Uncertain Innovation and Growth Traps

Franziska Bender

August 2025

Abstract

Does the inherent uncertainty of innovation make economies vulnerable to persistent stagnation? I study this question in an expanding-variety endogenous growth model with two features absent from standard frameworks: aggregate uncertainty about the returns to R&D, and learning from market outcomes. Firms cannot observe the conditions determining their future profits and instead infer them from the noisy outcomes of other innovators. Because firms do not account for the value of the information their R&D reveals to others, this learning creates an informational externality that can produce growth traps – periods where innovation and growth remain persistently low even after fundamentals have recovered. The mechanism is belief persistence: pessimistic firms invest little, which chokes off the signals that would reveal improved conditions, so beliefs recover only sluggishly. Strikingly, these traps arise even when uncertainty, in itself, encourages R&D, highlighting the crucial role of belief dynamics in shaping long-run growth.

Keywords: Endogenous growth, Growth traps, Aggregate uncertainty, Bayesian belief updating.

JEL Codes: O41, O31, D80

1 Introduction

Innovation is the engine of economic growth, and its inherent uncertainty is a critical feature of the R&D process. Standard models of endogenous growth reflect this, typically by incorporating firm-level risks such as a stochastic arrival rate of innovations (Klette and Kortum, 2004) or uncertainty over an innovation's ultimate quality and value (Acemoglu et al., 2018; Akcigit and Kerr, 2018). While these frameworks capture important micro-level risks, the uncertainty is characteristically idiosyncratic; it affects individual firms or projects independently, and in the aggregate, these risks are assumed to wash out. Furthermore, firms in these models are typically assumed to operate with full information, they are assumed to know the true probability distributions governing innovation. Consequently, these models typically do not account for two features that are central to real-world R&D decisions: aggregate uncertainty that affects firms in a related way, and ambiguity about the true probability distribution of returns, which necessitates that firms form subjective beliefs.

It is reasonable to assume, however, that these latter forms of uncertainty are highly relevant for firms' R&D decisions. Consider for example the decision problem faced by a firm contemplating a large R&D investment in 'green' technology. The firm recognize a shift in consumer preferences toward sustainability, but the true depth of this new market is unknown. Will consumers pay a significant premium for an electric vehicle? Is the demand for new energy storage solutions robust enough to justify a large upfront cost? The answers depend on a complex mix of consumer sentiment, future regulations, and evolving supply chains – factors that create a shared, aggregate uncertainty for the entire sector. The firm must therefore navigate a situation of ambiguity, forming its own judgments about the potential returns rather than relying on known probabilities. Faced with this ambiguity, firms rarely make decisions in isolation; instead, they form their judgments based on whatever information is available, for example by closely watching the outcomes of other related firms. When an early competitor launches a new electric vehicle, its reception – whether it becomes a market success or a costly failure – provides a crucial, albeit noisy, public signal about underlying factors that affect the profitability of similar innovations. This example highlights that the uncertainty relevant for innovation is not just idiosyncratic; firms may also face significant aggregate uncertainty and must form beliefs about profitability based on limited information rather than on known probability distributions.

In this paper, I develop a model designed to capture these key features. I extend a standard expanding-variety model of endogenous growth by incorporating (i) aggregate uncertainty about the returns to R&D and (ii) learning from market outcomes. In the model, firms invest in R&D to create new product varieties, the profitability of which depends on a stochastic state variable that is not directly observable. This state variable can be thought of as capturing factors, such as the underlying preferences of consumers, changes in the regulatory environment, the cost of key inputs, or the latent potential of a new general-purpose technology. These factors affect the returns to R&D and create aggregate uncertainty as they affect the outcomes of firms in the economy or in a certain sector in a related manner. The assumption that the state variable is not directly observed reflects the ambiguity. In the model firms must form beliefs, and they do so by observing the market outcomes of other innovating firms, which are similarly affected by these latent conditions.

This learning mechanism gives rise to an informational externality. When firms invest and their subsequent market success (or failure) is observed, they inadvertently reveal information about the underlying economic state. This information is valuable to other firms, helping them refine their own beliefs and make more informed R&D decisions. However, innovating firms do not internalize this social value of the information they produce. This informational spillover is distinct from the two externalities typically emphasized in the endogenous growth literature: the negative "business-stealing" effect of creative destruction and the positive "knowledge spillover" effect where firms build on the technological discoveries of others. The informational externality identified here is conceptually distinct from a knowledge spillover: it does not make subsequent innovation cheaper or of intrinsically higher quality, but rather allows for better-informed investment decisions by leading to more precise beliefs about the prevailing market conditions.

The learning mechanism I introduce is adapted from Fajgelbaum et al. (2017), who study this type of learning and the resulting informational externality in a business cycle framework with physical capital investment. A central result of their model is the emergence of *uncertainty traps*, a self-reinforcing cycle where high uncertainty suppresses investment, and the resulting lack of new information keeps uncertainty high. Building on their insights, this paper investigates whether such traps can also emerge in the context of R&D investment. This is a crucial extension, as a persistent slowdown in innovation would have lasting implications for long-run growth, associated with even higher welfare costs. Furthermore, shifting the focus to innovation is particularly interesting because the relationship between uncertainty and R&D investment is theoretically ambiguous. For R&D projects, the combination of limited downside risk with the potential for a high-payoff discovery can create an incentive to invest more when uncertainty is high, potentially counteracting the mechanism that leads to an uncertainty trap. A central question of this paper is therefore whether such trap dynamics still emerge in an endogenous growth model, accounting for the possibility that the direct effect of uncertainty on R&D investment is positive.

The central finding of this paper is that these growth traps do emerge, but through a mechanism driven by belief persistence rather than a self-reinforcing cycle of uncertainty. I show that even when higher uncertainty encourages R&D investment, ceteris paribus, a period of pessimistic beliefs can lead to a prolonged innovation slump. When firms are pessimistic, R&D investment is low. This, in turn, stifles the flow of public signals, causing beliefs to revert only sluggishly. As a result, firms can remain pessimistic and therefore reluctant to invest, long after underlying market conditions have recovered.

The model presented here is highly stylized and not intended to provide quantitative predictions about the prevalence or welfare costs of growth traps in reality. Rather, its primary contribution is to provide a proof of concept: to demonstrate that information-based trap dynamics can emerge in an endogenous growth context. Establishing this theoretical possibility is important, particularly because these traps can form even when uncertainty has a positive effect on R&D – a condition that might otherwise be expected to prevent such dynamics.

The remainder of the paper is structured as follows. After a brief review of the related literature, Section 2 develops the model, first presenting a benchmark case with full information before introducing the learning mechanism. Section 3 presents the main results, demonstrating how the informational externality leads to the emergence of growth traps. It also explores the endogenous relationship between uncertainty and R&D, discusses potential policy implications and the limitations of the model. Section 4 concludes.

1.1 Related Literature

The paper primarily contributes to the endogenous growth literature by incorporating aggregate uncertainty and belief formation, features that are typically absent from standard frameworks that focus on idiosyncratic risk and assume full information (e.g., Klette and Kortum, 2004; Lentz and Mortensen, 2008; Acemoglu et al., 2018; Akcigit and Kerr, 2018). The analysis is set within an expanding-variety framework, similar to recent work by Gersbach et al. (2018) and Gabrovsky (2020).

A related literature studies firm dynamics under parameter uncertainty and learning from market outcomes where that uncertainty is firm-specific. A seminal example is Jovanovic (1982), where firms are uncertain about their own efficiency and learn by observing their costs. More recently, Arkolakis et al. (2018) model firms that learn about idiosyncratic demand parameters by observing market prices. A key difference in the model presented here is that the parameter uncertainty pertains to an aggregate variable rather than a firm-specific one, allowing firms to learn from the outcomes of their peers.

The social learning mechanism in this paper is therefore most closely related to the work of Fajgelbaum et al. (2017), who is uncertainty about the aggregate returns to capital investment within a real business cycle framework. The key departure from their paper, is to embed this mechanism of social learning under aggregate uncertainty into an endogenous growth framework. This change of setting is crucial because the nature of R&D investment differs from the physical capital investment at the center of their model. While higher uncertainty often suppresses physical investment due to an 'option value of waiting', the effect of uncertainty on R&D investment can be positive. The potential for a high-payoff discovery can create a powerful incentive for innovation that counteracts the desire to wait. A central question of this paper is therefore to explore the aggregate dynamics that emerge from this tension. We investigate whether the unique nature of R&D is sufficient to overcome the informational frictions that cause 'uncertainty traps', or if similar trap-like dynamics can ultimately emerge in an endogenous growth context.

More broadly the paper is also related to parts of the business cycle literature, in particular the literature that studies the persistence of business cycles or hysteresis (Benigno and Fornaro, 2018; Queralto, 2020; Anzoategui et al., 2019; Cerra et al., 2023) and the literature studying the interplay between uncertainty and business cycles (Bloom et al., 2018; Bachmann and Moscarini, 2011; Taschereau-Dumouchel, 2022).

2 The Model

The foundation of the analysis is a standard expanding-variety model of endogenous growth. In this framework, long-run growth is driven by the innovation activities of firms. Firms undertake research and development (R&D) investments to create new product varieties. The successful introduction of a new variety grants the innovating firm monopoly power over its production, and the continual expansion of the set of varieties drives aggregate economic growth.

To this standard framework, I introduce two principal features: (i) aggregate uncertainty regarding the returns to innovation, and (ii) a social learning mechanism through which firms form beliefs about these returns. I model aggregate uncertainty by assuming that the 'quality' of a newly created variety, which determines its profitability, is stochastic. Specifically, the quality of a successful innovation depends on two components: an unobserved, persistent fundamental variable, θ, and an idiosyncratic random shock. The fundamental variable θ represents the aggregate state of innovative potential and is assumed to follow an AR(1) process, introducing aggregate fluctuations in the expected returns to R&D.

Since firms cannot directly observe the fundamental component θ, they face parameter uncertainty and must form subjective beliefs about its current value. This is where the second key feature, learning from market outcomes, comes into play. When a firm successfully innovates and brings a new variety to market, its subsequent performance provides a public, albeit noisy, signal about the current state of θ. Other firms in the economy can observe these outcomes and use them to update their own beliefs about the aggregate innovative environment.

This social learning process links firms' R&D decisions through an informational channel. Specifically, when more firms innovate, more public information is generated, which in turn reduces the uncertainty faced by all other potential innovators. Because individual firms do not internalize the value of the information their R&D activities provide to others, this process creates an informational externality. This externality raises the possibility of self-reinforcing dynamics, where an initial wave of pessimism could lead to low R&D investment, which in turn stifles the flow of new information, thus perpetuating high uncertainty and deterring future innovation – a dynamic that could culminate in a *growth trap*.

2.1 Households

Households in this economy supply labor and choose consumption and savings over time to maximize their expected lifetime utility, given by:

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t u(C_t)\right],$$

where $\beta \in (0,1)$ is the discount factor and $u(\cdot)$ is a standard concave utility function. For simplicity, households inelastically supply one unit of labor ($L_t = 1$), earning a wage W_t. They also own the intermediate firms and receive profits in the form of dividends, D_t. Each period, households make the decision to allocate their income between consumption (C_t) and savings (I_t). The budget constraint is therefore:

$$C_t + I_t = W_t + D_t$$

The household's optimal trade-off between consuming today versus saving for the future gives rise to the stochastic discount factor, $M_{t,t+s} = \beta^s \frac{u'(C_{t+s})}{u'(C_t)}$, which firms use to discount future profits. The marginal utility of consumption, $\lambda_t = u'(C_t)$, is thus the key component for valuing investment projects.

