

$$X_n^* = \frac{(\widehat{\beta}_0 - \gamma k \varphi \sigma^2)\mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*}}{(\gamma k^2\sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma-1)\nu_0 \mathcal{C}_0(X_n^*)}. \tag{A.60}$$

When $n = 1$, we are back to the one-agent economy and we recover the social optimum of Proposition 3. Since the left-hand side of (A.55) strictly increases in n, the equilibrium X_n^* strictly increases in n and $X_n^* > X^*$, $\forall n \geq 2$ (see panel (a) of Figure 1 for an illustration).

We provide here another proof of the result that X_n^* increases with n (argument used in the text after Proposition 4.) The equilibrium X_n in an economy with n agents solves:

$$\frac{(\gamma-1)(n-1)}{X_n^*} = \frac{1}{\mathcal{U}_0(X_n^*)} \frac{\partial \mathcal{U}_0(X_n^*)}{\partial X_n^*}. \tag{A.61}$$

Using the fact that $\mathcal{U}_0(X_n^*)$ takes strictly negative values when $\gamma > 1$, we obtain

$$\frac{\partial \mathcal{U}_0(X_n^*)}{\partial X_n^*} = \frac{(\gamma-1)(n-1)}{X_n^*}\mathcal{U}_0(X_n^*) < 0. \tag{A.62}$$

Consider an economy with $n+1$ agents. For any agent i, a marginal increase in x_i changes $\mathcal{U}_0^i(x_1, ..., x_{n+1})$ by (this is the equivalent of (26), written here for $n+1$ agents):

$$\frac{\partial \mathcal{U}_0^i(x_1, ..., x_{n+1})}{\partial x_i} = \frac{(1-\gamma)(X_{n+1}-x_i)}{x_i X_{n+1}}\theta_i^{1-\gamma}\mathcal{U}_0(X_{n+1}) + \theta_i^{1-\gamma}\frac{\partial \mathcal{U}_0(X_{n+1})}{\partial X_{n+1}}. \tag{A.63}$$

Assume now that the $n+1$ agents experiment at the level X_n^*, that is, $x_j = X_n^*/(n+1)$ for all j. Replacing this into (A.63) and using (A.61) to replace $\partial \mathcal{U}_0(X_n^*)/\partial X_n^*$ yields

$$\frac{\partial \mathcal{U}_0^i(x_1, ..., x_{n+1})}{\partial x_i} = \frac{(1-\gamma)n}{X_n^*}\theta_i^{1-\gamma}\mathcal{U}_0(X_n^*) - \frac{(1-\gamma)(n-1)}{X_n^*}\theta_i^{1-\gamma}\mathcal{U}_0(X_n^*) \tag{A.64}$$

$$= \frac{1-\gamma}{X_n^*}\theta_i^{1-\gamma}\mathcal{U}_0(X_n^*) > 0. \tag{A.65}$$

Getting back to (32),

$$\mathcal{U}_0^i(x_1, ..., x_{n+1}) = \left(\frac{x_i}{\sum_{j=1}^{n+1} x_j}\right)^{1-\gamma}\mathcal{U}_0(X_n^*), \tag{A.66}$$

in the economy with $n+1$ agents in which we start from the aggregate equilibrium $X_{n+1} = X_n^*$, a marginal change in x_i decreases the total pie $\mathcal{U}_0(X_{n+1})$ according to (A.62), but also increases the consumption share $x_i/\sum_{j=1}^{n+1} x_j$. This latter effect dominates such that (A.65) is satisfied. Thus, every agent has an incentive to increase her own experimentation from the level $X_n^*/(n+1)$ so that in equilibrium the aggregate level of experimentation increases in n. □

A.8 Proofs of Propositions 5 and 6

Propositions 5 and 6 follow from standard asset pricing theory. In this setting, agents have common knowledge about the underlying parameters, have a common prior about β, and the aggregate consumption stream δ is publicly observable. Thus, agents have homogeneous beliefs.

The following Lemma characterizes the dynamic process for the stochastic discount factor in an economy in which the aggregate level of experimentation is X_n^*.

Lemma A.1 *The stochastic discount factor, defined as $\xi_t \equiv e^{-\rho t}(\delta_t/\delta_0)^{-\gamma}$, follows*

$$\begin{aligned}\frac{d\xi_t}{\xi_t} = &-\left[\rho + \gamma\left(\bar{f} + \widehat{\beta}_t X_n^* - \frac{c}{2}(X_n^*)^2\right) - \frac{\gamma(\gamma+1)}{2}\sigma^2(1 + k^2(X_n^*)^2 + 2k\varphi X_n^*)\right]dt \\ &- \gamma\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}\,d\widehat{W}_t.\end{aligned} \tag{A.67}$$

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

$$r_t^f = \rho + \gamma\left(\bar{f} + \widehat{\beta}_t X_n^* - \frac{c}{2}(X_n^*)^2\right) - \frac{\gamma(\gamma+1)}{2}\sigma^2(1 + k^2(X_n^*)^2 + 2k\varphi X_n^*), \tag{A.68}$$

$$\zeta(X_n^*) = \gamma\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}. \tag{A.69}$$

The proof follows standard results in asset pricing (Duffie 2010, Dumas and Luciano 2017, Ch. 12, Munk 2013, Ch. 8). Assuming time-additive expected utility, the stochastic discount factor is defined from the optimal consumption plan of the aggregate consumer as

$$\xi_t = e^{-\rho t} \frac{u'(c_t)}{u'(c_0)}. \tag{A.70}$$

In our case agents consume fixed shares of the aggregate consumption and observe the economy under the same filtration. Thus, they all share the same stochastic discount factor. Given the CRRA assumption, the dynamics of ξ can then be expressed as in (A.67) by means of Itô's Lemma. The continuously compounded risk-free rate is the negative of the drift of the stochastic discount factor, whereas the market price of risk is the negative of the diffusion of the stochastic discount factor. This yields (A.68)-(A.69). □

Proof of Proposition 5 To prove Proposition 5, consider any agent i. In an economy with n agents, agent i consumes a share θ_i of the aggregate consumption and thus her value function written at time t and under an optimal policy follows from Proposition 2:

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t, \theta_i, X_n^*) = \frac{e^{-\rho t}(\theta_i \delta_t)^{1-\gamma}}{1-\gamma} \mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*), \tag{A.71}$$

where $\mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)$, defined in (16) and written here at time t, is the price-dividend ratio:

$$\mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*) \equiv \int_t^T \exp\left[\kappa(X_n^*, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2}(X_n^*)^2 \nu_t (s-t)^2\right] ds. \tag{A.72}$$

Denote by $\mathcal{W}_t^i$ the wealth of agent i at time t. Since the risk-free asset is in zero-net supply, $\mathcal{W}_t^i = \theta_i \delta_t \mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)$ and $J(\mathcal{W}_t^i, \widehat{\beta}_t, \nu_t, t, X_n^*)$ can further be written:

$$J(\mathcal{W}_t^i, \widehat{\beta}_t, \nu_t, t, X_n^*) = \frac{e^{-\rho t}(\mathcal{W}_t^i)^{1-\gamma}}{1-\gamma} \mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)^{\gamma}. \tag{A.73}$$

From the analysis in Merton (1973) and Brennan (1998), agent i's optimal portfolio demand can then be written as (define $\pi_t(X_n^*)$ as the risk premium and $\sigma_{P,t}(X_n^*)$ as the stock diffusion):

$$\phi_t = \frac{J_{\mathcal{W}^i}}{-J_{\mathcal{W}^i\mathcal{W}^i}\mathcal{W}_t^i} \frac{\pi_t(X_n^*)}{\sigma_{P,t}(X_n^*)^2} + \frac{J_{\mathcal{W}^i\widehat{\beta}}}{-J_{\mathcal{W}^i\mathcal{W}^i}\mathcal{W}_t^i} \frac{\mathrm{Cov}_t(d\widehat{\beta}_t, dP_t/P_t)}{\mathrm{Var}_t(dP_t/P_t)}. \tag{A.74}$$

