ISSN 0956-8549

The Option Value of Waiting for Institutional Improvement: Real Options and Growth with Endogenous Beliefs

By

David C Webb

FINANCIAL MARKETS GROUP DISCUSSION PAPER NO. 976

August 2026

The Option Value of Waiting for Institutional Improvement:
Real Options and Growth with Endogenous Beliefs

David C Webb
London School of Economics
d.c.webb@lse.ac.uk

April 6, 2026

Abstract

This paper explains why institutional reforms often fail to stimulate investment despite improving long-run fundamentals. We develop a real options framework where uncertainty operates through three compounding channels: direct institutional volatility, the option value of waiting for conditions to improve, and uncertainty about whether announced reforms will actually succeed. This third channel, belief uncertainty about government credibility, amplifies traditional wait-and-see effects by raising the characteristic root β^* that governs the option multiplier, compounding the standard real-options premium multiplicatively. Even moderate uncertainty about reform success raises investment hurdle rates by 7–15 percentage points, while crisis-level uncertainty can nearly double required returns. We identify a "zone of maximum paralysis" at intermediate institutional quality where these combined effects peak, and show that low-credibility governments face a paradox: ambitious announcements deepen paralysis rather than stimulating growth. Calibrating to historical episodes (Poland 1990–91, UK 2016–20), the model matches observed investment collapses of 30–60%. The framework explains why gradualism often dominates "big bangs" and why external anchors like IMF programmes succeed, they reduce belief uncertainty more than they improve fundamentals.

JEL Codes: D81, E22, O43, D83, E61.

Keywords: Real options; Institutional reform; Belief uncertainty; Investment irreversibility; Government credibility.

1 Introduction

In July 1990, Poland announced the most radical institutional reforms in post-communist Europe. Despite widespread agreement that the reforms were necessary and would ultimately succeed, investment collapsed by 35% in the subsequent year. Firms explicitly cited "waiting for the new rules to become clear" as their reason for delay. This pattern, reform announcement followed by investment paralysis, has repeated across dozens of countries: Mexico during NAFTA negotiations (1991–1994), Egypt during its multiple reform attempts (2004–2020), and India during its 2016 demonetisation and GST reforms. Between 2010 and 2020, countries announcing major structural reforms experienced an average 2.8 percentage point decline in investment rates in the subsequent year, even when international observers rated the reforms as credible and necessary.

This paper provides a theoretical explanation for why institutional reform, even reform that will ultimately succeed—can paradoxically suppress investment and growth. Building on Bernanke [1983]'s seminal insight that uncertainty creates a value of waiting when investment is irreversible, we extend the real options framework of McDonald and Siegel [1986] and Dixit and Pindyck [1994] to institutional uncertainty. Our key insight is that reform uncertainty operates through three compounding channels: (i) conventional uncertainty about current institutional quality, (ii) an option value of waiting for institutions to improve, and (iii) uncertainty about whether announced reforms will actually materialise. This third channel, which depends critically on government credibility, can dominate the other two and explains why well-intentioned reforms often fail to stimulate immediate economic activity.

A distinctive feature of our framework is that institutional quality follows a bounded process, confined to $[0, \bar{\theta}]$ by a Jacobi diffusion whose volatility vanishes at both endpoints [Karlin and Taylor, 1981]. This boundedness is economically essential, not a technical convenience. Institutions cannot deteriorate indefinitely, even the most dysfunctional state retains minimal capacity for contract enforcement, nor improve without limit, since $\bar{\theta}$ represents the frontier of feasible institutional quality given a country's underlying constraints. More importantly, boundedness generates the paper's central non-linearity: because the diffusion coefficient $\nu\sqrt{\theta(\bar{\theta} - \theta)}$ peaks at intermediate institutional quality, the option value of waiting is largest precisely in the middle of the institutional range, where both deterioration and improvement remain live possibilities. This is the "zone of maximum paralysis." An unbounded specification, standard in the real options literature, would miss this hump-shaped pattern entirely.

Our theoretical framework yields four main contributions. First, we show that the investment threshold under institutional uncertainty exceeds the break-even level by a factor of

$\frac{\beta^*}{\beta^*-1} \times \left(1 + \frac{\sigma_\alpha^2}{\nu^2}\gamma\right)$, where the first term, governed by β^*, captures the traditional real option value and the second, scaled by σ_α^2, captures belief uncertainty about reform success. This multiplicative structure means that even modest σ_α^2 can create prohibitive investment hurdles. Second, we demonstrate that government credibility fundamentally alters the effect of reform announcements. Using a Bayesian learning framework, we show that low-credibility governments face a dilemma: ambitious announcements increase σ_α^2 more than they raise $\hat{\alpha}$, making bold reforms counterproductive. We derive an optimal announcement rule exhibiting striking non-monotonicity in $\hat{\alpha}$, low-credibility governments should barely deviate from market priors regardless of their true intentions, while high-credibility governments can afford to be aspirational. Third, we identify a "compound precautionary effect" that can trap economies in low-investment equilibria even when everyone expects institutional improvement, through a feedback loop in which waiting validates itself by making reforms appear to fail. Fourth, the framework generates three sharp, testable predictions—a credibility-dependent sign reversal in investment responses, an irreversibility gradient across sectors, and a discontinuous pattern of investment recovery, that constitute a coherent empirical agenda distinguishing our mechanism from existing alternatives.

A methodological note on Sections 6 and 6.2: the historical episodes considered there (Poland 1990–91 and the post-Brexit period 2016–20) are *illustrative simulations*, not empirical validations. Each maps the model's structural parameters to historically documented values to demonstrate that the framework can generate quantitative magnitudes consistent with observed investment behaviour. Formal identification, isolating the credibility mechanism from credit shocks, demand effects, and political instability is left to the empirical agenda outlined in Section 5.

The policy implications are stark. Reducing the volatility of the reform path, through external anchors, gradual implementation, or credibility-building institutions may be more important than the reforms themselves. The model explains why IMF programmes and EU accession processes succeed where similar domestic reforms fail, and why low-credibility governments should under-promise and over-deliver rather than announce ambitious targets.