2.2 Final Good Production

The final good, Y_t, is produced by a representative and perfectly competitive firm. This firm's production technology combines labor, L_t, with a range of differentiated intermediate inputs, $x_{j,t}$. Each intermediate variety j is characterized by a specific quality level, $q_{j,t}$. The aggregate production function takes following form:

$$Y_t = L_t^{1-\alpha}\sum_{j=0}^{M_t} q_{j,t}^{1-\alpha} x_{j,t}^{\alpha}\ , \tag{1}$$

where M_t is the measure of available intermediate varieties at time t, and $\alpha \in (0,1)$ is the expenditure share of intermediate goods[1]. Taking the prices of intermediate goods, $P_{j,t}$, and the wage rate as given, the final good producer maximizes static profits. The first-order conditions from this optimization problem yield the firm's conditional demand function for each intermediate variety j:

$$x_{j,t} = L_t \left(\frac{\alpha}{P_{j,t}} \right)^{\frac{1}{1-\alpha}} q_{j,t} \tag{2}$$

The demand for intermediate good of variety j increases in its quality $q_{j,t}$ and decreases in it's price $P_{j,t}$.

2.3 Intermediate Goods Production and Profits

Each variety of intermediate good is produced by a single monopolistic firm that owns the patent for that specific variety. These patents are the result of successful past R&D investments. Production is simple: one unit of the final good is required to produce one unit of any intermediate good, implying a constant marginal cost of one.

The producer of variety j chooses its price, $P_{j,t}$, to maximize its operating profits, taking the demand function $x_{j,t}(P_{j,t})$ from the final good sector as given. The firm's problem is:

$$\max_{P_{j,t}} \ \pi_{j,t} = \max_{P_{j,t}} \ [P_{j,t} - 1] x_{j,t}$$

The solution to this problem yields a constant markup over marginal cost: $P^*_{j,t} = 1/\alpha$. Substituting this optimal price back into the demand function the production of intermediate good of variety j is

$$x_{j,t} = \alpha^{\frac{2}{1-\alpha}} q_{j,t} L_t \tag{3}$$

and then into the profit expression, the maximized operating profits for the producer of variety j are:

$$\pi_{j,t} = \left(\frac{1}{\alpha} - 1 \right) \alpha^{\frac{2}{1-\alpha}} q_{j,t} L_t \equiv \bar{\omega} q_{j,t} L_t \ ,$$

with $\bar{\omega} \equiv \left(\frac{1}{\alpha} - 1\right) \alpha^{\frac{2}{1-\alpha}}$. A key feature of this result is that a firm's operating profits depend on the quality of its variety $q_{j,t}$. They are however independent of the number or quality of other varieties in the market.

2.4 Entry and Uncertainty about the Returns to R&D

Every period outside firms can spend η on R&D to invent a new variety. A firm that invests today will invent a new variety with certainty[2] and can enter production in the next period. They face uncertainty about the quality of the new variety $q_{j,t+1}$ which is given by

$$q_{j,t+1} = \exp(\theta_{t+1}) \tag{4}$$

It depends on a fundamental variable θ_{t+1}, that is unobserved by firms, and that follows an AR(1) process:

$$\theta_{t+1} = \rho \theta_t + \varepsilon^{\theta}_{t+1}, \qquad \text{where} \qquad \varepsilon^{\theta}_t \sim N\Big(0, (1-\rho^2)\sigma^2_\theta\Big) \tag{5}$$

[1]See Appendix A.3

[2]Here we deviate from the typical assumption that there is uncertainty about the arrival of a successful innovation.

Innovating firms are assumed to operate for a single period, after which their patent expires and the monopoly rent is randomly allocated to another firm. Potential entrants make their R&D investment decision by maximizing the expected value of innovating. This value is the discounted profit from production, net of the initial R&D cost.

A crucial element of the R&D decision is the information set available to firms and how they form expectations about the profitability of new varieties. To isolate the role of social learning and potential informational externalities, I analyze two specifications. The first is a benchmark model with Full Information, where firms observe the current state θ_t but still face uncertainty about the future state θ_{t+1}. The second specification, Learning from Market Outcomes, assumes firms do not observe the true state θ_t. Instead, they start with prior beliefs and update them by observing the performance of recent innovators. As other externalities in the model are common to both specifications, this comparison highlights how the informational externality through learning from market outcomes can shape investment decisions and the path of the economy.

2.5 Full Information

In the full-information benchmark, firms perfectly observe the current state of innovative potential, θ_t. While this removes any ambiguity about the present, firms still face the crucial uncertainty over the future state, θ_{t+1}, which will determine the profitability of their R&D investment.

A potential entrant's decision to invest is based on the net present value of creating a new variety. The value of investing in R&D at a cost of η units of the final good is given by:

$$V_t = -\eta + \frac{1}{\eta}\mathbb{E}_t\Big[M_{t,t+1}\bar{\omega}q_{t+1}L_{t+1}\Big|\theta_t\Big],$$

The first term, $-\eta$, represents the cost of R&D. The second term is the expected benefit: the profit from the new variety ($\bar{\omega}q_{t+1}L_{t+1}$) discounted by the household's stochastic discount factor ($M_{t,t+1}$). The expectation is formed conditional on the observed state θ_t. In equilibrium, firms will enter and undertake R&D as long as the value of doing so is non-negative. This process continues until the expected profit is driven to zero, which is captured by the free-entry condition:

$$Z_t\Big(-\eta + \mathbb{E}_t\Big[M_{t,t+1}\bar{\omega}q_{t+1}L_{t+1}\Big|\theta_t\Big]\Big) = 0 \tag{6}$$

This condition pins down the aggregate R&D expenditure in the economy, denoted by Z_t. It ensures that if any R&D investment occurs ($Z_t > 0$), the value of creating a new variety must be competed down to exactly zero. If the expected value is negative, no firm will invest, and aggregate R&D expenditure is zero ($Z_t = 0$).

2.6 Learning from Market Outcomes

In this specification, firms face an additional layer of uncertainty: they do not observe the true state of innovative potential, θ_t. Instead, they must learn about it by observing the market.

The learning mechanism works as follows. At the beginning of period t, the cohort of firms that invested in R&D at $t-1$ enters the market. The profitability of these new firms provides a noisy public signal, S_t, about the current state, θ_t:

$$S_t = \theta_t + \varepsilon_t^S \quad \text{with} \quad \varepsilon_t^S \sim N\Big(0, \big(z_{t-1}\gamma_S\big)^{-1}\Big) \tag{7}$$

A key assumption of the model is that the precision of this signal, $z_{t-1}\gamma_S$, is increasing in the R&D

intensity of the previous period ($z_{t-1} \equiv Z_{t-1}/A_{t-1}$)[3], where aggregate R&D spending is scaled by the economy's productivity level, A_{t-1}. This captures the idea that a higher level of past innovation generates more information, reducing uncertainty for the entire economy. This creates an informational externality: when making their investment decisions at $t-1$, firms do not account for the fact that their entry will reveal valuable information about the state that informs the R&D decisions made by the next cohort of firms.

Potential entrants use this signal to update their beliefs. They begin the period with a prior belief about θ_t based on information from the previous period $\mathcal{I}_{t-1}$. This prior is characterized by a mean and a precision, formally $\theta_t|\mathcal{I}_{t-1} \sim (\mu_{t-1}, \gamma_{t-1})$. After observing S_t they update this belief using Bayes' rule and then project it forward to form an expectation about the future state, $\theta_{t+1}|\mathcal{I}_t \sim (\mu_t, \gamma_t)$. This process is described by the standard Kalman filter updating equations for the mean (μ_t) and precision (γ_t) of their beliefs about the next-period state:[4]

$$\mu_t = \rho \frac{\gamma_t \mu_{t-1} + z_{t-1}\gamma_S S_t}{\gamma_{t-1} + z_{t-1}\gamma_S} \tag{8}$$

$$\gamma_t = \left(\frac{\rho^2}{\gamma_{t-1} + z_{t-1}\gamma_S} + (1-\rho^2)\sigma_\theta^2 \right)^{-1} \tag{9}$$

The R&D investment decision is analogous to the full-information case, but it is now based on these subjective beliefs rather than the true state. The value of investing is given by:

$$V_t = -\eta + \frac{1}{\eta}\mathbb{E}_t\Big[M_{t,t+1}\bar{\omega}q_{t+1}L_{t+1}\Big|\mu_t\, \gamma_t\Big],$$

The firm's expectation is now conditional on its belief about the future state, summarized by the mean (μ_t) and precision γ_t. The free-entry condition operates as before, pinning down aggregate R&D spending Z_t:

$$Z_t\Big(-\eta + \mathbb{E}_t\Big[M_{t,t+1}\bar{\omega}q_{t+1}L_{t+1}\Big|\mu_t, \gamma_t\Big]\Big) = 0 \tag{10}$$

2.7 Aggregate Dynamics and Resource Constraint

This section describes the aggregate environment of the economy, including the laws of motion for key state variables and the economy-wide resource constraint.

Net output (GDP) is the sum of value added across all sectors. Given the optimal production decision of intermediate and final good producers this is[5]:

$$\text{GDP}_t \equiv Y_t - \sum_{j=0}^{M_t} x_{j,t} = \Psi L_t A_t,$$

with $\Psi \equiv \alpha^{\frac{2\alpha}{1-\alpha}} - \alpha^{\frac{2}{1-\alpha}}$. The variable $A_t \equiv \sum_{j=0}^{M_t} q_{j,t}$ represents the aggregate level of technology, which is the quality-weighted sum of all existing intermediate varieties.

The level of technology evolves according to the amount of R&D investment undertaken in the previous period. With a mass of Z_{t-1}/η entering and introducing new varieties of quality q_t, the law of motion for

[3]The specification of the signal (equation 7) is a tractable, reduced-form representation of an environment where firms observe the individual outcomes (e.g., profits) of producing firms directly. As shown in Appendix A.1, under certain assumptions this setup is equivalent to one where the average profitability of the new cohort serves as the public signal, and the precision of this signal is naturally increasing in the number of entrants.

[4]See Appendix A for a formal derivation of the updating rules.

[5]See Appendix A.2 for a derivation of net output (GDP).

aggregate technology is

$$A_t = A_{t-1} + \frac{Z_{t-1}}{\eta} q_t \tag{11}$$

The growth rate of technology, $g_t \equiv (A_t - A_{t-1})/A_{t-1}$, can therefore be expressed as a function of the R&D intensity, $z_{t-1} \equiv Z_{t-1}/A_{t-1}$, and the quality of new innovations:

$$g_t = \frac{1}{\eta} z_{t-1} q_t \tag{12}$$

Finally, the aggregate resource constraint ensures that total output is allocated between consumption and investment in R&D:

$$\text{GDP}_t = C_t + Z_t \tag{13}$$

2.8 Equilibrium

A recursive competitive equilibrium consists of a set of policy functions, pricing functions, and laws of motion for the aggregate state variables such that all agents optimize given their information, and all markets clear. The key policy function is for the normalized R&D intensity, z_t, which depends on the relevant state variables in each specification of the model.

An equilibrium is defined by the following conditions:

1. Profit Maximization: The pricing function for intermediate goods producers is $P^* = 1/\alpha$.
2. Free entry: The R&D policy function satisfies the free-entry condition.
   - In the full information specification the policy function is $z(\theta)$, which depends on the true state θ.
   - In the Learning from Market Outcomes specification, the policy function is $z(\mu, \gamma)$, which depends on the mean and precision of firms' beliefs.
3. Household optimization: Household optimization is satisfied, as reflected in the firm's use of the correct stochastic discount factor to value future profits.
4. Market Clearing: The aggregate resource constraint holds in all periods.
5. Rational Beliefs: In the Learning from Market Outcomes specification, firms' beliefs are updated according to the Kalman filter equations.
6. Laws of Motion: The aggregate state variables evolve according to their specified processes. The fundamental state, θ, follows its exogenous AR(1) process, and aggregate technology, A, evolves according to its law of motion.

2.9 Numerical Solution

The model is solved numerically to find the time-invariant policy functions for R&D intensity, $z^*(\theta)$ and $z^*(\mu, \gamma)$ for each specification. To do so, it is useful to express the free-entry conditions in a more operational form. This requires specifying a functional form for the household's utility to derive an explicit expression for the stochastic discount factor.