To compute $\mathrm{Cov}_t(d\widehat{\beta}_t, dP_t/P_t)/\mathrm{Var}_t(dP_t/P_t)$, we need to characterize the equilibrium price of the risky asset claim to aggregate consumption at time t, P_t, which equals $P_t = \delta_t \mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)$. Compute the following partial derivatives of the price-dividend ratio:

$$\mathcal{P}_t \equiv \frac{\partial \mathcal{P}}{\partial t} = -1 - \kappa(X_n^*, \widehat{\beta}_t)\mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*) - (1-\gamma)^2 (X_n^*)^2 \nu_t G(\widehat{\beta}_t, \nu_t, t, X_n^*), \tag{A.75}$$

$$\mathcal{P}_{\widehat{\beta}} \equiv \frac{\partial \mathcal{P}}{\partial \widehat{\beta}_t} = (1-\gamma) X_n^* G(\widehat{\beta}_t, \nu_t, t, X_n^*), \tag{A.76}$$

$$\mathcal{P}_{\widehat{\beta}\widehat{\beta}} \equiv \frac{\partial^2 \mathcal{P}}{\partial \widehat{\beta}_t^2} = (1-\gamma)^2 (X_n^*)^2 H(\widehat{\beta}_t, \nu_t, t, X_n^*), \tag{A.77}$$

$$\mathcal{P}_\nu \equiv \frac{\partial \mathcal{P}}{\partial \nu_t} = \frac{(1-\gamma)^2}{2}(X_n^*)^2 H(\widehat{\beta}_t, \nu_t, t, X_n^*), \tag{A.78}$$

where we define the functions

$$G(\widehat{\beta}_t, \nu_t, t, X_n^*) \equiv \int_t^T (s-t)\exp\left[\kappa(X_n^*, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2}(X_n^*)^2 \nu_t (s-t)^2\right] ds \tag{A.79}$$

$$H(\widehat{\beta}_t, \nu_t, t, X_n^*) \equiv \int_t^T (s-t)^2 \exp\left[\kappa(X_n^*, \widehat{\beta}_t)(s-t) + \frac{(1-\gamma)^2}{2}(X_n^*)^2 \nu_t (s-t)^2\right] ds. \tag{A.80}$$

We can then write

$$\frac{J_{\mathcal{W}^i}}{-J_{\mathcal{W}^i\mathcal{W}^i}\mathcal{W}_t^i} = \frac{1}{\gamma} \quad \text{and} \quad \frac{J_{\mathcal{W}^i\widehat{\beta}}}{-J_{\mathcal{W}^i\mathcal{W}^i}\mathcal{W}_t^i} = \frac{\mathcal{P}_{\widehat{\beta}}}{\mathcal{P}}. \tag{A.81}$$

Apply Itô's lemma to $P_t = \delta_t \mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)$,

$$dP_t = \delta_t \mathcal{P}\frac{d\delta_t}{\delta_t} + \delta_t \mathcal{P}_{\widehat{\beta}} d\widehat{\beta}_t + \delta_t \mathcal{P}_\nu d\nu_t + \delta_t \mathcal{P}_t dt + \frac{1}{2}\left[\delta_t \mathcal{P}_{\widehat{\beta}\widehat{\beta}}(d\widehat{\beta}_t)^2 + 2\mathcal{P}_{\widehat{\beta}}(d\delta_t)(d\widehat{\beta}_t)\right], \tag{A.82}$$

to obtain the dynamics of the stock price:

$$\frac{dP_t}{P_t} = \mu_{P,t}dt + \sigma_{P,t}(X_n^*)d\widehat{W}_t, \tag{A.83}$$

with

$$\mu_{P,t} \equiv \bar{f} + \widehat{\beta}_t X_n^* - \frac{c(X_n^*)^2}{2} - \kappa(X_n^*, \widehat{\beta}_t) - \frac{1}{\mathcal{P}(\widehat{\beta}_t, \nu_t, t)} + \gamma(1-\gamma)(X_n^*)^2 \nu_t \frac{G(\widehat{\beta}_t, \nu_t, t, X_n^*)}{\mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)}, \tag{A.84}$$

$$\sigma_{P,t}(X_n^*) \equiv \sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*} + (1-\gamma)\frac{(X_n^*)^2\nu_t}{\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}}\frac{G(\widehat{\beta}_t, \nu_t, t, X_n^*)}{\mathcal{P}(\widehat{\beta}_t, \nu_t, t, X_n^*)}, \tag{A.85}$$

where we recognize the wealth duration at time t, $\frac{G(\widehat{\beta}_t,\nu_t,t,X_n^*)}{\mathcal{P}(\widehat{\beta}_t,\nu_t,t,X_n^*)} = \mathcal{D}_t(X_n^*)$. Getting back to (A.74):

$$\phi_t = \frac{1}{\gamma}\frac{\pi_t(X_n^*)}{\sigma_{P,t}(X_n^*)^2} + \frac{\mathcal{P}_{\widehat{\beta}}}{\mathcal{P}}\frac{X_n^*\nu_t}{\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}\sigma_{P,t}(X_n^*)} \tag{A.86}$$

$$= \frac{1}{\gamma}\frac{\pi_t(X_n^*)}{\sigma_{P,t}(X_n^*)^2} + (1-\gamma)\frac{(X_n^*)^2\nu_t}{\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}\sigma_{P,t}(X_n^*)}\mathcal{D}_t(X_n^*). \quad \square \tag{A.87}$$

Proof of Proposition 6 The equilibrium diffusion of stock returns has been determined in (A.85). To obtain the risk premium as in (46), multiply the market price of risk from (A.69), $\zeta(X_n^*) = \gamma\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}$, with $\sigma_{P,t}(X_n^*)$:

$$\pi_t(X_n^*) = \gamma\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}\left[\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*} + (1-\gamma)\frac{(X_n^*)^2\nu_t}{\sigma\sqrt{1 + k^2(X_n^*)^2 + 2k\varphi X_n^*}}\mathcal{D}_t(X_n^*)\right] \tag{A.88}$$

$$= \gamma\sigma^2(1 + k^2(X_n^*)^2 + 2k\varphi X_n^*) - \gamma(\gamma-1)(X_n^*)^2\nu_t\mathcal{D}_t(X_n^*). \quad \square \tag{A.89}$$

A.9 Discretization of the Continuous-Time Setup

This appendix derives a discretization of our continuous-time setup (see Section 3.2).

$$\delta_{t+\Delta} = \delta_t e^{\left[\bar{f}+\widehat{\beta}_t X_n^* - \frac{c}{2}(X_n^*)^2 - \frac{1}{2}\sigma^2(1+k^2(X_n^*)^2+2k\varphi X_n^*)\right]\Delta + \sqrt{\Delta}\sigma\sqrt{1+k^2(X_n^*)^2+2k\varphi X_n^*} z_{t+\Delta}}, \tag{A.90}$$

$$\widehat{\beta}_{t+\Delta} = \widehat{\beta}_t + \sqrt{\Delta}\frac{X_n^*}{\sigma\sqrt{1+k^2(X_n^*)^2+2k\varphi X_n^*}}\nu_t z_{t+\Delta}, \tag{A.91}$$

$$\nu_{t+\Delta} = \nu_t - \frac{(X_n^*)^2\nu_t^2\Delta}{\sigma^2(1+k^2(X_n^*)^2+2k\varphi X_n^*)}, \text{ where } z_{t+\Delta} \sim i.i.d.\ N(0,1) \text{ and } \Delta = \text{time step in years.} \tag{A.92}$$

A.10 Appendix for Section 4.1 (Infinite Horizon)

Proof of Proposition 7 In the setup with obsolescence and infinite horizon, the unobservable variable, β_t, is now time-varying according to (48). Hence, in Theorem A.1, $a_0 = 0$, $a_1 = -\lambda$, and $b_1 = b_2 = 0$. The observable process is δ_t. Applying Itô's lemma on $\ln \delta_t$ yields

$$d\ln\delta_t = \left[\bar{f} + \beta_t X - \frac{c}{2}X^2 - \frac{1}{2}\sigma^2(1+kX)^2\right]dt + \sigma(1+kX)dW_t, \tag{A.93}$$

and thus

$$A_0 = \bar{f} - \frac{c}{2}X^2 - \frac{1}{2}\sigma^2(1+kX)^2, \quad A_1 = X, \quad B = \sigma(1+kX). \tag{A.94}$$

Because a_0, a_1, b_1, b_2, A_0, A_1, and B are constants, all the conditions of Theorem A.1 are satisfied. Direct application of (A.5) then yields

$$d\widehat{\beta}_t = -\lambda\widehat{\beta}_t + \frac{\nu_t X}{\sigma(1+kX)}d\widehat{W}_t, \tag{A.95}$$

where $\widehat{W}_t$ is a standard Brownian motion with respect to the agent's filtration $\{\mathcal{F}_t^\delta\}$.