The remainder of the paper proceeds as follows. Section 2 presents the central insight. Section 3 develops the formal framework. Section 4 analyses growth implications. Section 5 presents testable predictions and an empirical agenda. Section 6 examines policy implications, with Section 6.2 applying the framework to historical episodes. Section 7 concludes. A complete notation list is provided in Appendix C.

2 The Central Insight

Standard growth models often struggle to explain why investment remains depressed even when institutional reforms are on the horizon. This paper proposes that poor current institutions, when paired with the potential for improvement, create a compound precautionary motive operating through three channels.

The first is direct: high institutional volatility $\sigma^2(\theta_{\text{low}})$ means that committing capital immediately carries substantial risk of further deterioration before improvement arrives. The second is more subtle: because institutions may improve, waiting acquires a positive option value, a firm that delays can deploy capital more productively in a better institutional environment tomorrow. These two channels are familiar from the real options literature.

The critical innovation is a third channel that interacts with and amplifies the first two. Markets cannot directly observe whether announced reforms will succeed; they must infer the true reform parameter $\tilde{\alpha}$ from the observed evolution of θ_t. This inference problem creates an additional layer of option value qualitatively distinct from the first two: firms are waiting not just for institutions to improve, nor just for uncertainty about the institutional path to resolve, but for uncertainty about the reform process itself, about whether the government can and will deliver what it has promised, to become clearer. When government credibility is low, this third channel can dominate the other two, making even well-intentioned and ultimately successful reforms counterproductive in the short run.

These three channels raise the cost of capital from the risk-free rate r to an effective rate:

$$r_{\text{eff}} = r + \underbrace{\pi(\theta)}_{\text{risk premium}} + \underbrace{\Omega(\theta, \nu)}_{\text{option value of waiting}} + \underbrace{\Psi(\sigma^2_{\alpha,t})}_{\text{belief uncertainty premium}} \tag{1}$$

The additive form of this decomposition might suggest the three premia simply stack. In terms of the required premia themselves, this is correct. But the decomposition understates the true interaction, because belief uncertainty and the standard option value of waiting compound multiplicatively at the level of the investment multiplier. Because Ψ operates through the characteristic root β^*, which governs the entire multiplier $\beta^*/(\beta^* - 1)$, a rise in $\sigma^2_{\alpha,t}$ scales up an option value already elevated by institutional volatility, rather than merely adding to it. This compounding generates non-linearities, credibility traps, and sign reversals in investment responses to reform announcements that standard real-options models cannot produce, as Propositions 1 and 3 make precise.

2.1 The Compound Precautionary Effect

The dominant quantitative effect operates not through the multiplier itself but through the break-even threshold $\theta_{NPV=0} \propto (r^*)^{1/(1-\eta)}$: as r^* rises with belief uncertainty, the minimum institutional quality at which any investment is economically viable rises independently of and prior to the application of the multiplier. Even moderate belief uncertainty can therefore produce substantial increases in the investment threshold, as Table 1 illustrates. The full multiplicative interaction between the two convexity channels, $V_{\theta\theta} > 0$ and $V_{\alpha\alpha} > 0$, is formalised via the Feynman–Kac representation in Appendix A.3.

3 Formal Framework

3.1 The Investment Problem

A firm considers whether to sink an irreversible cost I to install capital k that generates flow output $A(\theta_t)k^\eta$ per unit time, where θ_t is current institutional quality. The decision is a stopping problem: the firm chooses a stopping time τ to maximise the expected net present value of investment:

$$\sup_{\tau} \mathbb{E}\big[e^{-r\tau}\big(V(\theta_\tau) - I\big)\big] \tag{2}$$

where

$$V(\theta) = \frac{A(\theta)k^\eta}{r - g(\hat{\alpha}, \theta)} \tag{3}$$

is the value of an installed project and $g(\hat{\alpha}, \theta) \equiv \hat{\alpha}(\bar{\theta}/\theta - 1)$ is the institutional drift rate, increasing in the gap $(\bar{\theta} - \theta)$ and vanishing as $\theta \to \bar{\theta}$. We assume $r > g(\hat{\alpha}, \theta)$ throughout. Agents maintain a belief distribution $\tilde{\alpha} \sim \mathcal{N}(\hat{\alpha}_t, \sigma^2_{\alpha,t})$ over the unobserved reform parameter $\tilde{\alpha}$, updated continuously as the path $\{\theta_s\}_{s \leq t}$ is observed; the formal learning dynamics are developed in Section 3.4. Capital k is treated as fixed at the investment date; the sole decision is when to invest.

The optimal policy takes the form of a threshold rule: invest at the first passage time

$$\tau^* = \inf\{t \geq 0 : \theta_t \geq \theta^*\}. \tag{4}$$

Sections 3.2–3.3 characterise the optimal threshold θ^* and show how belief uncertainty $\sigma^2_{\alpha,t}$ raises it above the certainty benchmark.

Definition 1 (Investment Metrics)**.** *The* investment threshold θ^* *is the minimum institutional quality at which firms initiate investment. The* break-even threshold $\theta_{NPV=0}$ *is the institutional quality at which net present value equals zero under certainty:* $V(\theta_{NPV=0}) = I$.

The investment multiplier *is the ratio $\theta^*/\theta_{NPV=0}$, also expressed as $\beta^*/(\beta^* - 1)$ in the real options framework. The* required premium *is*

$$\textit{Investment Premium} = \left(\frac{\theta^*}{\theta_{NPV=0}} - 1\right) \times 100\% = \left(\frac{\beta^*}{\beta^* - 1} - 1\right) \times 100\%. \tag{5}$$

An investment multiplier of 2.0 means firms require institutional quality to be twice the break-even level.