I assume a CRRA utility function, $u(C) = \frac{C^{1-\sigma}-1}{1-\sigma}$, and a constant labor supply $L = 1$. The aggregate resource constraint is then $\Psi A_t - Z_t = C_t$. Using the assumption on the utility function and the aggregate resource constraint the stochastic discount factor can be written in terms of the growth rate of the economy g and R&D intensity z:

$$
\begin{aligned}
M_{t,t+1} = \beta \frac{u'(C_{t+1})}{u(C_t)} &= \beta \frac{(\Psi A_{t+1} - Z_{t+1})^{-\sigma}}{(\Psi A_t - Z_t)^{-\sigma}} \\
&= \beta \Big(\frac{A_{t+1}}{A_t}\Big)^{-\sigma} \left(\frac{\Psi - z_{t+1}}{\Psi - z_t}\right)^{-\sigma} \\
&= \beta (1 + g_{t+1})^{-\sigma} \left(\frac{\Psi - z_{t+1}}{\Psi - z_t}\right)^{-\sigma}
\end{aligned}
$$

Substituting this expression back into the free-entry conditions yields the key equilibrium equations that the policy functions must satisfy.

For the Full Information specification, the condition is:

$$
-\eta + \mathbb{E}_t \left[\beta (1 + g_{t+1})^{-\sigma} \left(\frac{\Psi - z_{t+1}}{\Psi - z_t}\right)^{-\sigma} \bar{\omega} q_{t+1} \middle| \theta_t \right] \le 0 \quad \text{with equality if} \quad z_t > 0, \tag{14}
$$

For the Learning from Market Outcomes specification, the condition is:

$$
-\eta + \mathbb{E}_t \left[\beta (1 + g_{t+1})^{-\sigma} \left(\frac{\Psi - z_{t+1}}{\Psi - z_t}\right)^{-\sigma} \bar{\omega} q_{t+1} \middle| \mu_t, \gamma_t \right] \le 0 \quad \text{with equality if} \quad z_t > 0, \tag{15}
$$

With these conditions, the policy functions can be found using policy function iteration over a discretized state space. The algorithm proceeds as follows. First, the continuous state variables – θ for the full information model, and (μ, γ) for the learning model – are discretized onto a finite grid. The iteration begins with an initial guess for the policy function, z^0, for all points on the grid. For each guess $z^k(\cdot)$, a new policy function, $z^{k+1}(\cdot)$, is computed at each grid point by solving the free-entry condition using a numerical root-finding algorithm. The calculation first checks if R&D is profitable; if the value of entry at z_t is negative, the optimal policy is the corner solution $z^{k+1} = 0$. Otherwise, the interior solution is found. This process is repeated until the policy function converges.

Once the converged policy function $z^*(\cdot)$ is obtained, it can be used in conjunction with the law of motion for technology (equation 11) to simulate the dynamic path of the economy.

2.10 Parameter

The model is solved numerically, which requires specific values for its parameters. Given that the model is highly stylized, the goal of this exercise is not to perform a formal calibration to match specific data moments or to derive precise quantitative implications. Instead, the aim is to provide a clear illustration of the key economic mechanisms at play. The parameter values used in this analysis are summarized in Table 1.

Table 1: Parameter

Parameter	**Value**	**Description**
α	0.6	Parameter of the production function
β	0.95	Households discount rate
σ	2	Coefficient of relative risk aversion
ρ	0.9	Persistence of the AR(1) process for θ
σ_θ	0.229	Standard deviation of the ergodic distribution of θ
η	0.048	Cost of R&D
γ_S	200	Scale parameter for precision of signals

The parameter α is set to 0.6. This choice is guided by its implications for two economically meaningful quantities in the model: the markup charged by intermediate goods producers and the aggregate labor share. The chosen value of α implies a markup, given by $\mu = (\frac{1}{\alpha} - 1)$, of approximately 67%, which is within the range of empirical estimates for firm markups (Syverson, 2025). It also implies that the labor share of income[6] is 62.5%, which is consistent with observed aggregate data.

For household preferences, the discount rate β is set to 0.95. Utility is assumed to follow a standard CRRA specification, with the coefficient of relative risk aversion σ set to 2.

The parameters governing the exogenous process for innovative potential, θ_t, are crucial for the model's learning dynamics. For learning to be valuable, the state must be persistent $\rho > 0$, as this ensures that information gathered today is useful for predicting future market conditions. However, if the process is too noisy (high σ_θ), the value of knowing the current state is diminished, which reduces the incentive to learn. The parameters are chosen to ensure that significant learning is possible, the persistence parameter is set to $\rho = 0.9$, and the standard deviation of the ergodic distribution is set to $\sigma_\theta = 0.229$.

The R&D cost, η, is chosen to target the average growth rate in the full information benchmark. Specifically, $\eta = 0.048$ is set to yield an average growth rate of approximately 5%. While this target is higher than realistic rates of economic growth, the choice is deliberate. As the results in section **??** will show, the introduction of learning and informational frictions systematically lowers the growth rate. Setting a higher benchmark target ensures that the economy with learning still exhibits positive growth, allowing for a meaningful analysis of the trap dynamics without the economy simply shutting down.

Finally, the parameter γ_S scales the amount of information revealed by R&D investment. The precision of the public signal is given by $z_{t-1}\gamma_S$, meaning that higher values of γ_S imply that any given level of R&D intensity is more informative. The choice of γ_S must be high enough for investment to be informative, but not so high that a marginal amount of investment would make the signal almost perfectly revealing. The value is set to γ_S to ensure that R&D investment provides a meaningful but imperfect signal about the underlying state of the economy.

3 Results

3.1 Optimal R&D Investment

Figure 1 shows the optimal function, z^*, for both the full information benchmark and the main specification with learning from market outcomes.

The left panel of figure 1 illustrates the R&D policy in the full information benchmark. The decision rule is a threshold policy based on the observed fundamental, θ_t. Below a certain threshold, it is not optimal to invest, as firms expect the quality of their innovation, and therefore its profitability, to be too low to justify the R&D cost. Above this threshold, R&D investment is positive and strictly increasing in θ_t. This is because a higher fundamental directly increases firms' expectations about the quality and profitability of a new variety.

The right panel of figure 1 shows the policy function for the model with learning from market outcomes. Here firms do not observe the fundamental, θ, directly. Instead the optimal R&D policy is a function of both the mean (μ_t) and the precision (γ_t) of firms' beliefs. For any given level of precision, the policy has a similar threshold structure: firms only invest if they are sufficiently optimistic (i.e., if μ_t is above a certain threshold).

An important observation for the subsequent analysis is the role of subjective uncertainty. The figure shows that higher uncertainty (a lower precision, γ_t) encourages firms to begin investing at a lower

[6] The labor share of income is $1/(1 + \alpha)$, see Appendix A.4.

Figure 1: Optimal R&D Investment

(a) Full Information: $z^*(\theta_t)$ (b) With Learning: $z^*(\mu_t, \gamma_t)$

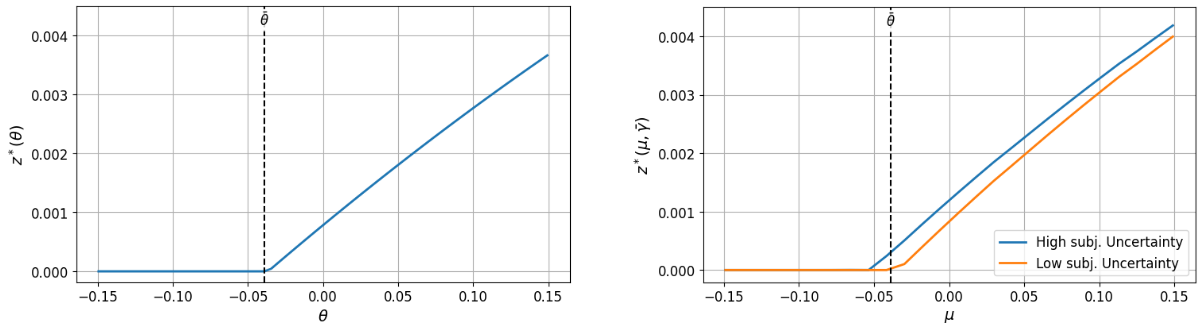

Notes: The figure compares the optimal R&D intensity policy functions from the two model specifications. Panel (a) shows the policy function for the Full Information model, where R&D intensity is a function of the true state, θ. Panel (b) shows the policy function for the model With Learning, where R&D intensity is a function of firms' beliefs . The two lines in panel (b) show the policy function for a high and low level of subjective uncertainty (i.e., a low and high precision, γ_t, of beliefs), plotted against the mean of beliefs, μ_t.

threshold of mean beliefs, μ_t. Furthermore, for any level of beliefs above this threshold, firms invest more when their subjective uncertainty is higher[7]. This seemingly counter-intuitive result can be understood through the lens of real options theory. R&D investment is an option on a future profit stream with limited downside risk—the most a firm can lose is its initial investment cost. When uncertainty is high, the potential upside (discovering a very high-quality innovation) increases, making the R&D option more valuable and encouraging investment.

The positive effect of subjective uncertainty on the investment policy raises a crucial question. If higher uncertainty encourages firms to invest more, does this rule out the possibility of a self-reinforcing trap where high uncertainty stifles R&D? As the next section will show, the answer is no. The dynamics of belief updating can still lead to coordination failures and persistent low-growth outcomes.

3.2 Emergence of Growth Traps

This section uses numerical simulations to demonstrate how the informational externality in the model can lead to the emergence of persistent low-growth traps. The analysis proceeds in two steps. First, a single simulation path is used to provide an illustration of the mechanism and build intuition for how a trap can form. Second, to show that these traps are a robust feature of the model with learning – and not an artifact of a single realization of shocks – the analysis is extended to a large number of simulations.

The growth rate of the economy is determined by the interplay of two key drivers. The first is the endogenous R&D investment decision of firms, z_t. The second is the exogenous state of innovative potential, θ_t, which determines the quality of new varieties and thus the returns to R&D. Figure 2 plots the exogenous path of the fundamental state, θ_t, for the simulation. High values of θ_t correspond to periods of high innovative potential, where the expected quality of new products and the returns to R&D are high. The simulation begins with a period of high innovative potential, followed by a prolonged, exogenously driven slump where the returns to R&D are low. Eventually, the fundamental state recovers.

Given the exogenous path for innovative potential, θ_t, figure 3 shows the resulting growth rate of the economy for the two model specifications: the benchmark case where firms observe the state directly, and

[7]While the figure focuses on the region around the investment threshold, this relationship is not globally monotonic. For very high levels of optimism (i.e., for values of μ_t in the far right tail of its distribution), the effect reverses, and higher uncertainty leads to slightly lower investment. However, this region of the state space is rarely visited in simulations, and the magnitude of the effect is quantitatively much smaller than the positive effect of uncertainty on investment for moderate levels of μ_t. Furthermore, the dynamics of a potential growth trap are driven by the economy's behavior near this investment threshold, making the policy function in this region the most relevant for the subsequent analysis.