Eq. (A.6) further implies

$$\frac{d\nu_t}{dt} = -2\lambda\nu_t - \frac{X^2\nu_t^2}{\sigma^2(1+kX)^2}, \tag{A.96}$$

whose solution is Eq. (53) in the text. □

Proof of Proposition 8 We follow the same steps as in the proof of Proposition 2. The expected lifetime utility is defined in (A.13). We now compute the expectation $\mathbb{E}_0[\delta_t^{1-\gamma}]$. Write

$$\begin{bmatrix} d\ln\delta_t \\ d\widehat{\beta}_t \end{bmatrix} = \left(\begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1+kX)^2 - \frac{c}{2}X^2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & X \\ 0 & -\lambda \end{bmatrix}\begin{bmatrix} \ln\delta_t \\ \widehat{\beta}_t \end{bmatrix}\right)dt + \begin{bmatrix} \sigma(1+kX) \\ A(t) \end{bmatrix}d\widehat{W}_t, \tag{A.97}$$

where $A(t)$ is a function of time that results from Proposition 7 and Eq. (53):

$$A(t) = \frac{2\lambda\nu_0 X\sigma(1+kX)}{(e^{2\lambda t}-1)\nu_0 X^2 t + 2\lambda e^{2\lambda t}\sigma^2(1+kX)^2}. \tag{A.98}$$

Define:

$$K_0 \equiv \begin{bmatrix} \bar{f} - \frac{1}{2}\sigma^2(1+kX)^2 - \frac{c}{2}X^2 \\ 0 \end{bmatrix}, \tag{A.99}$$

$$K_1 \equiv \begin{bmatrix} 0 & X \\ 0 & -\lambda \end{bmatrix}, \tag{A.100}$$

$$H_0(t) \equiv \begin{bmatrix} \sigma(1+kX) \\ A(t) \end{bmatrix} \begin{bmatrix} \sigma(1+kX) \\ A(t) \end{bmatrix}' = \begin{bmatrix} \sigma^2(1+kX)^2 & \sigma(1+kX)A(t) \\ \sigma(1+kX)A(t) & A(t)^2 \end{bmatrix}. \tag{A.101}$$

Let $t, s \in [0, T]$ such that $t \leq s \leq T$, and define $\tau = s - t$. In order to compute the expectation $\mathbb{E}_t[\delta_s^{1-\gamma}] = \mathbb{E}_t[e^{(1-\gamma)\ln \delta_s}]$, we conjecture an exponential-affine solution of the form

$$\mathbb{E}_t\left[\delta_s^{1-\gamma}\right] = e^{\alpha_0(\tau)+\alpha_1(\tau)\ln\delta_t+\alpha_2(\tau)\widehat{\beta}_t}, \tag{A.102}$$

for some coefficient functions $\alpha_j(\cdot)$, $j = 0, 1, 2$, which satisfy (A.21)-(A.22) with boundary conditions $\alpha_0(0) = 0$, $\alpha_1(0) = 1-\gamma$, and $\alpha_2(0) = 0$. This is a system of Riccati ordinary differential equations (Duffie et al., 2003). The solution of (A.21) is:

$$\alpha_1(\tau) = 1-\gamma \tag{A.103}$$

$$\alpha_2(\tau) = (1-\gamma)X\frac{1-e^{-\lambda\tau}}{\lambda}, \tag{A.104}$$

which can be now inserted in the remaining Riccati equation (A.22). After replacing $A(t)$ from (A.98) and t by $s - \tau$, then further using (13) to replace ν_0 with a function of ν_t, we obtain a solution for α_0 for any $t, s \in [0, T]$ with $s \geq t$ and $\tau = s - t$:

$$\alpha_0(\tau) = \frac{1}{2}(\gamma-1)\left[-2\bar{f} + cX^2 + \gamma\sigma^2(1+kX)^2\right]\tau + \frac{1}{2}X^2\nu_t(\gamma-1)^2\left(\frac{1-e^{-\lambda\tau}}{\lambda}\right)^2. \tag{A.105}$$

Written at time 0 and with $\tau = t$, the solution of the Riccati system (A.21)-(A.22) is

$$\alpha_0(t) = \frac{1}{2}(\gamma-1)\left[-2\bar{f} + cX^2 + \gamma\sigma^2(1+kX)^2\right]t + \frac{1}{2}X^2\nu_0(\gamma-1)^2\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2, \tag{A.106}$$

$$\alpha_1(t) = 1-\gamma, \tag{A.107}$$

$$\alpha_2(t) = (1-\gamma)X\frac{1-e^{-\lambda t}}{\lambda}, \tag{A.108}$$

and we notice that taking the limit $\lambda \to 0$ yields the baseline case solution (Proposition 2). After replacing this solution into the conjecture (A.102) and multiplication with $e^{-\rho t}$, we obtain:

$$e^{-\rho t}\mathbb{E}_0\left[\delta_t^{1-\gamma}\right] = \exp\left[(1-\gamma)\ln\delta_0 + \kappa(X, \widehat{\beta}_0, t) + \frac{(\gamma-1)^2X^2\nu_0}{2}\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2\right], \tag{A.109}$$

where $\kappa(X, \widehat{\beta}_0, t)$ is now defined as in (54):

$$\kappa(X, \widehat{\beta}_0, t) \equiv \left[(1-\gamma)\left(\bar{f} - \frac{c}{2}X^2 - \gamma\frac{\sigma^2(1+kX)^2}{2}\right) - \rho\right]t + (1-\gamma)\widehat{\beta}_0X\frac{1-e^{-\lambda t}}{\lambda}. \tag{A.110}$$

Replacing (A.109) into (A.13) yields Eq. (55) of Proposition 8:

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\int_0^{\infty} \exp\left[\kappa(X,\widehat{\beta}_0,t) + \frac{(\gamma-1)^2}{2}X^2\nu_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2\right]dt. \tag{A.111}$$

Part (b) of Proposition 8 follows from standard asset pricing theory, as in Proposition 2. The equilibrium price-dividend ratio equals

$$\mathcal{P}_0(X) \equiv \int_0^{\infty} \exp\left[\kappa(X,\widehat{\beta}_0,t) + \frac{(\gamma-1)^2}{2}X^2\nu_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2\right]dt, \tag{A.112}$$

and we also notice that

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X). \quad \square \tag{A.113}$$