Definition 2 (Government Credibility)**.** *The government is either* competent*, able to implement announced reforms, or* incompetent*, unable to do so. Government type is unobservable by the private sector. We define* credibility $\rho_t \in [0, 1]$ *as the private sector's posterior probability at time t that the government is competent, given the observed history of institutional outcomes $\{\theta_s\}_{s \le t}$.*

3.2 Institutional Dynamics: The Jacobi Diffusion and Belief Formation

We model institutional quality as a bounded mean-reverting process, reflecting the natural constraints that institutions cannot fall below a minimum level of dysfunction nor exceed some ideal ceiling $\bar{\theta}$.

Why the Jacobi diffusion? The choice of a Jacobi (Wright–Fisher) diffusion is not a technical convenience; it captures two economically essential features. First, the diffusion coefficient $\nu\sqrt{\theta(\bar{\theta} - \theta)}$ vanishes at both boundaries, so that the process is automatically confined to $[0, \bar{\theta}]$ without requiring reflecting or absorbing boundary conditions to be imposed *ad hoc*. This reflects the economic reality that institutional deterioration and improvement are both self-limiting: even the most dysfunctional state retains some minimal enforcement capacity, and no country can exceed a feasible institutional frontier $\bar{\theta}$ determined by its underlying constraints. Second, the diffusion coefficient is hump-shaped in θ, it peaks at the midpoint $\bar{\theta}/2$ and falls to zero at both ends. This generates the paper's central nonlinearity: the option value of waiting, which scales with local variance, is largest precisely at intermediate institutional quality where both deterioration and further improvement remain live possibilities. An unbounded specification, such as arithmetic or geometric Brownian motion, imposes constant or monotone variance and would miss this hump-shaped pattern entirely. The Jacobi diffusion is the natural process for a state variable known to live on a compact interval [Karlin and Taylor, 1981].

The true process follows:[1]

$$d\theta_t = \tilde{\alpha}(\bar{\theta} - \theta_t)\,dt + \nu\sqrt{\theta_t(\bar{\theta} - \theta_t)}\,dB_t \tag{6}$$

Agents observe the path $\{\theta_t\}_{t\geq 0}$ continuously but cannot directly observe the true reform parameter $\tilde{\alpha}$. They maintain a prior $\tilde{\alpha} \sim \mathcal{N}(\alpha_0, \sigma_\alpha^2)$ and update to a posterior $\tilde{\alpha} \sim \mathcal{N}(\hat{\alpha}_t, \sigma_{\alpha,t}^2)$ at time t, where $\hat{\alpha}_t$ denotes the posterior mean and $\sigma_{\alpha,t}^2$ the posterior variance. The posterior mean evolves optimally via the continuous-time Kalman filter [Liptser and Shiryaev, 2001, Øksendal, 2003]:

$$d\hat{\alpha}_t = \frac{\sigma_{\alpha,t}^2}{\nu^2\theta_t(\bar{\theta} - \theta_t)}\Big[d\theta_t - \hat{\alpha}_t(\bar{\theta} - \theta_t)\,dt\Big] \tag{7}$$

This framework builds on Décamps and Mariotti [2004], though we focus on learning about government credibility rather than project quality.

Calibration of λ. The uncertainty-adjusted discount rate derived below takes the form $r^* = r + \lambda\sigma_\alpha^2$, where

$$\lambda = \frac{(\bar{\theta} - \theta_0)^2}{2\theta_0\bar{\theta}} \tag{8}$$

is the institutional gap parameter. Three properties make this calibration transparent. First, λ is dimensionally consistent with σ_α^2, so $\lambda\sigma_\alpha^2$ adds directly to r. Second, $\lambda \to 0$ as $\theta_0 \to \bar{\theta}$: belief uncertainty becomes irrelevant when institutions are already at their frontier. Third, λ is monotone increasing in the fractional gap $(\bar{\theta} - \theta_0)/\bar{\theta}$—uncertainty about reform speed matters most when institutions are far from steady state. Crucially, λ is not a free parameter: it is determined entirely by the observable ratio $\theta_0/\bar{\theta}$. For the baseline calibration $\theta_0 = 0.60\bar{\theta}$, $\lambda = (0.40)^2/(2\times 0.60) = 0.133$. World Bank Governance Indicator data place most transition economies in the fractional-gap range 0.30–0.50, making $\lambda \in [0.075, 0.208]$ the empirically relevant range and the baseline value conservative.

Structurally, λ is a first-order approximation to the local curvature of the belief-value mapping at θ_0, derived in Appendix A.1 from a second-order Taylor expansion of $V(\theta, \tilde{\alpha})$ around $\hat{\alpha}$; it is therefore exact to order $O(\sigma_\alpha^2)$ and inherits no additional error beyond the

[1] Our framework shares the methodological approach of Dixit and Pindyck [1994], Chapter 5, positing a power function solution in the waiting region, deriving a characteristic equation, and applying value-matching and smooth-pasting conditions to determine the optimal threshold. However, the two settings are mathematically distinct. Dixit and Pindyck employ mean-reverting processes with diffusion coefficients that are either constant or proportional to the state variable, yielding support on $(0, \infty)$. Our institutional quality process uses a Jacobi (Wright–Fisher) diffusion with term $\nu\sqrt{\theta(\bar{\theta} - \theta)}$, which vanishes at both boundaries and confines θ to $[0, \bar{\theta}]$. The resulting HJB equation differs from theirs, and their specific functional-form results do not transfer directly. What carries over is the general solution methodology, applied here via the local approximation described in Appendix A.1.

certainty-equivalent approximation established in Lemma 1. Alternative specifications sharing the same boundary behaviour—vanishing as $\theta_0 \to \bar{\theta}$ and increasing in the institutional gap—yield identical comparative statics, since it is these two properties, not the precise functional form, that drive the paper's results.