Figure 2: Evolution of θ over Time

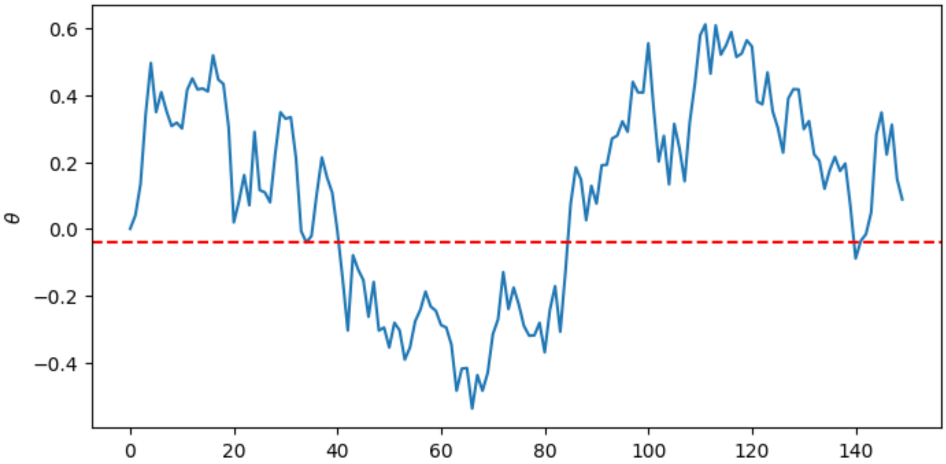

Notes: The figure shows a simulated time series path for the fundamental state of innovative potential, θ_t, over $T = 150$ periods. The dashed red line represents the investment threshold, $\bar{\theta} = -0.039$, from the full information benchmark.

the case where firms must form beliefs based on market outcomes. In the full information benchmark (blue) firms' investment decisions respond directly and immediately to the fundamental state. When θ_t is high, firms invest, and the economy grows. When θ_t falls below the investment threshold (around period 40), R&D and growth come to a halt. As soon as the fundamental state recovers and crosses the threshold again (around period 85), investment and growth promptly resume.

The dynamics of the growth rate are different when firms do not observe the state, but learn from market outcomes (orange). During the initial period of high fundamentals, the growth path closely tracks the benchmark. However, the response to the subsequent slump and recovery is starkly different. The prolonged period of low fundamentals leads to a sharp decline in R&D investment. While investment does not fall to zero as it does in the benchmark, it remains at a very low level and, crucially, fails to recover even when the fundamental state improves. This leads to a persistent low-growth trap, highlighted by the shaded area.

Figure 3: Simulated Growth Dynamics: Full Information vs. Learning

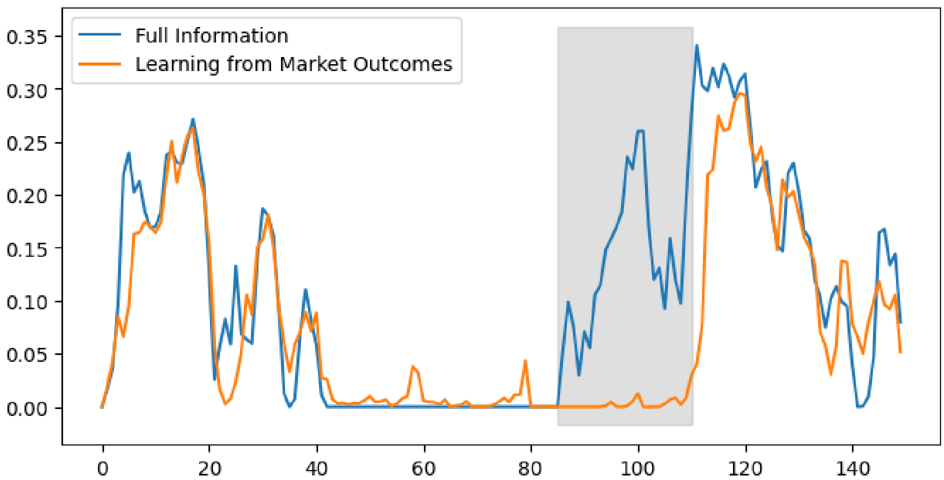

Notes: The figure plots the simulated equilibrium growth rates for the Full Information model and the model with Learning from Market Outcomes, conditional on the same exogenous path for the fundamental state, θ_t. The shaded gray area identifies a growth trap, defined as a span of at least five consecutive periods where the R&D intensity in the learning model is less than 20% of the level in the full information benchmark.

The underlying mechanism for this trap is revealed in Figure 4, which plots the path of the fundamental state, θ_t, against the endogenous dynamics of firms' beliefs, μ_t. In the initial phase of the simulation, R&D investment is high, which generates precise public signals about the state. As a result, the mean of firms' beliefs tracks the true fundamental state relatively closely. However, when the economy enters the slump and investment falls, the flow of information dries up. The public signals become very imprecise, causing a gap to open between the true state and firms' beliefs. This gap persists even as the fundamental state begins to recover. Firms' beliefs remain pessimistic for an extended period because there is not enough R&D investment to generate the precise signals that would reveal the true improvement in market conditions. Only when a sufficient number of firms begin to invest again does the information flow resume, allowing beliefs to catch up to the fundamental reality.

Figure 4: Dynamics of Beliefs

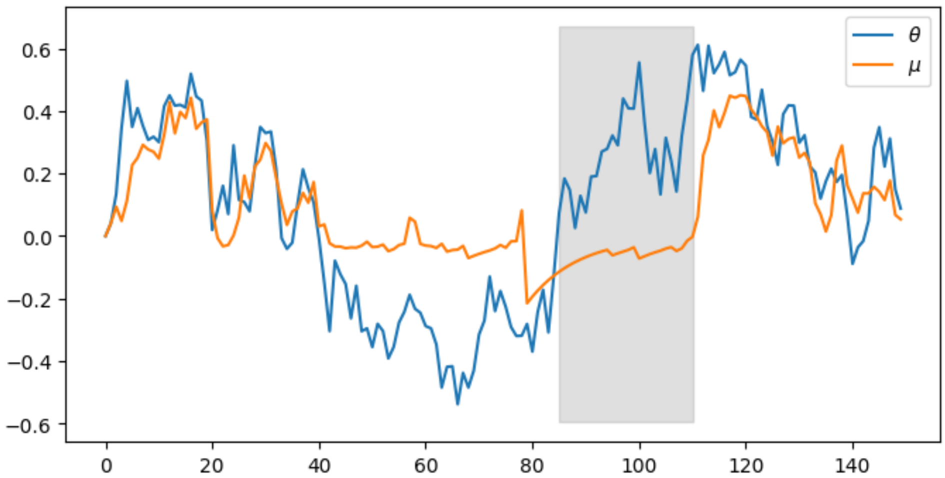

Notes: The figure plots the simulated time series path for the true fundamental state, θ_t, against the mean of firms' beliefs, μ_t. The shaded gray area identifies the growth trap period, as defined in Figure 3.

This dynamic highlights the role of the information spillover. When investment is high, each innovating firm contributes to a public good: a more precise understanding of the aggregate state. Because individual firms do not internalize the value of the information their R&D provides to others, this spillover represents a positive externality. During the initial high-growth phase, this externality allows the market as a whole to form accurate beliefs, fostering efficient investment. The growth trap, therefore, can be understood as a coordination failure where the production of this public good collapses, leaving all firms to navigate a period of high uncertainty with outdated, pessimistic beliefs.

This outcome, a prolonged period of low growth, is similar to the uncertainty traps described in Fajgelbaum et al. (2017), although the underlying mechanism is crucially different. In their model of capital investment, traps emerge from wait-and-see dynamics: higher uncertainty increases the option value of waiting, which suppresses costly and irreversible investment, leading to a self-reinforcing episode of high uncertainty and low economic activity. In the context of R&D, however, the emergence of a trap is surprising. The effect of uncertainty on R&D may be - and in the model presented here is - positive[8], a condition that would typically break the self-reinforcing cycle of high uncertainty and low investment seen in other models.

The trap in this model is therefore not driven by uncertainty directly deterring investment. Instead, it arises from the persistence of pessimistic beliefs when the flow of new information ceases during a

[8]The relationship between uncertainty and R&D investment is the subject of an ongoing debate in both the theoretical and empirical literature, with findings pointing to both positive and negative effects. The structure of this model is deliberately chosen to be consistent with the growth-options channel, which generates a positive relationship between uncertainty and R&D. This provides a stronger test for the emergence of a trap. If a trap can form even when uncertainty encourages investment - a condition that would typically break the self-reinforcing cycle seen in other models - it highlights the strength of the belief-persistence mechanism. A negative relationship would only amplify the trap-like dynamics.

downturn. This persistent pessimism dominates the positive 'growth-option' effect of uncertainty. The growth-option channel is only powerful when firms are optimistic enough to be near the investment threshold; if the mean of their beliefs remains far below this point, the investment-enhancing effect of uncertainty is never activated. This dynamic also means the trap is not permanent. The positive relationship between firms subjective uncertainty and R&D provides a slow but reliable endogenous escape mechanism: in the absence of new information, beliefs slowly mean-revert, and the high uncertainty eventually encourages a few firms to invest, restarting the flow of information and allowing the economy to escape the trap[9].

3.3 Robustness of Growth Traps

The previous section used a single simulation to illustrate the mechanism through which a growth trap can emerge. To demonstrate that this outcome is a robust feature of the model with learning, and not just an artifact of a particular sequence of shocks, this section analyzes the results from a large number of simulations. The model was simulated 100,000 times, each for 150 periods, under both the full information and the learning specifications. For each simulation, the average growth rate was recorded.

Figure 5: Distribution of Average Growth: Full Information vs. Learning

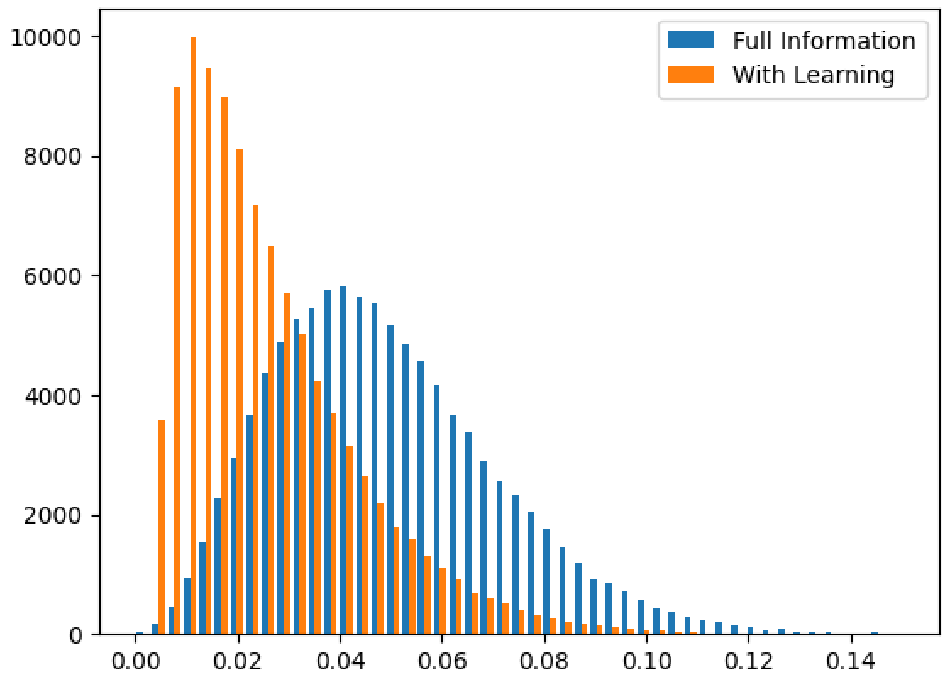

Notes: The figure shows the distribution of average growth rates from 100,000 simulations of the economy, each run for 150 periods. The blue histogram corresponds to the Full Information benchmark, and the orange histogram corresponds to the model With Learning.

Figure 5 shows the distribution of these average growth rates for both versions of the model. Since both economies are subject to aggregate uncertainty, the average growth rate is stochastic and depends on the specific realization of the fundamental shocks. The distribution for the full information benchmark (in blue) provides a natural reference point. The distribution for the economy with learning (in orange) is clearly shifted to the left, with its mass concentrated at lower average growth rates. This confirms that the growth traps illustrated previously are a systematic feature of the model with learning from market outcomes, which consistently reduces the economy's long-run growth performance.

[9]The uncertainty traps described in Fajgelbaum et al. (2017) are "near permanent", for the economy to escape a trap a sufficiently large and positive temporary shock to either the fundamental or to beliefs is necessary.