Transversality condition We derive a necessary and sufficient condition for the expected lifetime utility in (A.111) and the price-dividend ratio in (A.112) to be bounded. When $\lambda > 0$, the term $(1-e^{-\lambda t})/\lambda$ equals $1/\lambda$ as $t \to \infty$. Given this, to obtain finite values for the utility and price-dividend ratio, the exponent in (A.111)-(A.112),

$$\kappa(X,\widehat{\beta}_0,t) + \frac{(\gamma-1)^2}{2}X^2\nu_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2, \tag{A.114}$$

must reach $-\infty$ as $t \to \infty$. Using the expression for $\kappa(X,\widehat{\beta}_0,t)$ in (A.110), we obtain

$$\lim_{t\to\infty} \kappa(X,\widehat{\beta}_0,t) + \frac{(\gamma-1)^2}{2}X^2\nu_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2 \tag{A.115}$$

$$= \left[(1-\gamma)\left(\bar{f} - \frac{c}{2}X^2 - \gamma\frac{\sigma^2(1+kX)^2}{2}\right) - \rho\right]\infty + \frac{(1-\gamma)\widehat{\beta}_0 X}{\lambda} + \frac{(\gamma-1)^2X^2\nu_0}{2\lambda^2}, \tag{A.116}$$

which reaches $-\infty$ only if (57) is satisfied:

$$(1-\gamma)\left(\bar{f} - \frac{c}{2}X^2 - \gamma\frac{\sigma^2(1+kX)^2}{2}\right) - \rho < 0. \tag{A.117}$$

For similar transversality conditions in the existing literature, see Brennan and Xia (2001, Theorem 1) and Dumas, Kurshev, and Uppal (2009, Lemma 6). $\square$

Proof of concavity of expected lifetime utility Write (A.111) as

$$\mathcal{U}_0(X) = \frac{\delta_0^{1-\gamma}}{1-\gamma}\int_0^{\infty} e^{y(\widehat{\beta}_0,\nu_0,X,t)}dt, \tag{A.118}$$

and consider the second partial derivative of the function $y(\widehat{\beta}_0,\nu_0,X,t)$ with respect to X,

$$\frac{\partial^2 y(\widehat{\beta}_0,\nu_0,X,t)}{\partial X^2} = (\gamma-1)(c+\gamma\sigma^2k^2)t + (\gamma-1)^2\nu_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right)^2 > 0. \tag{A.119}$$

This implies that $U_0(X)$ is concave in X, following the same reasoning as in Appendix A.4. It further implies that the price-dividend ratio is log-convex in X and thus $\partial \ln \mathcal{P}_0(X)/\partial X$ is increasing in X. One can show that $\partial \ln \mathcal{P}_0(X)/\partial X$ increases from $(1-\gamma)[\widehat{\beta}_0 \mathcal{D}_0^E(0) - \gamma k \sigma^2 \mathcal{D}_0(0)]$ to ∞ as X increases from 0 to ∞, an important result for the existence and uniqueness of an equilibrium level of experimentation with competition (Proposition 9). □

Proof of Proposition 9 Any agent i's lifetime expected utility, $\mathcal{U}_0^i$, can be written as

$$\mathcal{U}_0^i(x_1, ..., x_n) = \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{(\theta_i \delta_t)^{1-\gamma}}{1-\gamma} dt \right] = \theta_i^{1-\gamma} \mathcal{U}_0 \left(\sum_{j=1}^n x_j \right), \tag{A.120}$$

where $\mathcal{U}_0(\cdot)$ is the function defined and characterized in Proposition 8.

To find the Nash equilibrium, write the first-order condition for agent i's maximization problem as

$$0 = \frac{\partial \mathcal{U}_0^i(x_1, ..., x_n)}{\partial x_i} = \frac{(1-\gamma)(X_n - x_i)}{x_i X_n} \theta_i^{1-\gamma} \mathcal{U}_0(X_n) + \theta_i^{1-\gamma} \frac{\partial \mathcal{U}_0(X_n)}{\partial X_n}, \tag{A.121}$$

which, after dividing by $\theta_i^{1-\gamma} \mathcal{U}_0(X_n)$ and replacing $\mathcal{U}_0(X_n)$ by $\frac{\delta_0^{1-\gamma}}{1-\gamma} \mathcal{P}_0(X_n)$, yields

$$\frac{(\gamma - 1)(X_n - x_i)}{x_i X_n} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{A.122}$$

We rule out asymmetric equilibria using the same logic as in Section 2.2. In a symmetric equilibrium, X_n solves

$$\frac{(\gamma - 1)(n-1)}{X_n} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{A.123}$$

The equilibrium price-dividend ratio $\mathcal{P}_0(X)$ is log-convex in X. This implies that the right-hand side of (A.123) strictly increases in X_n. We have showed above that $\partial \ln \mathcal{P}_0(X_n)/\partial X_n$ takes values from $(1-\gamma)[\widehat{\beta}_0 \mathcal{D}_0^E(0) - \gamma k \sigma^2 \mathcal{D}_0(0)]$, which is finite, to ∞ as X increases from 0 to ∞. When $n \geq 2$, the left-hand side strictly decreases in X_n, taking values from ∞ to 0. Thus, any equilibrium that satisfies (A.123) is unique, and X_n solves

$$X_n^* = \frac{\widehat{\beta}_0 \mathcal{D}_0^E(X_n^*) - \gamma k \sigma^2 \mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*}}{(\gamma k^2 \sigma^2 + c) \mathcal{D}_0(X_n^*) + (\gamma - 1) \nu_0 \mathcal{C}_0^E(X_n^*)}, \tag{A.124}$$

where $\mathcal{D}_0^E(X)$, $\mathcal{C}_0^E(X)$, and $\mathcal{D}_0(X)$ are defined as

$$\mathcal{D}_0^E(X) = \frac{1}{\mathcal{P}_0(X)} \int_0^\infty \frac{1 - e^{-\lambda t}}{\lambda} \exp \left[\kappa(X, \widehat{\beta}_0, t) + \frac{(\gamma-1)^2}{2} X^2 \nu_0 \left(\frac{1 - e^{-\lambda t}}{\lambda} \right)^2 \right] dt, \tag{A.125}$$

$$\mathcal{C}_0^E(X) = \frac{1}{\mathcal{P}_0(X)} \int_0^\infty \left(\frac{1 - e^{-\lambda t}}{\lambda} \right)^2 \exp \left[\kappa(X, \widehat{\beta}_0, t) + \frac{(\gamma-1)^2}{2} X^2 \nu_0 \left(\frac{1 - e^{-\lambda t}}{\lambda} \right)^2 \right] dt. \tag{A.126}$$

$$\mathcal{D}_0(X) = \frac{1}{\mathcal{P}_0(X)} \int_0^\infty t \exp \left[\kappa(X, \widehat{\beta}_0, t) + \frac{(\gamma-1)^2}{2} X^2 \nu_0 \left(\frac{1 - e^{-\lambda t}}{\lambda} \right)^2 \right] dt. \quad \square \tag{A.127}$$

A.11 Alternative Setup with Learning from $\left\{\delta^S,\ \delta^E\right\}$

When the agent learns from $\left\{\delta^S,\ \delta^E\right\}$, she perfectly observes dW_t and therefore it is important in this case to have $\varphi \neq \pm 1$, otherwise g is fully revealed. When $X > 0$, we have:

$$\frac{d\delta_t^S}{\delta_t^S} = \left(\bar{f} - \frac{c}{2}X^2\right)dt + \sigma dW_t, \tag{A.128}$$

$$\frac{d\delta_t^E}{\delta_t^E} = \left(g - \frac{k^2\sigma^2}{2}X\right)dt + k\sigma\left(\varphi dW_t + \sqrt{1-\varphi^2}dW_t^E\right), \tag{A.129}$$

which we can write

$$d\ln\delta_t^S = \left(\bar{f} - \frac{1}{2}\sigma^2 - \frac{c}{2}X^2\right)dt + \sigma dW_t, \tag{A.130}$$

$$d\ln\delta_t^E = \left[g - \frac{1}{2}k^2\sigma^2(1+X)\right]dt + k\sigma\left(\varphi dW_t + \sqrt{1-\varphi^2}dW_t^E\right), \tag{A.131}$$

where g is unobservable. Learning about g is equivalent with learning about β (see (7)). Filtering theory (Liptser and Shiryaev, 2001, theorem 12.7, p. 36) then implies:

$$d\widehat{g}_t = \frac{\nu_t}{k\sigma\sqrt{1-\varphi^2}}d\widehat{W}_t^E, \quad \text{with } d\widehat{W}_t^E \equiv dW_t^E + \frac{g-\widehat{g}_t}{k\sigma\sqrt{1-\varphi^2}}dt, \tag{A.132}$$

$$d\nu_t = -\frac{\nu_t^2}{k^2\sigma^2(1-\varphi^2)}dt. \tag{A.133}$$

Thus, Proposition 1 becomes

$$\frac{d\delta_t}{\delta_t} = \left(\bar{f} + \widehat{\beta}_t X - \frac{c}{2}X^2\right)dt + \sigma(1+\varphi kX)dW_t + k\sigma X\sqrt{1-\varphi^2}d\widehat{W}_t^E, \tag{A.134}$$

$$d\widehat{\beta}_t = \frac{\nu_t}{k\sigma\sqrt{1-\varphi^2}}d\widehat{W}_t^E, \tag{A.135}$$

$$d\nu_t = -\frac{\nu_t^2}{k^2\sigma^2(1-\varphi^2)}dt. \tag{A.136}$$

From here, everything follows as in the paper, with the main difference that the dynamics of ν_t do not depend on the level of experimentation.