3.3 The Investment Threshold with Belief Uncertainty

State space reduction: from 2D to 1D. The joint dynamics of $(\theta_t, \hat{\alpha}_t)$ under the filtered probability measure exhibit instantaneous covariance $d\langle\theta, \hat{\alpha}\rangle_t = \sigma^2_{\alpha,t}\,dt$, as established in Appendix B, equation (B.3). A fully general treatment would therefore require solving a partial differential equation in the joint state space $(\theta, \hat{\alpha})$, including the cross-derivative term $\sigma^2_\alpha(\bar{\theta} - \theta)W_{\theta\hat{\alpha}}$ arising from this covariance. The following lemma establishes that this reduction to the one-dimensional state θ is rigorous at the order of approximation maintained throughout.

Lemma 1 (State Space Reduction)**.** *Let $W(\theta, \hat{\alpha}; \varepsilon)$ denote the value of waiting, where $\varepsilon \equiv \sigma^2_\alpha$. Under a perturbation expansion in ε (detailed in Appendix B), the cross-derivative term $\sigma^2_\alpha(\bar{\theta} - \theta)W_{\theta\hat{\alpha}}$ and the second-order belief term $W_{\hat{\alpha}\hat{\alpha}}$ enter the investment threshold only at order $O(\sigma^4_\alpha)$. To first order in σ^2_α, the value of waiting satisfies the one-dimensional HJB equation*

$$r^*W = \hat{\alpha}(\bar{\theta} - \theta)W_\theta + \tfrac{1}{2}\nu^2\theta(\bar{\theta} - \theta)W_{\theta\theta}, \qquad r^* = r + \lambda\sigma^2_\alpha, \tag{9}$$

with belief uncertainty entering exclusively through the adjusted discount rate r^*. The omitted $O(\sigma^4_\alpha)$ terms are established in Appendix B to be negligible across the empirically relevant parameter ranges of Table A.1.

Proof. See Appendix B, Sections B.1–B.3. The key steps are: (i) the full two-dimensional generator of $(W_{\theta\hat{\alpha}})$ is written out explicitly in equation (B.4); (ii) a perturbation expansion $W = W^{(0)} + \varepsilon W^{(1)} + \varepsilon^2 W^{(2)} + O(\varepsilon^3)$ is substituted into the HJB; (iii) at order ε^0 the standard one-dimensional equation is recovered with root β_0; (iv) at order ε^1 the cross-derivative and $W_{\hat{\alpha}\hat{\alpha}}$ terms are zero because $W^{(0)}$ depends only on θ, and the correction $W^{(1)}$ satisfies the modified equation (B.7) with adjusted root β^*; (v) the two-dimensional terms enter only at order $\varepsilon^2 = O(\sigma^4_\alpha)$. □

To incorporate belief uncertainty while maintaining tractability, we adopt the certainty-equivalent approach justified by Lemma 1, where uncertainty about reform success affects investment through the augmented discount rate $r^* = r + \lambda\sigma^2_\alpha$.[2]

[2] This certainty-equivalent representation avoids solving the full two-dimensional HJB by projecting onto

Intuition: why belief uncertainty amplifies hurdles. Consider a typical infrastructure project in a reforming economy. Under certainty about institutional evolution, irreversibility alone might require a 40% return premium—the standard real options effect captured by $\beta/(\beta-1)$. When investors cannot determine whether the government's reform promises will materialise, an additional layer of option value emerges: the value of waiting to learn whether reforms will actually succeed.

At the level of the investment multiplier, the two sources of uncertainty interact multiplicatively. For small belief uncertainty, a first-order expansion of $\beta^*/(\beta^*-1)$ around r yields:

$$\text{Multiplier} \approx \underbrace{\left(1+\frac{\tilde{\nu}^2}{2r}\right)}_{\text{base real option}} \times \underbrace{\left(1+\frac{\lambda\sigma_\alpha^2}{2r}\right)}_{\text{belief amplification}} \tag{10}$$

where $\tilde{\nu}$ is the locally approximated diffusion coefficient evaluated at a representative interior threshold $\theta^* \in (0,\bar{\theta})$, defined formally in Appendix A.1, Step 3.

Before stating the main result, let $A(\theta)$ denote the productivity of installed capital at institutional quality θ, where $A : [0,\bar{\theta}] \to \mathbb{R}_+$ is strictly increasing and continuously differentiable, with $A(0) = 0$ and $A(\bar{\theta}) = \bar{A} < \infty$. Flow output is $A(\theta)k^\eta$ with $\eta \in (0,1)$ the capital share.

Proposition 1 (Investment Threshold with Belief Uncertainty)**.** *Under belief uncertainty about the reform parameter $\tilde{\alpha}$, and given the state space reduction of Lemma 1, the optimal investment threshold is:*

$$\theta^* = \theta_{NPV=0} \cdot \frac{\beta^*}{\beta^*-1} \tag{11}$$

where

$$\theta_{NPV=0} = \left(\frac{rI}{A(\theta^*)k^\eta}\right)^{1/(1-\eta)} \tag{12}$$

is the break-even threshold at which net present value equals zero under certainty. The characteristic root β^* solves the modified characteristic equation

$$\tfrac{1}{2}\tilde{\nu}^2\beta(\beta-1) + \tilde{\alpha}\beta - r^* = 0 \tag{13}$$

with uncertainty-adjusted discount rate $r^* = r + \lambda\sigma_\alpha^2$, yielding positive root:

$$\beta^* = \frac{1}{2} - \frac{\hat{\alpha}}{2\nu^2} + \sqrt{\left(\frac{1}{2}-\frac{\hat{\alpha}}{2\nu^2}\right)^2 + \frac{2r^*}{\nu^2}} \tag{14}$$

---

the one-dimensional state θ. It is standard in the incomplete-information investment literature [Gennotte, 1986, Brennan, 1998] and, as Lemma 1 confirms, is exact to first order in σ_α^2. The perturbation expansion in Appendix B establishes that omitted terms affect the investment threshold only at order $O(\sigma_\alpha^4)$.

Since $\partial\beta^*/\partial r^* > 0$ and $\partial[\beta^*/(\beta^*-1)]/\partial\beta^* < 0$, belief uncertainty unambiguously raises the investment threshold through $\theta_{NPV=0}$, with the multiplier effect partially offsetting.