To quantify the prevalence and duration of these traps more directly, the number of "trap periods" was calculated for each of the 100,000 simulations. Conceptually, a growth trap is a situation where the fundamental state of the economy is high enough to warrant R&D investment, but firms remain pessimistic and do not invest due to a lack of information. To operationalize this, a period is defined as part of a trap if R&D intensity is less than 20% of the level that would occur under full information, and this condition holds for at least five consecutive periods[10]. The relative threshold captures the idea of a significant investment shortfall, while the persistence requirement ensures that the definition distinguishes a true trap from a brief, natural lag in adjustment as the economy recovers from a downturn. Applying this definition to the 100,000 simulations reveals that growth traps are a frequent occurrence, emerging in 92% of all simulations with an average duration of 19 periods.

The central takeaway is that in a model with aggregate uncertainty where firms learn from the market, average growth is systematically lower than in an equivalent full information benchmark. This difference is driven by an informational externality: individual firms do not internalize that their investment reveals valuable information to others, leading to the frequent emergence of persistent growth traps. A sensitivity analysis, presented in Appendix B, shows that this is a robust feature of the model. While the quantitative magnitude of the growth gap and the frequency of traps depend on the specific parameter values, their existence is a consistent outcome across a wide range of plausible parameterizations. This does not imply however, that all parameter combinations will lead to such outcomes, nor can I make a quantitative claim about their prevalence in reality. The primary contribution of this highly stylized model is to demonstrate that growth traps *can* emerge, even under conditions where uncertainty, ceteris paribus, encourages R&D investment. Given the potentially large welfare costs associated with prolonged periods of low innovation, establishing this theoretical possibility underscores the importance of considering belief dynamics for long-run growth.

3.4 The Endogenous Relationship Between Uncertainty and R&D

The analysis so far has established that aggregate uncertainty about the returns to R&D, in combination with the informational spillover, leads to the emergence of growth traps. This section delves deeper into the relationship between R&D investment and subjective uncertainty, revealing a dynamic that connects this model to a wider debate in the business cycle literature.

Figure 6 plots the relationship between R&D investment and subjective uncertainty from the model simulations. Both the contemporaneous relationship (Panel a) and the relationship between investment at time t and uncertainty in the next period at time $t+1$ (Panel b) are clearly negative. Periods of high R&D investment are strongly associated with periods of low uncertainty.

This negative aggregate relationship appears, at first glance, to contradict the firm-level incentive discussed in Section 3.1. The R&D policy function (Figure 1b) showed that, ceteris paribus, higher uncertainty increases the option value of R&D and encourages more investment. The key to reconciling these two findings lies in the endogeneity of uncertainty. In this model, uncertainty is not an exogenous force; rather, it is determined by the amount of information generated by firms' collective actions. When more firms invest, they produce more public signals about the state of innovative potential, which in turn reduces the uncertainty faced by the next cohort of firms. This mechanism is visible in Panel (6b), which shows that high investment in period t is a strong predictor of low uncertainty in period $t+1$. Because both investment and uncertainty are persistent, this dynamic also results in the negative contemporaneous relationship seen in Panel (6a).

[10]The specific choices for the relative R&D threshold (20%) and the minimum duration (5 periods) are necessarily somewhat arbitrary. However, the qualitative results concerning the frequency and persistence of traps are robust to reasonable alternative specifications of these parameters.

Figure 6: Relationship between R&D Investment and Uncertainty

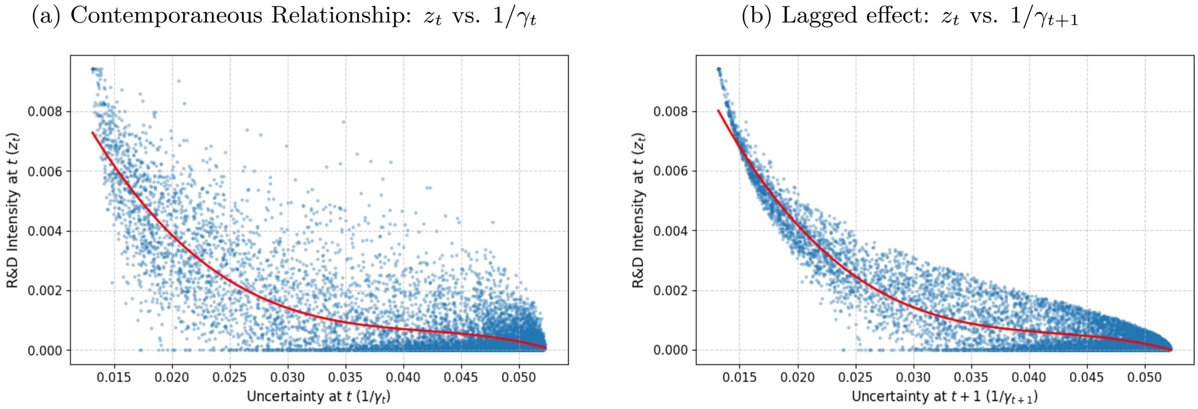

Notes: The figure displays the relationship between R&D investment (z_t) and economic uncertainty ($1/\gamma_t$) using data from 10,000 model simulations. Each blue dot represents the outcome from a single simulation. The red line is a third-degree polynomial trend fitted to the data. Panel (a) shows the contemporaneous relationship between investment and uncertainty. Panel (b) plots the investment rate at time t against uncertainty at time $t+1$.

This result speaks to an active debate in the business cycle literature regarding the causality between uncertainty and economic outcomes. A robust empirical finding is that aggregate uncertainty is countercyclical, rising when output is low. One view argues that exogenous uncertainty shocks are themselves a key driver of business cycles (Bloom et al., 2018). An alternative view, however, posits that uncertainty is an endogenous mechanism that amplifies and propagates economic shocks (see for example; Van Nieuwerburgh and Veldkamp, 2006; Bachmann and Moscarini, 2011; Taschereau-Dumouchel, 2022). The mechanism in this paper aligns with the latter. A negative shock to fundamentals leads to lower R&D, which in turn causes uncertainty to rise as the flow of information dries up. This endogenous increase in uncertainty is a necessary condition for growth traps to form; it allows pessimistic beliefs to persist and ultimately prolongs the period of low growth. In this model, uncertainty is not the trigger for a downturn, but an endogenous amplifier that gives it its persistence.

3.5 Discussion of Potential Policy Implications

The model presented is highly stylized, and its primary contribution is to demonstrate a theoretical mechanism rather than to make quantitative claims about the real-world prevalence of growth traps. Consequently, deriving specific recommendations is beyond the scope of this paper. However, the model offers interesting and novel perspectives on innovation policy that emerge from its central mechanism: an informational externality that manifests as persistent, belief-driven growth traps.

The potential policy implications of this model are particularly interesting because the informational externality at its core differs from those in standard growth models in two fundamental ways. First, unlike ‘knowledge spillovers’ where past R&D contributes to a public stock of knowledge that makes future innovation cheaper or of higher quality, the spillover here is purely informational. It does not alter the intrinsic potential of an innovation; rather, it helps firms form more precise beliefs about exogenous market conditions, leading to better-informed investment decisions. Second, the welfare cost of this externality manifests differently. Instead of uniformly depressing the economy’s balanced growth rate, it creates the potential for intermittent but severe growth traps – prolonged periods of stagnation where welfare losses are concentrated. These unique features – a spillover affecting beliefs rather than technology, and a welfare cost that is state-dependent and acute – raise novel questions for the design of innovation policy.

Given that the welfare costs are concentrated in discrete periods, policy could be designed to target growth traps directly. State-dependent interventions, such as countercyclical R&D subsidies or direct public investment, could be employed with the primary aim of generating information to lead the economy out of a trap. A successful public project could reveal underlying market conditions, reducing uncertainty regarding the returns to R&D in a specific field and thereby crowding in private investment, with potentially large welfare effects. However, such a strategy faces a critical difficulty: policymakers are subject to the same informational constraints as firms and cannot observe the true state of the economy. This means that state-dependent policies carry the significant risk of subsidizing costly failed investments if the underlying fundamentals are genuinely unfavorable, making the design of such interventions a formidable challenge.

Ultimately, this shared informational constraint is a realistic feature of innovation policy. Policymakers, like firms, must make decisions under ambiguity. Formally modeling the design of optimal policy under such conditions would be an avenue for future research.

3.6 Limitations and Avenues for Future Research

The primary contribution of this paper is to provide a theoretical demonstration of how informational frictions can lead to growth traps. The model is intentionally stylized to isolate this mechanism and, as such, is not calibrated to quantitatively assess the real-world frequency or duration of such traps, nor the precise welfare costs associated with the informational externality. This limitation stems from several simplifying assumptions made for tractability. Among these are the use of an expanding-variety framework for growth and the limited patent duration of a single period, after which the monopoly over a product is randomly reassigned.

The choice of an expanding-variety model, where growth is driven by the creation of new product lines, contrasts with quality-ladder framework with richer firm dynamics, which are standard in many quantitative analyses (e.g.; Acemoglu et al., 2018; Akcigit and Kerr, 2018). In quality-ladder models, firms improve existing products, embodying the concept of "creative destruction" and introducing a negative business-stealing externality. Integrating the learning mechanism presented here into such a framework would be a valuable extension for future research. Similarly, the assumption of a single-period patent duration significantly simplifies the model's numerical solution. A richer model would relax this assumption to describe firms' R&D incentives more accurately over a longer horizon.

Finally, a central limitation lies in the specific exogenous process for the state variable, θ, which affects the returns to R&D. This variable is intended to capture all factors that influence the profitability of innovation but are beyond a firm's control and affect all firms similarly. The assumption of a single, aggregate state variable is a useful simplification for illustrating the mechanism of information-driven growth traps. However, this abstraction also raises important questions that point towards potential extensions and variations of the model.

First, it is difficult to conceptualize a single fundamental that affects all innovating firms across an entire economy in the same way. It is perhaps more intuitive to think of factors that create aggregate uncertainty within a specific sector or for firms using a common technology. For instance, a broad shift in consumer preferences towards environmentally friendly products would alter the innovation possibilities for firms developing green technologies, just as a change in data privacy regulations would reshape the R&D landscape for software companies A natural extension, therefore, would be a multi-sector model where each sector faces its own aggregate uncertainty. Such a framework would allow for an analysis of how inter-sector dynamics might amplify or mitigate the emergence of growth traps. In this light, the model presented in this paper can be interpreted as a representation of firms within a single market, abstracting from the broader interactions across the economy.

Second, the process for θ evolves exogenously, regardless of whether firms invest in R&D. The validity of this assumption depends on what θ is meant to represent. It is plausible for factors like shifts in consumer preferences, which can alter the returns to R&D independently of firm activity. However, it is also conceivable that a feedback loop exists, where successful innovations in a field increase the incentives for subsequent R&D, making the state endogenous to investment. This dynamic is not captured in the current model. A compelling variation could involve a state variable that only evolves when innovation occurs. For example, Callander et al. (2022) build a model where the outcomes of research follow a Brownian motion, and the path is only discovered through the research process itself[11]. Incorporating such a mechanism into an endogenous growth framework to study the macroeconomic effects of the resulting information externalities presents an interesting avenue for future research.

4 Conclusion

This paper extends a standard expanding-variety model with two primary features – aggregate uncertainty about the returns to R&D and social learning from market outcomes – to explore how firms' belief formation can shape long-run economic growth. In the model, firms do not know the true probability distribution of returns from their R&D investments. Instead, they must form subjective beliefs, using the observed outcomes of other firms as a valuable source of information.

This learning mechanism introduces an informational spillover that is distinct from the externalities typically studied in endogenous growth models. When firms invest, they do not internalize the fact that their outcomes reveal valuable information to others. This differs from traditional "knowledge spillovers", where a larger existing knowledge stock allows firms to create better or cheaper innovations. The informational spillover at the center of this model does not enable firms to produce higher-quality goods or innovate at a lower cost; rather, it provides crucial information about market conditions that helps them make better investment decisions. As Fajgelbaum et al. (2017) show for capital investment, this type of spillover can lead to *uncertainty traps* – a self-reinforcing cycle where high uncertainty depresses investment, and low investment in turn increases uncertainty by stifling the flow of new information.