A.12 Appendix for Section 4.2 (Competition and Inequality)

Any agent i's lifetime expected utility, $\mathcal{U}_0^i$, can be written as

$$\mathcal{U}_0^i(x_1,...,x_n) = \mathbb{E}_0\left[\int_0^T e^{-\rho t}\frac{(\theta_i\delta_t)^{1-\gamma}}{1-\gamma}dt\right] = \theta_i^{1-\gamma}\mathcal{U}_0\left(\textstyle\sum_{j=1}^n x_j\right), \tag{A.137}$$

where $\mathcal{U}_0(\cdot)$ is the function defined and characterized in Proposition 2 and θ_i is now defined as in (60). The first-order condition for agent i's maximization problem is

$$0 = \frac{\partial \mathcal{U}_0^i(x_1, ..., x_n)}{\partial x_i} = \frac{(1-\gamma)(1+X_n-x_i)}{x_i(1+X_n)}\theta_i^{1-\gamma}\mathcal{U}_0(X_n) + \theta_i^{1-\gamma}\frac{\partial \mathcal{U}_0(X_n)}{\partial X_n}, \tag{A.138}$$

which, after dividing by $\theta_i^{1-\gamma}\mathcal{U}_0(X_n)$ and replacing $\mathcal{U}_0(X_n)$ by $\frac{\delta_0^{1-\gamma}}{1-\gamma}\mathcal{P}_0(X_n)$, yields

$$\frac{(\gamma-1)(1+X_n-x_i)}{x_i(1+X_n)} = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{A.139}$$

We rule out asymmetric equilibria using the same arguments as in Section 2.2. In a symmetric equilibrium, X_n solves

$$(\gamma-1)\left(\frac{n-1}{X_n} + \frac{1}{X_n(1+X_n)}\right) = \frac{\partial \ln \mathcal{P}_0(X_n)}{\partial X_n}. \tag{A.140}$$

From Section 2.2, we know that the equilibrium price-dividend ratio $\mathcal{P}_0(X)$ is log-convex in X and that $\partial \ln \mathcal{P}_0(X_n)/\partial X_n$ takes values from $(1-\gamma)(\widehat{\beta}_0 - \gamma k\sigma^2)\mathcal{D}_0(0)$, which is finite, to ∞. When $n \geq 2$, the left-hand side strictly decreases in X_n, taking values from ∞ to 0. Thus, any equilibrium that satisfies (A.140) is unique. Using (A.58), the equilibrium aggregate level of experimentation solves (62), which is an implicit equation in X_n:

$$X_n^* = \frac{(\widehat{\beta}_0 - \gamma k\sigma^2)\mathcal{D}_0(X_n^*) + \frac{n-1}{X_n^*} + \frac{1}{X_n^*(1+X_n^*)}}{(\gamma k^2\sigma^2 + c)\mathcal{D}_0(X_n^*) + (\gamma-1)\nu_0\mathcal{C}_0(X_n^*)}. \tag{A.141}$$

A.13 Appendix for Section 4.3 (Dynamic Experimentation)

A.13.1 Proof of Proposition 11

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

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + \widehat{\beta}_t X_t)dt + \sigma(1+kX_t)d\widehat{W}_t, \tag{A.142}$$

with

$$d\widehat{\beta}_t = \frac{X_t}{\sigma(1+kX_t)}\nu_t d\widehat{W}_t, \tag{A.143}$$

$$d\nu_t = -\frac{X_t^2}{\sigma^2(1+kX_t)^2}\nu_t^2 dt. \tag{A.144}$$

Now that the agent can choose the experimentation level X_t at any time t, the agent's expected lifetime utility of consumption J satisfies the partial differential equation

$$0 = \max_X \left[\mathcal{L}J(\delta_t, \widehat{\beta}_t, \nu_t, t) + e^{-\rho t}\frac{c_t^{1-\gamma}}{1-\gamma}\right], \quad \text{where} \quad \mathcal{L}J = \frac{\mathbb{E}[dJ]}{dt}, \tag{A.145}$$

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

In equilibrium consumption equals total output and therefore

$$0 = \max_X \frac{e^{-\rho t}\delta_t^{1-\gamma}}{1-\gamma} + J_t + \delta_t(\bar{f} + \widehat{\beta}_t X)J_\delta + \frac{X^2\nu_t^2}{2\sigma^2(1+kX)^2}(J_{\beta\beta} - 2J_\nu) + \frac{1}{2}\delta_t^2\sigma^2(1+kX)^2 J_{\delta\delta} + \delta_t\nu_t X J_{\delta\beta}. \tag{A.146}$$

With CRRA utility, we make the usual conjecture

$$J(\delta_t, \widehat{\beta}_t, \nu_t, t) = e^{-\rho t}\frac{\rho_t^{1-\gamma}}{1-\gamma}\mathcal{P}(\widehat{\beta}_t, \nu_t, t), \tag{A.147}$$

and thus the partial differential equation (A.146) becomes

$$0 = \max_X \mathcal{P}_t + \kappa(X, \widehat{\beta}_t)\mathcal{P} + \frac{X^2\nu_t^2}{2\sigma^2(1+kX)^2}(\mathcal{P}_{\beta\beta} - 2\mathcal{P}_\nu) + (1-\gamma)X\nu_t\mathcal{P}_\beta + 1, \tag{A.148}$$

with boundary condition $\mathcal{P}(\widehat{\beta}_T, \nu_T, T) = 0$ and $\kappa(X, \widehat{\beta}_t)$ defined as in (14), with $\varphi = 1$ and $c = 0$. The first order condition for X is

$$\begin{aligned} 0 &= \kappa_X\mathcal{P} + (1-\gamma)\nu_t\mathcal{P}_\beta + \frac{X\nu_t^2}{(1+kX)^3\sigma^2}(\mathcal{P}_{\beta\beta} - 2\mathcal{P}_\nu) && \text{(A.149)} \\ &= (\gamma - 1)\left(k(1+kX)\gamma\sigma^2 - \widehat{\beta}_t\right)\mathcal{P} + (1-\gamma)\nu_t\mathcal{P}_\beta + \frac{X\nu_t^2}{(1+kX)^3\sigma^2}(\mathcal{P}_{\beta\beta} - 2\mathcal{P}_\nu) && \text{(A.150)} \end{aligned}$$

This is a quartic equation in X, which we solve numerically for our illustrations in Section 4.3. Re-arranging this equation yields Equation (64) in the text. □

A.13.2 Asset Prices with Dynamic Experimentation

Proposition A.12 *In an economy with dynamic experimentation, the risk-free rate and the market price of risk are given by*

$$\begin{aligned} r_t^f &= \rho + \gamma(\bar{f} + X_t^*\widehat{\beta}_t) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kX_t^*)^2, && \text{(A.151)} \\ \zeta(X_t^*) &= \gamma\sigma(1+kX_t^*), && \text{(A.152)} \end{aligned}$$

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

$$\begin{aligned} \pi_t(X_t^*) &= \gamma\sigma^2(1+kX_t^*)^2\left(1 + \frac{X_t^*\nu_t}{\sigma^2(1+kX_t^*)^2}\frac{\mathcal{P}_\beta}{\mathcal{P}}\right), && \text{(A.153)} \\ \sigma_{P,t}(X_t^*) &= \sigma(1+kX_t^*)\left(1 + \frac{X_t^*\nu_t}{\sigma^2(1+kX_t^*)^2}\frac{\mathcal{P}_\beta}{\mathcal{P}}\right) && \text{(A.154)} \end{aligned}$$

Proof The stochastic discount factor follows

$$\frac{d\xi_t}{\xi_t} = -\left(\rho + \gamma(\bar{f} + X_t^*\widehat{\beta}_t) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1+kX_t^*)^2\right)dt - \gamma\sigma(1+kX_t^*)d\widehat{W}_t^\delta, \tag{A.155}$$

which yields the risk-free rate and the market price of risk from (A.151)-(A.152).