Proof Sketch. See Appendix A.1 for the full proof. Given Lemma 1, the value of waiting satisfies the one-dimensional HJB (9). A second-order Taylor expansion of $V(\theta,\tilde{\alpha})$ around $\hat{\alpha}$, using $\mathbb{E}[\tilde{\alpha}-\hat{\alpha}] = 0$, yields $\mathbb{E}_{\tilde{\alpha}}[V(\theta,\tilde{\alpha})] \approx V(\theta,\hat{\alpha}) + \frac{1}{2}V_{\alpha\alpha}\sigma_\alpha^2$. Since $V_{\alpha\alpha} < 0$ (equation (A.3)), uncertainty reduces the expected value of an installed project, creating the precautionary motive to delay. We seek r^* such that $A(\theta)k^\eta/r^* = \mathbb{E}_{\tilde{\alpha}}[A(\theta)k^\eta/(r-g(\tilde{\alpha},\theta))]$, which to second order yields $r^* = r+\lambda\sigma_\alpha^2$ with λ as in equation (8). Positing $W(\theta) = B\theta^\beta$ in a neighbourhood of θ^*, the local approximation of Appendix A.1, Step 3 gives the characteristic equation (13). Applying value-matching and smooth-pasting at θ^* gives equation (11); Lemma 2 below confirms this fixed-point equation has a unique solution $\theta^* \in (0,\bar{\theta})$.

□

Corollary 1 (Zone of Maximum Paralysis)**.** *The investment premium $\theta^*/\theta_{NPV=0} - 1$ is a strictly concave function of $\theta_0 \in (0,\bar{\theta})$, attaining its maximum at an interior point $\theta^{**} \in (0,\bar{\theta})$. Thus the combined effect of institutional volatility and belief uncertainty is largest at intermediate institutional quality—neither very poor institutions, where investment has already ceased, nor very good ones, where the reform option is nearly exhausted.*

Proof. The premium is $\beta^*/(\beta^*-1)-1$, where $r^* = r+\lambda(\theta_0)\sigma_\alpha^2$ and $\lambda(\theta_0) = (\bar{\theta}-\theta_0)^2/(2\theta_0\bar{\theta})$. The function $\lambda(\theta_0)$ is monotone decreasing in θ_0, so r^* and β^* both decrease as $\theta_0 \to \bar{\theta}$. Separately, the diffusion coefficient $\nu\sqrt{\theta_0(\bar{\theta}-\theta_0)}$ governing $\tilde{\nu}$ is hump-shaped with maximum at $\theta_0 = \bar{\theta}/2$. The product of these two effects—increasing near zero, decreasing near $\bar{\theta}$—implies the premium is hump-shaped with interior maximum θ^{**}. □

Lemma 2 (Existence and Uniqueness of the Investment Threshold)**.** *Define $M \equiv \beta^*/(\beta^*-1) > 1$, which depends only on $(r^*,\tilde{\nu},\hat{\alpha})$ and is therefore constant with respect to θ. The optimal threshold satisfies the fixed-point equation*

$$\theta^* = T(\theta^*), \qquad T(\theta) \equiv \left(\frac{r^*I}{A(\theta)k^\eta}\right)^{1/(1-\eta)} \cdot M. \tag{15}$$

This equation has a unique solution $\theta^* \in (0,\bar{\theta})$.

Proof. Since $A(\cdot)$ is strictly increasing and continuous on $(0,\bar{\theta}]$ with $A(0) = 0$, $T(\theta)$ is strictly decreasing and continuous, with $T(\theta) \to +\infty$ as $\theta \to 0^+$ and $T(\bar{\theta}) = (r^*I/\bar{A}k^\eta)^{1/(1-\eta)}\cdot M > 0$. The identity map $\theta \mapsto \theta$ is strictly increasing from 0 to $\bar{\theta}$. Since $T(\theta)-\theta$ is strictly decreasing and changes sign on $(0,\bar{\theta})$, the intermediate value theorem guarantees a unique crossing. □

Remark 1 (Relationship to Dixit and Pindyck [1994]). *The total investment premium decomposes as:*

$$\frac{\theta^*}{\theta_{NPV=0}} = \underbrace{\frac{\beta_0}{\beta_0 - 1}}_{\text{standard real option}} \times \underbrace{\frac{\beta^*/(\beta^* - 1)}{\beta_0/(\beta_0 - 1)}}_{\text{belief uncertainty amplification}} \tag{16}$$

Our contribution is extending this structure to incorporate belief uncertainty σ_α^2 through the adjusted discount rate r^*, yielding the amplification factor absent from their framework.

Example 1 (Quantitative Impact). *For empirically relevant parameters:*[3]

- *Baseline:* $r = 0.10$, $\tilde{\nu} = 0.15$, $\hat{\alpha} = 0.20 \Rightarrow \beta_0 = 2.31$, *multiplier* $= 1.76$
- *Institutional gap:* $(\bar{\theta} - \theta_0)/\theta_0 = 0.67 \Rightarrow \lambda = 0.133$
- *Belief uncertainty:* $\sigma_\alpha^2 = 0.10 \Rightarrow r^* = 0.10 + 0.133 \times 0.10 = 0.1133$
- *Result:* $\beta^* = 2.58$, *multiplier* $= 1.63$*; break-even threshold rises by* $\approx 20\%$*, so* θ^* *rises by approximately* 9%

The net increase in θ^* of roughly 9 percentage points arises purely from uncertainty about reform credibility, additional to the standard real options premium. Two features deserve emphasis. First, the multiplier itself falls slightly, from 1.76 to 1.63, as β^* rises, partially offsetting the threshold rise. Second, and more importantly, it is the rise in the break-even threshold rather than the change in the multiplier that drives the result. Section 4.5 establishes threshold dominance systematically.