The central finding of this paper is that trap dynamics can also occur in an endogenous growth model, with R&D at its core, even when the effect of uncertainty on investment is, ceteris paribus, positive. While a positive effect of uncertainty would seemingly break the self-reinforcing cycle seen in other models, this analysis demonstrates that traps can emerge nonetheless through a slightly different channel. It is not driven by high uncertainty directly deterring investment; instead, the mechanism is rooted in the persistence of pessimistic beliefs. During a downturn, these beliefs lead to low R&D investment, which stifles the flow of public information. This information scarcity causes beliefs to recover only sluggishly, trapping the economy in a low-growth period even as underlying market conditions improve. Consequently the traps in this growth model are not permanent, eventually beliefs recover. However with R&D investment and innovation as the engine of economic growth, even temporary traps might be associated with large welfare costs.

The model is highly stylized and not intended to quantify the prevalence of growth traps, welfare costs or to derive specific policy recommendations. Its aim is to demonstrate that these trap dynamics can emerge even under conditions where uncertainty encourages investment. This finding is significant because if such traps affect R&D – the primary driver of long term growth – the associated welfare costs could be substantial. Future research could build on this framework by integrating the learning mechanism into richer quantitative models of endogenous growth, which could then be used to assess the real-world existence and magnitude of these growth traps and derive specific policy implications. Ultimately, the analysis underscores that in a world of uncertainty, the mechanisms that govern how

[11] See also Callander (2011) and Callander and Matouschek (2019) for frameworks with a similar logic.

economic agents learn and form beliefs can be as important for long-run growth as the technological fundamentals themselves.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, Reallocation, and Growth. *American Economic Review*, 108(11):3450–3491.

Akcigit, U. and Kerr, W. R. (2018). Growth through Heterogeneous Innovations. *journal of political economy*, page 58.

Anzoategui, D., Comin, D., Gertler, M., and Martinez, J. (2019). Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence. *American Economic Journal: Macroeconomics*, 11(3):67–110.

Arkolakis, C., Papageorgiou, T., and Timoshenko, O. A. (2018). Firm learning and growth. *Review of Economic Dynamics*, 27:146–168.

Bachmann, R. and Moscarini, G. (2011). Business Cycles and Endogenous Uncertainty.

Benigno, G. and Fornaro, L. (2018). Stagnation Traps. *The Review of Economic Studies*, 85(3):1425–1470.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, 86(3):1031–1065.

Callander, S. (2011). Searching and Learning by Trial and Error. *American Economic Review*, 101(6):2277–2308.

Callander, S., Lambert, N., and Matouschek, N. (2022). Innovation and Competition on a Rugged Technological Landscape.

Callander, S. and Matouschek, N. (2019). The Risk of Failure: Trial and Error Learning and Long-Run Performance. *American Economic Journal: Microeconomics*, 11(1):44–78.

Cerra, V., Fatás, A., and Saxena, S. C. (2023). Hysteresis and Business Cycles. *Journal of Economic Literature*, 61(1):181–225.

Fajgelbaum, P. D., Schaal, E., and Taschereau-Dumouchel, M. (2017). Uncertainty Traps. *The Quarterly Journal of Economics*, 132(4):1641–1692.

Gersbach, H., Sorger, G., and Amon, C. (2018). Hierarchical growth: Basic and applied research. *Journal of Economic Dynamics and Control*, 90:434–459.

Jovanovic, B. (1982). Selection and the Evolution of Industry. *Econometrica*, 50(3):649–670.

Klette, T. J. and Kortum, S. (2004). Innovating Firms and Aggregate Innovation. *journal of political economy*, page 33.

Lentz, R. and Mortensen, D. T. (2008). An Empirical Model of Growth Through Product Innovation. *Econometrica*, 76(6):1317–1373.

Queralto, A. (2020). A model of slow recoveries from financial crises. *Journal of Monetary Economics*, 114:1–25.

Syverson, C. (2025). Markups and Markdowns. *Annual Review of Economics*, 17(1):57–76.

Taschereau-Dumouchel, M. (2022). Cascades and Fluctuations in an Economy with an Endogenous Production Network.

Van Nieuwerburgh, S. and Veldkamp, L. (2006). Learning asymmetries in real business cycles. *Journal of Monetary Economics*, 53(4):753–772.

A Derivations

Beliefs and Kalman Filter Derivation This appendix provides a derivation of the equations for the mean (μ_t) and precision (γ_t) of beliefs, which are presented in the main text. The derivation is based on the standard Kalman filter algorithm.

The unobserved state variable, θ_t, evolves according to the following process:

$$\theta_t = \rho\theta_{t-1} + \epsilon_t^\theta, \quad \text{where} \quad \epsilon_t^\theta \sim N(0, (1-\rho^2)\sigma_\theta^2) \tag{16}$$

Firms do not observe θ_t directly but receive a noisy signal, S_t:

$$S_t = \theta_t + \epsilon_t^S, \quad \text{where} \quad \epsilon_t^S \sim N(0, (z_{t-1}\gamma_S)^{-1}) \tag{17}$$

Let $\theta_{k|j} = \mathbb{E}[\theta_k|\mathcal{I}_j]$ denote the expectation of θ_k given the information set $\mathcal{I}_j$ available at time j, and let $P_{k|j} = \text{Var}(\theta_k|\mathcal{I}_j)$ be the corresponding variance.

Kalman Filter Algorithm The filtering process consists of a prediction step and an update step for each period.

Prediction Step

At the beginning of period t, before observing S_t, firms form a prior belief about θ_t based on information up to $t-1$. The mean and variance of this prior distribution are:

$$\theta_{t|t-1} = \rho\theta_{t-1|t-1} \tag{18}$$

$$P_{t|t-1} = \rho^2 P_{t-1|t-1} + (1-\rho^2)\sigma_\theta^2 \tag{19}$$

In the main text, we define the prior mean and precision as $\mu_{t-1} \equiv \theta_{t|t-1}$ and $\gamma_{t-1} \equiv P_{t|t-1}^{-1}$, respectively.

Update Step

After observing the signal S_t, the prior belief is updated to form a posterior belief about θ_t.

The Kalman gain, K_t, which determines the weight assigned to new information, is given by:

$$\begin{aligned} K_t &= \frac{P_{t|t-1}}{P_{t|t-1} + (z_{t-1}\gamma_S)^{-1}} \\ &= \frac{1/\gamma_{t-1}}{1/\gamma_{t-1} + 1/(z_{t-1}\gamma_S)} \\ &= \frac{z_{t-1}\gamma_S}{z_{t-1}\gamma_S + \gamma_{t-1}} \end{aligned}$$

The posterior mean, $\theta_{t|t}$, is a weighted average of the prior mean and the signal:

$$
\begin{aligned}
\theta_{t|t} &= \theta_{t|t-1} + K_t(S_t - S_{t|t-1}) \\
&= \theta_{t|t-1} + K_t(S_t - \theta_{t|t-1}) \\
&= \theta_{t|t-1} + \frac{z_{t-1}\gamma_S}{z_{t-1}\gamma_S + \gamma_{t-1}}(S_t - \theta_{t|t-1}) \\
&= \frac{\gamma_{t-1}\theta_{t|t-1} + z_{t-1}\gamma_S S_t}{z_{t-1}\gamma_S + \gamma_{t-1}}
\end{aligned}
$$

Using the substitution $\mu_{t-1} = \theta_{t|t-1}$, this becomes:

$$
\theta_{t|t} = \frac{\gamma_{t-1}\mu_{t-1} + z_{t-1}\gamma_S S_t}{z_{t-1}\gamma_S + \gamma_{t-1}}
$$

The posterior variance, $P_{t|t}$, is given by:

$$
\begin{aligned}
P_{t|t} &= P_{t|t-1} - K_t P_{t|t-1} \\
&= P_{t|t-1}(1 - K_t) \\
&= \frac{1}{\gamma_{t-1}}(1 - \frac{z_{t-1}\gamma_S}{z_{t-1}\gamma_S + \gamma_{t-1}}) \\
&= \frac{1}{\gamma_{t-1}}\left(\frac{\gamma_{t-1}}{z_{t-1}\gamma_S + \gamma_{t-1}}\right) \\
&= \frac{1}{z_{t-1}\gamma_S + \gamma_{t-1}}
\end{aligned}
$$

Thus, the posterior precision is $P_{t|t}^{-1} = \gamma_{t-1} + z_{t-1}\gamma_S$.

Prediction Step for $t + 1$

In the main text, equation (8) and (9), describe the mean and precision of beliefs about θ in $t + 1$, after observing the information in period t

They are derived by projecting the posterior beliefs at time t $(\theta_{t|t}, P_{t|t})$ forward to time $t + 1$.

$$
\mu_t \equiv \theta_{t+1|t} = \rho\theta_{t|t} = \rho\left(\frac{\gamma_{t-1}\mu_{t-1} + z_{t-1}\gamma_S S_t}{z_{t-1}\gamma_S + \gamma_{t-1}}\right)
$$

This corresponds to equation (8) in the main paper.

$$
\begin{aligned}
\gamma_t \equiv P_{t+1|t}^{-1} &= (\rho^2 P_{t|t} + (1 - \rho^2)\sigma_\theta^2)^{-1} \\
&= \left(\frac{\rho^2}{z_{t-1}\gamma_S + \gamma_{t-1}} + (1 - \rho^2)\sigma_\theta^2\right)^{-1}
\end{aligned}
$$

This corresponds to equation (9) in the main paper.

A.1 Equivalence to Learning Directly from Outcomes

Let $N_t = Z_{t-1}/\eta$ denote the number of firms that start production of new varieties in period t and let the profit of each firm i be $\pi_{i,t} = \exp(\theta_t + u_{it})$, where $u_{i,t} \sim N(0, \sigma_u^2)$ is a firm-specific idiosyncratic shock. Then the logarithm of the profit a firm i makes is a linear signal for the aggregate fundamental θ:

$$
\log(\pi_{i,t}) = \log(\bar{\omega}) + \theta_t + u_{i,t}
$$

Firms observe the outcomes (i.e. profits) of firms directly. Because the noise terms of the individual signals ($u_{i,t}$) are independent, this is equivalent to observing the average log profit which serves as the aggregate signal $\bar{S}_t$:

$$\begin{aligned}\bar{S}_t &\equiv \frac{1}{N_t}\sum_{i=1}^{N_t}\log(\pi_{i,t}) \\ &= \log(\bar{\omega}) + \theta_t + \frac{1}{N_t}\sum_{i=1}^{N_t} u_{i,t},\end{aligned}$$

where $\frac{1}{N_t}\sum_{i=1}^{N_t} u_{i,t}$ is the aggregate noise. The mean of the aggregate noise is $\mathbb{E}[\frac{1}{N_t}\sum_{i=1}^{N_t} u_{i,t}] = 0$, and the variance of the aggregate noise is $\text{Var}(\frac{1}{N_t}\sum_{i=1}^{N_t} u_{i,t}) = \frac{1}{N_t^2}\sum \text{Var}(u_{i,t}) = \frac{1}{N_t^2}(N_t\sigma^2_{u,t}) = \frac{\sigma^2_{u,t}}{N_t}$. The constant term $\log(\bar{\omega})$ does not contain any information and can be dropped. Therefore the aggregate signal from observing the individual firms realized profits is:

$$\bar{S}_t = \theta_t + \varepsilon_t^{\bar{S}}, \qquad \text{with} \quad \varepsilon_t^{\bar{S}} \sim N(0, \frac{\sigma^2_{u,t}}{N_t})$$

The precision of the aggregate signal is increasing in the number of firms that start production, in line with the motivation that more information about current market conditions is shared when more firms enter.