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

$$\frac{dP_t}{P_t} = \left(\bar{f} + X_t^* \widehat{\beta}_t - \kappa(X_t^*, \widehat{\beta}_t) - \frac{1}{\mathcal{P}} + \gamma X_t^* \nu_t \frac{\mathcal{P}_\beta}{\mathcal{P}}\right) dt + \sigma(1 + kX_t^*) \left(1 + \frac{X_t^* \nu_t}{\sigma^2(1 + kX_t^*)^2} \frac{\mathcal{P}_\beta}{\mathcal{P}}\right) d\widehat{W}_t, \tag{A.156}$$

from which we obtain the diffusion of stock returns. The risk premium is then given by

$$\pi_t(X_t^*) = \gamma\sigma^2(1 + kX_t^*)^2 \left(1 + \frac{X_t^* \nu_t}{\sigma^2(1 + kX_t^*)^2} \frac{\mathcal{P}_\beta}{\mathcal{P}}\right). \quad \square \tag{A.157}$$

A.14 Appendix for Section 4.4 (Recursive Preferences)

In this appendix, we consider an alternative setup in which agents derive utility from lifetime consumption and have stochastic differential utility (Epstein and Zin, 1989) with subjective discount rate β, relative risk aversion γ, and elasticity of intertemporal substitution ψ. The indirect utility function of any agent i is

$$J_{i,t} = \mathbb{E}_t \left[\int_t^\infty h(C_{i,s}, J_{i,s}) ds\right], \tag{A.158}$$

where the aggregator h is defined as in Duffie and Epstein (1992):

$$h(C_i, J_i) = \frac{\beta}{1 - 1/\psi} \left(\frac{C_i^{1-1/\psi}}{[(1-\gamma)J_i]^{1/\phi - 1}} - (1-\gamma)J_i\right), \quad \text{with } \phi \equiv \frac{1-\gamma}{1 - 1/\psi}. \tag{A.159}$$

All agents share the same preference parameters. We focus on the empirically relevant case when the coefficient of risk aversion is higher than one. The reciprocal of the elasticity of intertemporal substitution, $1/\psi$, represents agents' *aversion to intertemporal substitution* (Restoy and Weil, 2010). If $1/\psi = 0$, agents are indifferent to intertemporal substitution. The coefficient ϕ measures the departure from the time-additive isoelastic framework: when $\phi = 1$ (i.e., when $1/\psi = \gamma$), the preferences in (A.159) reduce to the standard CRRA utility representation. When $\psi > 1/\gamma$, agents prefer early resolution of uncertainty.

A difference with our baseline setup, in addition to the different utility specification, is that we consider here an infinite horizon problem. As we will show below, this will allow us to simplify the setup by eliminating one state variable.

The dynamics of the aggregate output stream in the economy is (for simplicity of exposition, we abstract away from opportunity costs):

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + X\mu_t)dt + \sigma dW_t. \tag{A.160}$$

where μ is unknown and plays the same role in β in (6). Because this is an infinite horizon economy, a constant μ would eventually be learned in finite time and technological uncertainty would be zero. To avoid this, we assume that μ_t has dynamics

$$d\mu_t = \lambda(\bar{\mu} - \mu_t)dt + \sigma_\mu dW_t^\mu, \tag{A.161}$$

with known parameters λ, $\bar{\mu}$, and σ_μ. Finally, we further simplify the setup to focus only on technological uncertainty and assume that disruption risk is zero ($k = 0$).

As in Section 2, we define agents' information filtration as $\{\mathcal{F}_t^\delta\}$, where $\mathcal{F}_t^\delta = \sigma(\delta_u : u \leq t)$. From agents' viewpoint, this partially observed economy is equivalent to a perfectly observed economy with

$$\frac{d\delta_t}{\delta_t} = (\bar{f} + X\widehat{\mu}_t)dt + \sigma d\widehat{W}_t, \tag{A.162}$$

$$d\widehat{\mu}_t = \lambda(\bar{\mu} - \widehat{\mu}_t)dt + \widehat{\sigma}_\mu d\widehat{W}_t, \tag{A.163}$$

where $\widehat{\sigma}_\mu \equiv X\bar{\nu}/\sigma$ and $d\widehat{W}_t \equiv dW_t + \frac{X(\mu - \widehat{\mu}_t)}{\sigma}dt$ represents the "surprise" component of the change in total output. ($\widehat{W}_t$ is a standard Brownian motion with respect to agents' filtration $\{\mathcal{F}_t^\delta\}$.) Using standard filtering results (Theorem A.1 in Appendix A.2), the posterior variance of the estimated expected growth of the new technology (i.e., the *technological uncertainty* ν) converges to a constant $\bar{\nu}$ that depends on the initial level of experimentation:

$$\bar{\nu} = \frac{\sigma\left(\sqrt{\lambda^2\sigma^2 + X^2\sigma_\mu^2} - \lambda\sigma\right)}{X^2} \geq 0. \tag{A.164}$$

As in the baseline setup, the representative agent's problem is to choose a level of experimentation that balances between expected growth gains and the imposed disturbance on future consumption growth. Solving for the optimal level of experimentation in equilibrium involves writing the Hamilton-Jacobi-Bellman (HJB) equation,

$$\max_{C,X}\{h(C,J) + \mathcal{L}J\} = 0, \quad \text{where} \quad \mathcal{L}J = \frac{\mathbb{E}[dJ]}{dt}. \tag{A.165}$$

We guess the following value function (Benzoni et al., 2011):

$$J(\delta, \widehat{\mu}; X) = \frac{\delta^{1-\gamma}}{1-\gamma}\left[\beta\mathcal{P}(\widehat{\mu}; X)\right]^\phi, \tag{A.166}$$

where $\mathcal{P}(\widehat{\mu}; X)$ is the price-dividend ratio. The benefit of assuming a steady-state value for ν (see, e.g., Dumas et al., 2009) is that the price-dividend ratio depends on one state variable only, $\widehat{\mu}$. This simplifies the approximation proposed below and allows us to show that our main results continue to hold. The drawback is that this setup has no dynamic implications related to time-variability in ν, such as Figure 3 in the baseline model.