3.4 Credibility Dynamics and Learning

Recall from Definition 2 that $\rho_t \in [0, 1]$ denotes the private sector's posterior probability that the government is competent to implement its announced reforms. This credibility evolves endogenously through Bayesian updating as the private sector observes whether reforms materialise. In continuous time, credibility evolves consistently with the Kalman filter of equation (7): as $\hat{\alpha}_t$ is updated upward by favourable institutional realisations, the posterior

[3] Baseline parameter ranges follow Dixit and Pindyck [1994], Chapter 5, adjusted for emerging market conditions based on Rodrik [1991]. Numerical results are robust to plausible variation in $\theta^*/\bar{\theta}$ across the range $[0.3, 0.7]$. The choice $r = 0.10$ reflects real rates in emerging markets during major reform episodes (IMF IFS, 1990–2020), the paper's target empirical context. At an OECD calibration of $r = 0.04$ the investment premia in Table 1 are *larger* in absolute terms, since a lower discount rate amplifies the option value of waiting; the results are therefore conservative rather than overstated at $r = 0.10$. The qualitative comparative statics are unaffected: the threshold exponent $1/(1 - \eta) \approx 1.49$ is independent of r, and the sign conditions $\partial\beta^*/\partial r^* > 0$ and $\partial[\beta^*/(\beta^* - 1)]/\partial\beta^* < 0$ hold for any $r > 0$.

probability of competence ρ_t rises accordingly. For clarity of exposition we also characterise the discrete-period updating rule:

$$\rho_{t+1} = \frac{\rho_t \cdot \Pr(s_t|\text{competent})}{\rho_t \cdot \Pr(s_t|\text{competent}) + (1-\rho_t) \cdot \Pr(s_t|\text{incompetent})} \tag{17}$$

where $s_t = 1$ if reforms succeed ($\theta_t \geq \theta_{t-1}^{\text{announced}}$) and $s_t = 0$ otherwise. The discrete rule (17) is the period-by-period analogue of the continuous updating embedded in (7): both are driven by the same signal—the gap between announced and realised institutional improvement and converge to the same steady-state beliefs as the observation interval shrinks to zero.

This creates the dynamic feedback central to our mechanism: early reform successes build credibility ($\rho_t \uparrow$), reducing $\sigma^2_{\alpha,t}$ and lowering θ^* through Proposition 1. Conversely, early failures trigger a credibility trap where $\rho_t \downarrow$, raising uncertainty and validating the decision to wait.[4]

4 Growth Implications

4.1 From Individual Firms to Aggregate Dynamics

The analysis thus far has focused on a representative firm's investment decision. The macroeconomic consequences depend on what happens when all firms simultaneously adopt wait-and-see strategies. This section aggregates individual option values to derive economy-wide growth effects, identifying mechanisms that make the aggregate outcome considerably worse than a simple summation of individual delays would suggest.

The first mechanism is the coordination dimension of waiting. In general equilibrium, individual rationality breaks down into collective self-defeat: when investment collapses across the economy, it validates each firm's decision to wait, since aggregate inactivity makes the reform environment appear less promising than it actually is. Widespread inactivity generates a signal that reforms are not working, which raises belief uncertainty $\sigma^2_{\alpha,t}$ for all firms, which raises investment thresholds, which in turn sustains the inactivity generating the pessimistic signal. The aggregate outcome is a self-fulfilling stagnation that no individual firm has an incentive to unilaterally break.

The second mechanism operates through the reform process itself. When investment collapses, governments lose the economic successes needed to demonstrate competence, and credibility ρ_t falls. This feeds directly back into belief uncertainty through equation (17):

[4] A full dynamic programming treatment of endogenous credibility is beyond our scope, but would show that governments face an intertemporal trade-off between ambitious announcements today and credibility capital tomorrow. See the discussion following Proposition 3.

lower ρ_t raises $\sigma^2_{\alpha,t}$, which raises individual investment thresholds, which further depresses investment that would otherwise have restored credibility.

The third complication is sectoral heterogeneity. Capital irreversibility varies substantially across the economy: infrastructure, heavy manufacturing, and chemicals, involve large, long-lived, highly specific investments whose option values are correspondingly high, while service industries face lower irreversibility costs. During reform periods this heterogeneity produces compositional shifts toward short-horizon, easily reversible projects, with consequences for long-run productivity that go beyond the immediate investment shortfall.

4.2 Aggregate Growth with Option Values and Beliefs

The aggregate growth rate reflects the weighted sum of sectoral investment decisions:

$$g = g^*(\bar{\theta}) - \Lambda(\nu, \theta_0, \bar{\theta}) - \Phi(\sigma^2_{\alpha,t}) \tag{18}$$

where $g^*(\bar{\theta})$ is the potential growth rate with reformed institutions, Λ captures the aggregate option value of waiting, and $\Phi(\sigma^2_{\alpha,t})$ represents the additional growth penalty from belief uncertainty.

The aggregate growth penalty Λ relates to firm-level option values through:

$$\Lambda(\nu, \theta_0, \bar{\theta}) = \int_i \omega_i \Omega_i(\theta, \nu)\, di + F(\{\Omega_i\}, \theta) \tag{19}$$

The first term, the weighted average of individual option values, captures Bernanke [1983]'s partial equilibrium effect. The second term F, our contribution, represents general equilibrium feedback: when investment collapses economy-wide, it (i) validates individual waiting decisions through coordination effects, and (ii) undermines reform credibility by making the government appear incompetent.