The aggregate signal $\bar{S}$ from observing the individual profits of entering firms directly is equivalent to the aggregate signal S_t in the main text if the variance of the noise term is the same, that is if

$$\begin{aligned}\frac{\sigma^2_{u,t}}{N_t} &= \frac{1}{z_{t-1}\gamma_S} \\ \frac{\sigma^2_{u,t}}{Z_{t-1}/\eta} &= \frac{1}{Z_{t-1}\gamma_S/A_{t-1}} \\ \sigma^2_{u,t} &= \frac{A_{t-1}}{\gamma_S\eta}\end{aligned}$$

Note that the precision of the signal in the main text depends on the R&D intensity $z_{t-1} \equiv \frac{Z_{t-1}}{A_{t-1}}$, which is the R&D expenditure scaled by the productivity level of the economy. This is a reasonable assumption to ensure that as the economy grows over time the amount of information about the economy from new entering firms depends on the number of entering firms relative to the size of the economy instead of the absolute number.

A.2 Derivation of Net Output

The optimal production of intermediate goods of variety j is determined by the demand of the final good producer (equation 2) at the optimal price level $P^* = 1/\alpha$ set by intermediate good producers.

$$x_{j,t} = \alpha^{\frac{2}{1-\alpha}} q_{j,t} L_t$$

Substituting this into the final good production function (equation 1) results in the optimal final good production

$$Y_t = L_t \alpha^{\frac{2\alpha}{1-\alpha}} \sum_{j=0}^{M_t} q_{jt}$$

The intermediate varieties $x_{j,t}$ are produced at a marginal cost of one unit of the final good. Therefore the net output, GDP, is the total final good production Y_t minus the cost of producing the intermediate varieties:

$$\begin{aligned}
\text{GDP}_t &= Y_t - \sum_{j=0}^{M_t} x_{j,t} \\
&= L_t \underbrace{\left(\alpha^{\frac{2\alpha}{1-\alpha}} - \alpha^{\frac{2}{1-\alpha}}\right)}_{\equiv \Psi} \sum_{j=0}^{M_t} q_{jt} \\
&= \Psi L_t \sum_{j=0}^{M_t} q_{jt} \\
&= \Psi L_t A_t,
\end{aligned}$$

A.3 Expenditure Share of Intermediate Goods

The final good producer acts as a price-taker in the market for intermediate goods. The firm's profit-maximization problem implies that it will demand each intermediate good up to the point where its marginal product equals its price, $P_{j,t}$:

$$\frac{\partial Y_t}{\partial x_{j,t}} = P_{j,t}$$

Given the final good production function (equation 1), the marginal product is:

$$\frac{\partial Y_t}{\partial x_{j,t}} = \alpha L_t^{1-\alpha} q_{j,t}^{1-\alpha} x_{j,t}^{\alpha-1}$$

The total expenditure on all intermediate products is $\sum_{j=0}^{M_t} P_{j,t} x_{j,t}$, where we can substitute the expression for the marginal product in for the price:

$$\begin{aligned}
\sum_{j=0}^{M_t} P_{j,t} x_{j,t} &= \sum_{j=0}^{M_t} \left(\alpha L_t^{1-\alpha} q_{j,t}^{1-\alpha} x_{j,t}^{\alpha-1}\right) x_{j,t} \\
&= \alpha \left(L_t^{1-\alpha} \sum_{j=0}^{M_t} q_{j,t}^{1-\alpha} x_{j,t}^{\alpha}\right) \\
&= \alpha Y_t
\end{aligned}$$

Thus, the total expenditure on intermediate goods is a constant share α of final output.

A.4 Derivation of the Labor Share

This appendix derives the labor share of national income (GDP) within the model's framework. The labor share is defined as the total wage bill divided by the economy's net output (GDP).

The Wage Bill

The final good, Y_t, is produced by a perfectly competitive firm with the production function:

$$Y_t = L_t^{1-\alpha} \sum_{j=0}^{M_t} q_{j,t}^{1-\alpha} x_{j,t}^{\alpha}$$

In a competitive market, factors of production are paid their marginal product. The marginal product of labor (MPL) is the partial derivative of the production function with respect to labor, L_t:

$$\frac{\partial Y_t}{\partial L_t} = (1-\alpha) L_t^{-\alpha} \sum_{j=0}^{M_t} q_{j,t}^{1-\alpha} x_{j,t}^{\alpha}$$

We can rewrite this by substituting in the definition of Y_t:

$$\frac{\partial Y_t}{\partial L_t} = (1-\alpha) \frac{Y_t}{L_t}$$

The total wage bill is the wage rate (equal to the MPL) multiplied by the total labor supply. As labor supply is normalized to one ($L_t = 1$), the wage bill is:

$$W_t = 1 \cdot \frac{\partial Y_t}{\partial L_t} = (1-\alpha) Y_t$$

Net Output (GDP)

Net output, or GDP, is defined as the value of final good production, Y_t, minus the cost of the intermediate goods used in its production. Since one unit of any intermediate good costs one unit of the final good to produce, the cost is the sum of the quantities of all intermediate goods, $\sum x_{j,t}$.

$$\text{GDP}_t = Y_t - \sum_{j=0}^{M_t} x_{j,t}$$

To find the total quantity of intermediate goods $\sum_{j=0}^{M_t} x_{j,t}$ we can use the fact that the total expenditure on intermediate goods is equal to:

$$\sum_{j=0}^{M_t} P_{j,t} x_{j,t} = \alpha Y_t,$$

as shown in Appendix A.3. Each monopolistic intermediate good producer sets a price $P_{j,t} = 1/\alpha$. Substituting this into the expenditure equation gives:

$$\sum_{j=0}^{M_t} \frac{1}{\alpha} x_{j,t} = \alpha Y_t$$

Solving for the total quantity of intermediates yields:

$$\sum_{j=0}^{M_t} x_{j,t} = \alpha^2 Y_t$$

Substituting this back into the definition of GDP, we get:

$$\text{GDP}_t = Y_t - \alpha^2 Y_t = (1-\alpha^2) Y_t$$

Labor Share of Income

The labor share of income is the total wage bill divided by the economy's net output (GDP):

$$\text{Labor Share} = \frac{(1-\alpha)Y_t}{(1-\alpha^2)Y_t} = \frac{1}{1+\alpha}$$

B Sensitivity to Parameter Values

This section assesses the robustness of the paper's main findings (section 3.2) to variations in the parameter values. The central result from the main analysis is that the informational frictions in the learning model systematically reduce long-run growth compared to the full information benchmark. To test whether and how this result depends on the specific parameterization, a sensitivity analysis is performed. Each key parameter is varied individually across a range of values, while all other parameters are held at their baseline levels. For each new parameter value, the economy is simulated 10,000 times over 150 periods for both the full information and the learning specifications.

The analysis confirms two key findings. First, the growth rate in the full information benchmark is consistently higher than the growth rate in the learning model and this gap is robust across a wide range of parameter values (Figure 7). Second, this difference in growth rates is consistent with the emergence of growth traps, which occur in a high percentage of simulations for most parameter choices (Figure 8). The remainder of this section discusses these findings in more detail.

Figure 7 shows the average growth rate for the two model specifications (i) the full information benchmark (blue) and the (ii) model where the state is unobserved and firms have to learn from market outcomes (orange). The growth rate in the learning model is consistently lower, though the quantitative magnitude of this growth gap varies, depending on the parameter values:

- γ_S: The scale parameter γ_S determines how effectively R&D intensity translates into public information. High values of γ_S mean that even a small amount of R&D generates a more precise signal, making social learning more effective. As shown in the top-left panel of Figure 7, as γ_S increases, the performance of the learning economy converges improves, causing the growth gap to shrink.

- ρ: The persistence of the fundamental state, ρ, also has a significant effect on the growth gap. A higher persistence allows for more extreme realizations of θ. Due to the convexity of the quality function ($q_t = \exp \theta_t$), the upside potential from very high values of θ outweighs the downside, where growth is bounded by zero. As a result, the average growth rate in the full information benchmark is increasing in ρ. In the learning model, however, the average growth rate seems to be largely independent of the persistence (for relatively persistent processes with $\rho \in (0.6, 0.95)$ as shown here). This suggests that the potential growth benefits of a more persistent process go unrealized, possibly because firms are more likely to be stuck in a growth trap missing the opportunity to innovate when the returns are highest.

- σ_θ: Similarly, increasing σ_θ, leads a mean-preserving spread of the distribution of θ conditional on beliefs. This is because σ_θ governs the variance of the shocks to the fundamental, a source of uncertainty that is common to both specifications. Given the convexity of the quality (and profits) this incentivizes R&D investment and therefore the growth rate is increasing in σ_θ in both cases.

- η: The cost of R&D, η, has a more complex effect on the growth gap. As expected, higher costs reduce growth in both models. For low values of η the gap between the full information and the learning model is very small, likely because the incentive for investment is very high leading to large investment and therefore precise information in the learning model. As the cost of R&D increases