Define the log price-dividend ratio:

$$I(\widehat{\mu}; X) \equiv \log \mathcal{P}(\widehat{\mu}; X). \tag{A.167}$$

Substituting (A.166) into the HJB equation (A.165) and imposing the market clearing condition $C = \delta$ yields the following ordinary differential equation for the log price-dividend ratio:

$$0 = \frac{\gamma - 1}{\phi}\left[-(\bar{f} + X\widehat{\mu}) + \frac{\gamma\sigma^2}{2}\right] - \beta + e^{-I} + \left[\lambda(\bar{\mu} - \widehat{\mu}) - (\gamma - 1)\sigma\widehat{\sigma}_\mu\right] I_{\widehat{\mu}} + \frac{\widehat{\sigma}_\mu^2}{2}\left(I_{\widehat{\mu}\widehat{\mu}} + \phi I_{\widehat{\mu}}^2\right). \tag{A.168}$$

To obtain an approximate solution for $\mathcal{P}(\widehat{\mu}; X)$, we use the customary log-linear approximation of the price-dividend ratio (Restoy and Weil, 2010; Bansal and Yaron, 2004; Beeler and Campbell, 2012; Benzoni et al., 2011):

$$\mathcal{P}(\widehat{\mu}; X) = e^{A(X)+B\times X\widehat{\mu}}. \tag{A.169}$$

We have checked the accuracy of this approximation using an alternative solution method with a high number of Chebyshev polynomials. The approximation (A.169) performed very well. Log-linear approximations of long-run risk models do not always guarantee reliable results (Pohl, Schmedders, and Wilms, 2018), but in our model with one state variable the approximation remains very accurate.

We conjecture the following form for the coefficients $A(X)$ and B:

$$\mathcal{P}(\widehat{\mu}; X) = \exp\left[b(1-1/\psi)\left(c_0 + c_1 X - \frac{c_2}{2}X^2 + X\widehat{\mu}\right)\right]. \tag{A.170}$$

This form results from a second-order approximation of $A(X)$ around $X = 0$. When $X = 0$ the price-dividend ratio is constant and the approximation is exact. A further justification for the second-order approximation of $A(X)$ results from the CRRA case: dividend strips (assets that pay the aggregate consumption only at time $s > t$) have exactly the form in Eq. (A.170).[26]

Lemma A.2 *If $\gamma > 1$, $\psi > 1/\gamma$, and $\psi \neq 1$, then $b > 0$.*

Lemma A.2 follows from Restoy and Weil (2010, Eqs. (3.1) and (4.4)). As long as the risk aversion is higher than one and agents have a preference for early resolution of uncertainty, the sign of the coefficient b is unambiguously positive. Eq. (A.170) further implies that the price-dividend ratio is convex in $\widehat{\mu}$, meaning that technological uncertainty induces over-valuation of the risky asset, as in our main setup (see Eq. (42) and its discussion).

The conjectured form (A.170) for the price-dividend ratio proves particularly convenient for the identification of the equilibrium level of experimentation, both in the representative-agent economy and in the economy with n agents.

Proposition A.13 (*a*) *At $t = 0$, the representative agent chooses an optimal level of experimentation that maximizes total welfare:*

$$X^* = \max\left[\frac{1}{c_2}(c_1 + \widehat{\mu}_0), 0\right]. \tag{A.171}$$

(*b*) *At $t = 0$, the aggregate level of experimentation in an economy with n competitors whose shares of consumption are given by* (23) *is:*

$$X_n^* = \max\left[\frac{1}{c_2}\left(c_1 + \widehat{\mu}_0 + \frac{1}{b}\frac{n-1}{X_n^*}\right), 0\right]. \tag{A.172}$$

Proof *The proof follows the same steps as in Proposition 4, Section 2.2.* □

[26]The second-order approximation of $A(X)$ is necessary because it pins down the optimal level of experimentation (see Proposition A.13). Without a second order term, the optimal level of experimentation is indeterminate (Devereux and Sutherland, 2011, make a similar argument in portfolio choice models).

Comparing (a) with Proposition 3 and (b) with Proposition 4, the same results hold: there exists a socially optimal level of experimentation that maximizes welfare; this level can only be reached in the representative-agent economy. Competition generates over-experimentation, and the aggregate level of experimentation increases with n. The economy with competition can again generate experimentation with a sub-optimal technology, as in Section 2.2.

One disadvantage of the log-linear approximation (A.170) is that the coefficients, b, c_1 and c_2, which dictate the equilibrium experimentation in Proposition A.13, do not clearly reveal the impact of technological uncertainty. Our numerical results suggest that the coefficient c_1 strongly depends on the parameters σ_μ and λ, which control the amount of technological uncertainty $\bar{\nu}$ (see Eq. (A.164)). If σ_μ is high and/or λ is low, then $c_1 + \widehat{\mu}_0$ can become negative and the active agent is better off not experimenting (c_2 is positive with our calibration).

Moving now to asset pricing, for any experimentation level X, the equilibrium risk premium in the economy is given by (see Restoy and Weil, 2010, Eq. (4.6), for a similar decomposition):

$$\pi_t(X) = \underbrace{\left(\gamma\sigma + (\gamma - 1)\frac{\phi - 1}{\phi} b\widehat{\sigma}_\mu X\right)}_{\text{Market price of risk}} \times \underbrace{\sigma\left(1 + \frac{\psi - 1}{\psi}\frac{\widehat{\sigma}_\mu}{\sigma} bX\right)}_{\text{Quantity of risk}}, \tag{A.173}$$

and the equilibrium stock return diffusion (the "quantity of risk") is

$$\sigma_{P,t}(X) = \sigma\left(1 + \frac{\psi - 1}{\psi}\frac{\widehat{\sigma}_\mu}{\sigma} bX\right). \tag{A.174}$$

Before analyzing these quantities, note that when $\gamma > 1$, ϕ can take the following values:

$$\begin{cases} \phi \in [0, 1), & \text{if} \quad \psi \in [0, 1/\gamma) \\ \phi \in [1, \infty), & \text{if} \quad \psi \in [1/\gamma, 1) \\ \phi \in (-\infty, 1 - \gamma), & \text{if} \quad \psi \in (1, \infty) \quad \text{(Bansal and Yaron, 2004)}. \end{cases} \tag{A.175}$$

Over each one of the three intervals above, ϕ is a strictly increasing function in ψ. The first case in (A.175) represents preference for *late* resolution of uncertainty. The next two cases represent preference for *early* resolution of uncertainty, with the last case being the calibration of the long-run risk model (Bansal and Yaron, 2004).

We will focus our discussion in a setting with $\gamma > 1$ and preference for early resolution of uncertainty.

1. **Case** $\psi \in [1/\gamma, 1) \;\Rightarrow\; \phi \in [1, \infty)$

   When the elasticity of intertemporal substitution is lower than one, the market price of risk increases with technological uncertainty. (In comparison, in the main model with CRRA utility, the market price of risk does not depend on technological uncertainty—see Footnote 20.) The quantity of risk, however, decreases with technological uncertainty. When competition is intense, high aggregate experimentation may generate a negative risk premium, as in our main model (Figure 5).

2. **Case** $\psi \in (1, \infty) \;\Rightarrow\; \phi \in [-\infty, 1 - \gamma)$

   When the elasticity of intertemporal substitution is higher than one, as in the typical calibration of the long-run risk model (Bansal and Yaron, 2004), both the market price of risk and the quantity of risk strictly increase with technological uncertainty, and this is further exacerbated for strong levels of experimentation. In this case, although the risky asset remains over-valued due to technological uncertainty, the risk premium in the economy always increases in experimentation.

These two cases highlight the important role played by the elasticity of intertemporal substitution ψ for the behavior of the risk premium in this alternative setup with stochastic differential utility. Empirical studies disagree about reasonable values for ψ. Some studies find ψ greater than one (Vissing-Jørgensen and Attanasio, 2003), other studies find ψ smaller than one (Campbell, 1999; Vissing-Jørgensen, 2002). While our main results related to over-experimentation and over-valuation of the asset hold in both cases, the relationship between the risk premium and technological uncertainty depends on the value of ψ.

To summarize, the main results in our paper are robust to an alternative utility function. But the analysis is less transparent due to the approximation of the price-dividend ratio and also because technological uncertainty is now constant, which prevents us to make statements about the dynamics of asset prices in the aftermath of experimentation.