Proposition 2 (Institutional-Credibility Trap)**.** *When ν is high and credibility is low (high $\sigma^2_{\alpha,0}$), the economy can be trapped in a low-investment equilibrium despite $\bar{\theta} > \theta_0$. Investment remains depressed even though eventual institutional improvement is expected, because low investment causes reforms to appear to fail, which worsens beliefs and raises $\sigma^2_{\alpha,t}$, further validating the decision to wait. This self-fulfilling pessimism sustains a persistent output gap $Y_{actual} \ll Y_{potential}(\bar{\theta})$.*

4.3 General Equilibrium: Aggregation and the Credibility Trap

The formal foundation of Proposition 2 requires closing the loop between individual optimisation, aggregate investment behaviour, and the evolution of reform credibility. Suppose the economy contains a unit mass of firms $i \in [0, 1]$, each characterised by an irreversibility parameter ν_i drawn from a distribution $F(\nu)$ with density $f(\nu)$. Given belief uncertainty $\sigma^2_{\alpha,t}$, firm i invests when θ_t reaches its threshold:

$$\theta^*_i(\sigma^2_{\alpha,t}) = \theta_{NPV=0}(r^*, \nu_i) \cdot \frac{\beta^*(r^*, \nu_i)}{\beta^*(r^*, \nu_i) - 1}, \quad r^* = r + \lambda\sigma^2_{\alpha,t}. \tag{20}$$

Aggregate investment is:

$$I_t = \int_0^1 \mathbf{1}\{\theta_t \geq \theta^*_i(\sigma^2_{\alpha,t})\}\, di = \int_{\nu:\, \theta_t \geq \theta^*(v, \sigma^2_{\alpha,t})} f(\nu)\, d\nu. \tag{21}$$

Credibility updates from aggregate signals:

$$\rho_{t+1} = H(\rho_t, I_t), \quad H_\rho > 0,\ H_I > 0. \tag{22}$$

Belief uncertainty is a function of credibility:

$$\sigma^2_{\alpha,t} = \Sigma(\rho_t), \quad \Sigma_\rho < 0. \tag{23}$$

Together, equations (21)–(23) constitute a closed system. A fixed point of the composed mapping $\Gamma : \rho \mapsto \Sigma(\rho) \mapsto I(\Sigma) \mapsto H(\rho, I)$ is a general equilibrium.[5]

Definition 3 (General Equilibrium)**.** *A general equilibrium is a sequence $\{\theta_t, I_t, \rho_t, \sigma^2_{\alpha,t}\}_{t\geq 0}$ such that: (1) each firm invests according to the threshold rule $\theta_t \geq \theta^*_i(\sigma^2_{\alpha,t})$; (2) aggregate investment I_t is given by* (21)*; (3) credibility evolves according to* (22) *and belief uncertainty satisfies* (23)*; and (4) θ_t evolves according to the Jacobi diffusion* (6)*.*

A steady state $(\theta^{ss}, I^{ss}, \rho^{ss}, \sigma^{2,ss}_\alpha)$ constitutes an *institutional-credibility trap* if I^{ss} is strictly below the investment rate consistent with convergence to $\bar{\theta}$, ρ^{ss} is low, $\sigma^{2,ss}_\alpha$ is high, and $\theta^{ss} < \bar{\theta}$. Any perturbation that reduces I_t propagates through (22) to reduce ρ_t, through (23) to raise $\sigma^2_{\alpha,t}$, and through (21) to raise thresholds and reduce I_t further,

[5] Existence of the trap steady state follows from Tarski's (1955) fixed-point theorem applied to Γ, which is monotone decreasing on the complete lattice $[0, 1]$ given the sign conditions $H_\rho > 0$, $H_I > 0$, $\Sigma_\rho < 0$. The composed map Γ is monotone decreasing in ρ since the decreasing link $\Sigma_\rho < 0$ dominates the composition. Stability follows from Evans and Honkapohja [2001]: the Bayesian updating rule defines a learning dynamic in their sense, and E-stability of the trap steady state follows from the sign structure above.

returning the system to the low-investment fixed point. Escape requires a coordinated shift large enough to move the system to the basin of attraction of the high-investment equilibrium.

4.4 The Option Value of Waiting: Convexity and Uncertainty

Investment falls in reforming economies because Jensen's inequality applied to a convex value function raises the expected value of waiting relative to investing today. The uncommitted firm holds a convex claim on future institutional quality; the committed firm holds a linear claim bearing the full downside without the corresponding upside. This asymmetry operates through two distinct convexity channels that interact multiplicatively at the level of the investment multiplier.

The first channel operates through institutional quality directly. When $V_{\theta\theta} > 0$, the marginal gain from a further improvement in institutions is increasing, so a mean-preserving increase in uncertainty raises the expected value of waiting. The second channel, belief uncertainty about $\tilde{\alpha}$, adds $V_{\alpha\alpha} > 0$: waiting allows the investor to learn, via the Kalman filter (7), whether the government's announced reform trajectory will actually materialise. These two convexity channels interact at the level of the investment multiplier multiplicatively: belief uncertainty raises r^*, which raises β^*, which raises the multiplier applied to the break-even threshold. The dominant quantitative effect, however, operates through the break-even threshold $\theta_{NPV=0} \propto (r^*)^{1/(1-\eta)}$—even moderate belief uncertainty produces substantial threshold increases because the exponent $1/(1-\eta) \approx 1.49$ amplifies small increments to r^*. The full PDE derivation of both convexity channels via the Feynman–Kac representation is in Appendix A.3.

4.5 Numerical Example and Parameter Sensitivity

We calibrate to parameters typical of a crisis-episode emerging market reform: $r = 0.10$, $\hat{\alpha} = 0.15$, $\tilde{\nu} = 0.15$, and an institutional gap of 40% of $\bar{\theta}$ ($\theta_0 = 0.60\bar{\theta}$, giving $\lambda = 0.133$). Without belief uncertainty ($\sigma_\alpha = 0$, $r^* = 0.10$): $\beta_0 = 1.82$, investment premium = 122%.