Figure 7: Sensitivity of Average Growth to Key Parameters

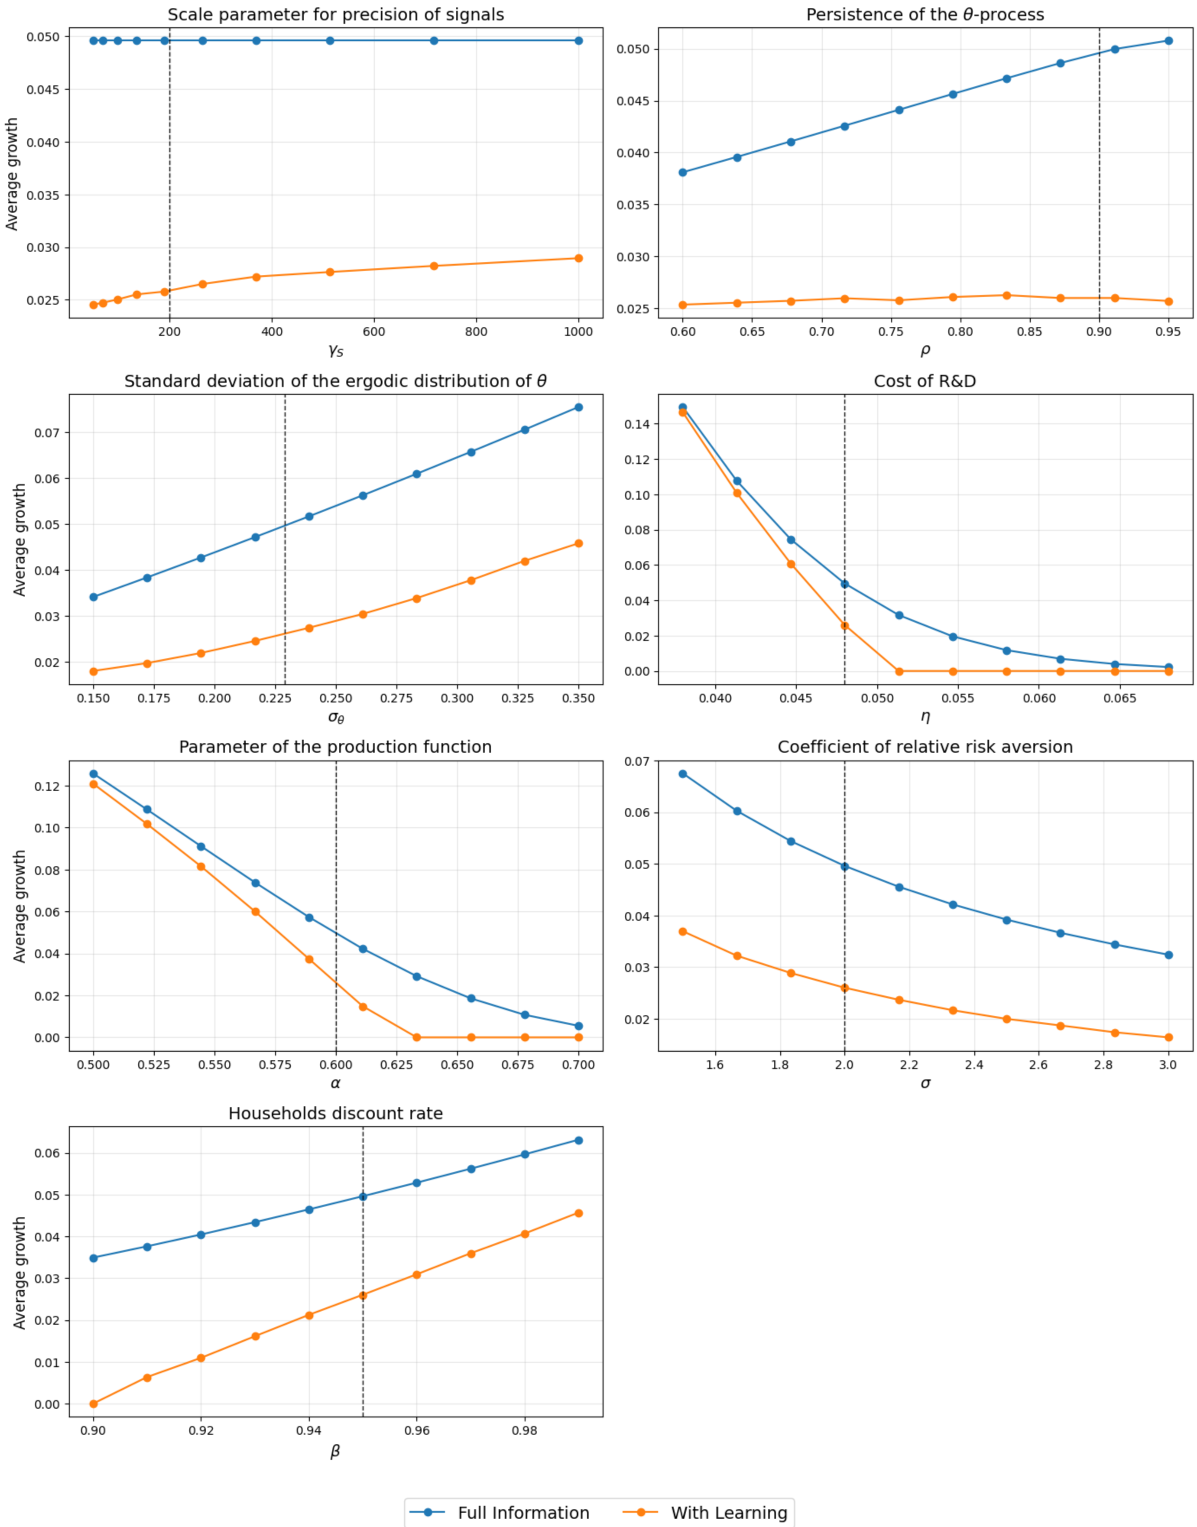

Notes: This figure shows the results of the sensitivity analysis. Each panel displays the average growth rate of the economy for both (i) the full information benchmark (blue) and the model specification with learning from market outcomes (orange) across a range of values for a single parameter, holding all other parameters at their baseline values. Each point is the average outcome of 10,000 simulations over 150 periods each. The vertical dashed line in each panel indicates the baseline parameter value used in the main analysis.

and investment incentives are reduced, the gap increases. In the learning model there is a "kink" for values of η higher than a certain threshold. This "kink" occurs because a high R&D cost can lead to a permanent trap. When η is sufficiently high, such that the unconditional belief that firms revert to in the absence of information is no longer optimistic enough to justify investment, the economy will be in a permanent trap. In contrast, firms in the full information benchmark still invest whenever θ_t is high enough, ensuring positive average growth even for higher costs.

- α: There is a similar pattern regarding the parameter α in the production function (equation 1). The parameter governs the markup of firms and their profits, with higher values of α implying lower markups. The growth rate is therefore decreasing in α as it reduces the incentives to invest in both model specifications, and similarly to the cost of R&D there is a certain threshold at which the incentives for firms in the learning model are too low to justify any investment when beliefs revert to the unconditional mean.

- σ, β: Finally, the parameters governing household preferences affect the incentives for R&D through the stochastic discount factor[12]. A higher coefficient of relative risk aversion, σ, makes households less willing to tolerate consumption volatility, which disincentivizes risky R&D investment. As a result, the growth rate is decreasing in σ in both specifications. The growth gap, however, persists and appears relatively constant in proportional terms. In contrast, a higher household discount factor, β, implies that households are more patient and willing to save for the future. This encourages investment and increases the growth rate in both models. As β increases, the growth gap narrows, as the stronger fundamental incentive to invest makes the informational friction less binding, but the gap remains significant even for very patient households.

As a final robustness check, Figure 8 directly examines the frequency of growth traps across the parameter space. This serves to confirm that the robust difference in growth rates is indeed driven by the emergence of these traps. The percentage of simulations in which a growth trap emerges is very high across most parameter ranges, confirming that the trap mechanism is a robust feature of the model. The results also align with the patterns observed for the difference in growth rates previously. For example, for very low R&D costs (η), the frequency of traps falls sharply, which is consistent with the near-disappearance of the growth gap under these same conditions. It is also worth noting that the relationship between some parameters and trap frequency is non-monotonic. This can occur for parameter values where the growth rate in the full information benchmark itself is very low. In these cases, there are fewer opportunities for a trap to form as defined, because the benchmark level of investment is already zero for a larger portion of the time.

It is important to place these robustness results in the proper context. While the analysis shows that the emergence of growth traps is not an artifact of a specific parameter choice, this does not imply that all parameter combinations will lead to such traps, nor does it make a quantitative claim about their prevalence in reality. The primary contribution of this paper is to demonstrate, within a stylized theoretical model, that growth traps can emerge from a purely informational friction, even when uncertainty does not deter R&D investment in the way it might affect capital investment. Given the potentially large welfare costs of prolonged stagnation in innovation, establishing the theoretical possibility of this mechanism is a crucial first step. It highlights the importance of social learning and belief dynamics in endogenous growth and provides a foundation for future research to build richer models that can be formally calibrated to assess the quantitative importance of these information-based traps.

[12]The stochastic discount factor used in the firms optimization problem reflects household preferences and is defined as $M_{t,t+1} = \beta \frac{u'(C_{t+1})}{u'(C_t)}$, and with CRRA utility it is $M_{t,t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\sigma}$

Figure 8: Sensitivity of Growth Trap Incidence to Key Parameters

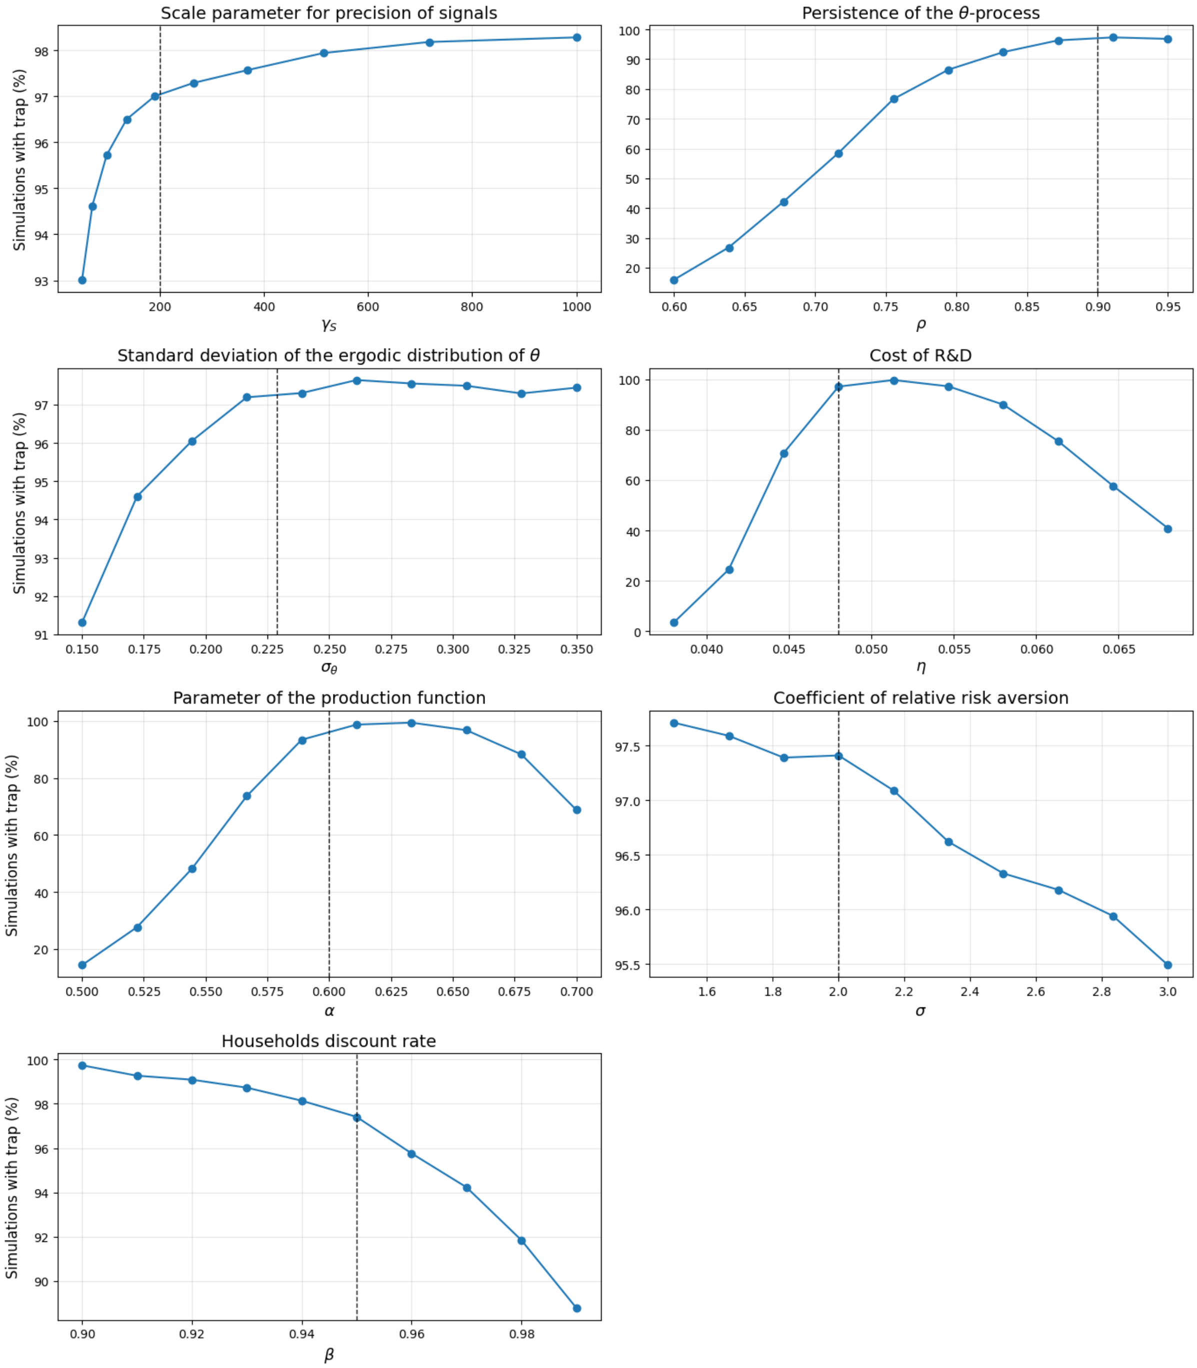

Notes: This figure shows the percentage of simulations in which a growth trap occurs. Each panel displays this across a range of values for a single parameter, holding all other parameters at their baseline values. A trap is defined as a sequence of at least five consecutive periods where the investment rate in the learning economy is less than 20% of the investment rate in the full information benchmark. Each point is calculated from 10,000 simulations over 150 periods. The vertical dashed line in each panel indicates the baseline parameter value.