A.15 Analysis of the Case $\gamma < 1$

We have analyzed our main results under the assumption $\gamma > 1$. In this appendix, we show that most of the results hold as long as $\gamma \neq 1$. When $\gamma < 1$, Propositions 2 and 3 remain valid, but in this case agents prefer more uncertainty (see footnote 13), and thus uncertainty *amplifies* experimentation. Competition continues to generate over-experimentation when $\gamma < 1$, and the level of experimentation increases in n (as in Proposition 4), but proof of uniqueness is no longer possible (when $\gamma < 1$, the agents' expected lifetime utilities are not globally concave). The results presented here are qualitatively similar to that of the case $\gamma > 1$, stronger in the case $\gamma < 1$, because agents' preference for uncertainty strengthens their incentives to experiment. When $\gamma < 1$, technological uncertainty continues to imply asset over-valuation and a downward average path for prices in the aftermath of experimentation, as in Figure 3, and negative return predictability still obtains, as in Table 1. Proposition 6, however, yields a different result: when $\gamma < 1$, the risk premium and the stock market volatility are now strictly increasing with the level of experimentation.

Proposition 2 holds for any positive value of γ. The optimal level of experimentation, X^*, is determined as in Proposition 3, but when $\gamma < 1$ the last term in the denominator of X^* is negative, meaning that the agent likes uncertainty. This further amplifies our results and increase the optimal level of experimentation relative to the case $\gamma > 1$.

The representative agent's lifetime expected utility in (15) is positive when $\gamma < 1$ and it is not strictly concave over its entire domain. The same holds for the expected lifetime utility of any active agent i in the setup with competition in (24). Thus, uniqueness is not guaranteed. Nevertheless, the aggregate level of experimentation is still determined as in Proposition 4. We illustrate this in panels (a) and (b) of Figure 10, which use $\gamma < 1$ in the calibration and replicate Figure 1 in the paper. The two panels confirm our main result that experimentation increases with the number of competitors, as predicted by Proposition 4.

Panel (c) of Figure 10 depicts the risk premium as a function of aggregate experimentation when $\gamma < 1$. In this case, the risk premium is strictly increasing in experimentation: agents now hold positive hedging demands and according to (44) the risk premium increases in $\mathcal{H}_t(X_n^*, \nu_t)$. (See also Proposition 6, where the last term in (46) now increases with experimentation.)

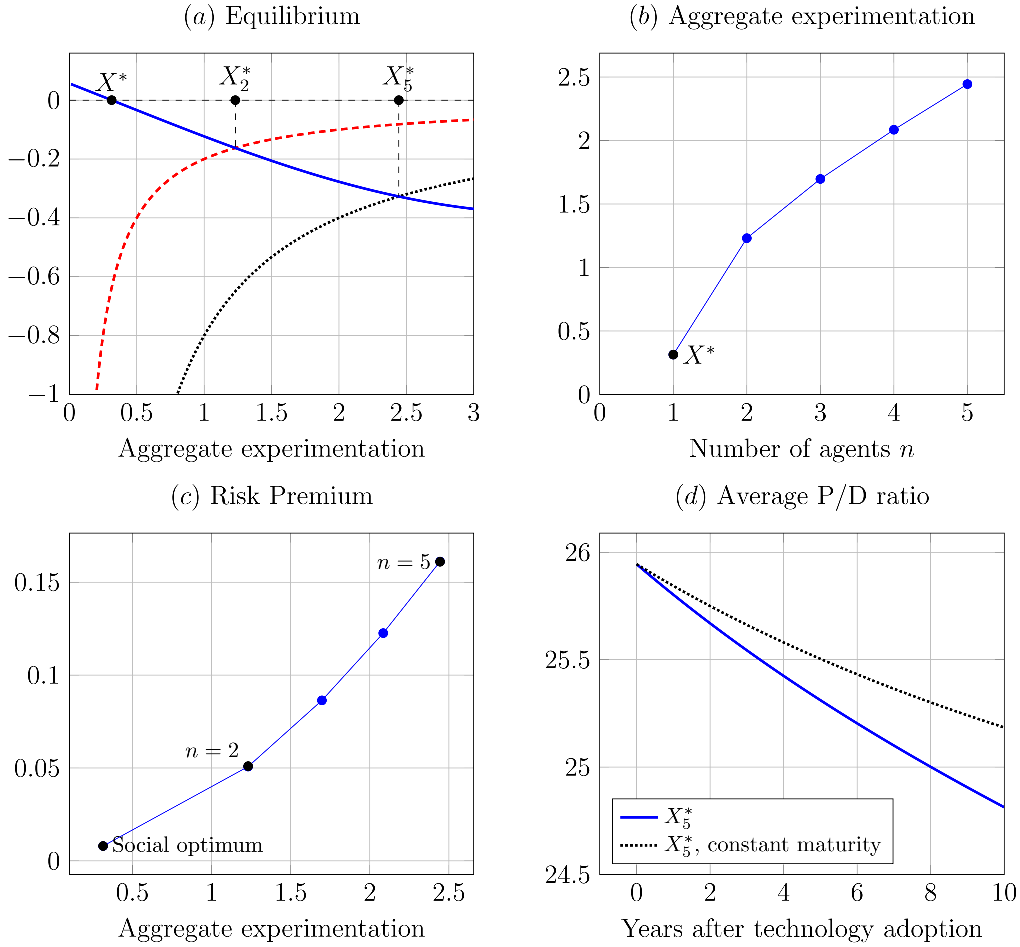

Figure 10: **Equilibrium Experimentation with Competition, the case $\gamma < 1$.** Panel (a) depicts the solution of (29), where the solid line represents the right-hand side and the dashed (dotted) line represent the left-hand side for $n = 2$ ($n = 5$). Panel (b) depicts the aggregate level of experimentation as a function of the number of agents in the economy. Panel (c) shows the risk premium as a function of the aggregate experimentation level in the economy, similar with Figure 5. Panel (d) is similar with panel (b) in Figure 3. The calibration used is: $\gamma = 0.8$, $\bar{f} = 0.03$, $\widehat{\beta}_0 = 0.015$, $\nu_0 = 0.03^2$, $\sigma = 0.05$, $k = 3$, $\rho = 0.03$, $T = 100$, $c = 0.02$, and $\delta_0 = 1$.

The over-valuation due to technological uncertainty holds when $\gamma < 1$. (The price-dividend ratio in (42) is *always* convex in β; the sole exception is the log-utility case $\gamma = 1$, when the price-dividend ratio is linear in β and there are no hedging demands.) Consequently, the result that the asset price on average goes down after strong experimentation remains valid. Panel (d) of Figure 10 depicts the average price paths after $t = 0$ in an economy with $n = 5$ competitors and is similar with panel (b) in Figure 3. Thus, our model continues to generate negative return predictability also in the case $\gamma < 1$, and Table 2 confirms. It presents results from simulated data at quarterly frequency, as in Table 1. The two panels of the table compare the first-best economy ($X^* = 0.314$, see also Figure 10, panel (a)) with an economy with competition and $n = 5$. Predictability is now present in both panels. We notice that average excess returns now increase with competition, in line with the risk premium result in panel (a) of Figure 10.

Social Optimum, $X^* = 0.314$ (medians of 1,000 simulations)

	4Q	12Q	20Q	40Q
Coefficient of Log (P/D)	-1.126	-3.258	-5.752	-10.85
t-stat	-1.939	-2.111	-2.371	-3.071
R-squared	0.028	0.082	0.138	0.268
Expected Excess Returns (annualized)	0.84%	0.72%	0.67%	0.50%

Competition, $n = 5$, $X_5^* = 2.444$ (medians of 1,000 simulations)

	4Q	12Q	20Q	40Q
Coefficient of Log (P/D)	-0.824	-2.265	-3.545	-6.653
t-stat	-2.092	-2.335	-2.504	-3.219
R-squared	0.033	0.097	0.151	0.262
Expected Excess Returns (annualized)	7.1%	5.5%	3.6%	4.4%

Table 2: **Return Predictability with the Price-Dividend Ratio (Simulations) when $\gamma < 1$.** This table is similar with Table 1 and reports the predictability of excess stock returns with the log price-dividend ratio. The two panels correspond to the social optimum and to the economy with $n = 5$ competitors. Both panels use a calibration with $\gamma < 1$ (see Figure 10).