The central finding is that the threshold effect dominates throughout: it accounts for 83–90% of the net increment in θ^* across all regimes, with the multiplier effect contributing the remainder and partially offsetting the threshold rise. These magnitudes align with observed hurdle rates during institutional transitions. Poland's shock therapy (1990–91) saw reported project hurdle rates of 40–60%, consistent with our crisis-regime predictions. IFC data show infrastructure IRRs of 25–35% in reforming economies versus 8–12% in stable OECD markets. Bank of England surveys document firm hurdle rates rising from 200–300bp above

Table 1: Investment premia by uncertainty regime

Regime	$\tilde{\nu}$	σ_α	r^*	Base premium	Threshold effect	Multiplier effect	Net increment in θ^*
Low	0.10	0.10	0.1013	22.2%	+1.5 pp	+0.3 pp	+1.8 pp
Moderate	0.15	0.20	0.1053	35.8%	+5.9 pp	+0.9 pp	+6.8 pp
High	0.20	0.30	0.1120	53.0%	+12.6 pp	+1.7 pp	+14.3 pp
Crisis	0.25	0.40	0.1213	74.2%	+23.1 pp	+2.5 pp	+25.6 pp

Notes: Base premium = $[\beta_0/(\beta_0 - 1) - 1] \times 100\%$ with $\sigma_\alpha = 0$. Threshold effect = $[(r^*/r)^{1/(1-\eta)} - 1] \times 100\%$, in percentage points added to $\theta^*/\theta_{NPV=0}$. Multiplier effect = $[\beta^*/(\beta^* - 1) - \beta_0/(\beta_0 - 1)] \times 100\%$; negative when $\beta^* > \beta_0$, partially offsetting the threshold effect. Net increment is the sum, in percentage points increase in θ^* above the no-uncertainty baseline. The threshold effect dominates in all regimes, accounting for 83–90% of the total. Parameters: $r = 0.10$, $\hat{\alpha} = 0.15$, $\eta = 0.33$, $\lambda = 0.133$.

WACC in normal times to 800–1200bp during high-uncertainty periods, directly paralleling the belief uncertainty premium $\Psi(\sigma_\alpha^2)$ in equation (1).

5 Empirical Implications and Future Evidence

The framework generates three predictions that are sharp enough to distinguish our mechanism from existing alternatives, and together constitute a coherent empirical agenda. What makes these predictions valuable is not merely qualitative consistency with the development literature, but structural content, specific signs, magnitudes, and cross-sectional patterns that competing theories cannot replicate.

5.1 Prediction 1: Non-Linearity of the Investment–Uncertainty Relationship

Our framework sharpens the standard prediction that investment falls as uncertainty rises in two ways. First, by Corollary 1, the effect is strongest at intermediate institutional quality, where the option value of waiting peaks: neither very poor institutions, where investment has already ceased, nor very good ones, where the reform option is nearly exhausted, generate maximum paralysis. Second, the non-linearity is asymmetric with respect to credibility: for a given increase in uncertainty, the investment response is larger for low-credibility governments because belief uncertainty amplifies the standard real options effect through the multiplicative structure of equation (10):

$$\frac{\partial I}{\partial \nu} = \left.\frac{\partial I}{\partial \nu}\right|_{\text{direct}} + \frac{\partial I}{\partial \theta^*} \cdot \frac{\partial \theta^*}{\partial \nu} < 0 \tag{24}$$

with amplification strongest at intermediate institutional quality. This is consistent with the cross-country evidence in Rodrik [1991] and the uncertainty shock literature of Bloom [2009], though neither framework generates the credibility-dependent asymmetry.

5.2 Prediction 2: Credibility-Dependent Sign Reversal in Reform Announcements

This is our most distinctive prediction and the one most clearly absent from existing frameworks. Reform announcements should have asymmetric effects on investment depending on the announcing government's prior credibility ρ_0:

$$\frac{\partial I}{\partial \alpha^{\text{announce}}} = f(\rho_0, \sigma^2_{\text{prior}}) \tag{25}$$

Investment responses should be positive for high-credibility governments ($\rho_0 > 0.7$), while low-credibility announcements may trigger investment declines. This prediction is sharply at odds with Julio and Yook [2012] and Handley and Limaõ [2015], which treat credibility as uniform and therefore predict similar investment responses to the same reform regardless of who announces it.

5.3 Prediction 3: Irreversibility Gradient across Sectors

Sectors with more irreversible capital, that is infrastructure, manufacturing, chemicals, should exhibit stronger wait-and-see behaviour because the option value of waiting scales with the cost of reversal. Empirically:

$$\frac{I}{Y} = \alpha_0 + \alpha_1\theta + \alpha_2\sigma^2_\theta + \alpha_3(\bar{\theta} - \theta) \cdot \sigma^2_\theta + \alpha_4\sigma^2_{\alpha,t} \tag{26}$$

where $\alpha_3 < 0$ captures the option value effect and $\alpha_4 < 0$ captures belief uncertainty. The investment gap between high- and low-irreversibility sectors should widen during reform episodes as $\sigma^2_{\alpha,t}$ rises, and narrow as credibility accumulates.

5.4 Empirical Agenda and Identification Strategy

While this paper is theoretical, the three predictions above generate empirical signatures with genuine discriminating power between our framework and existing alternatives. This section sketches how the central mechanism could be identified.

The credibility-dependent sign reversal. A minimal empirical design estimates:

$$I_{c,t} = \alpha_c + \gamma_t + \beta_1 A_{c,t} + \beta_3(A_{c,t} \times \rho_{c,t-1}) + \varepsilon_{c,t} \tag{27}$$

where $I_{c,t}$ is investment in country c, $A_{c,t}$ indicates a reform announcement, and $\rho_{c,t-1}$ is a predetermined credibility proxy instrumented using historical institutional quality in the tradition of Acemoglu et al. [2001]. The model predicts $\beta_1 < 0$ and $\beta_3 > 0$: announcements reduce investment for low-credibility governments and stimulate it for high-credibility ones. No competing framework generates this sign reversal. Policy uncertainty à la Bloom [2009] predicts uniformly negative responses regardless of who makes the announcement. Julio and Yook [2012] predicts responses driven by electoral cycles rather than credibility. Credit constraint models predict responses driven by balance sheet conditions rather than announcement content. The interaction term β_3 therefore provides a clean exclusion restriction against all three alternatives.

The irreversibility gradient. The investment gap between high- and low-irreversibility sectors should widen during reform announcement periods and narrow as credibility accumulates. This distinctive sectoral pattern is observable in firm-level datasets such as Compustat or Orbis. Credit shocks produce more uniform sectoral effects because they operate through the cost of capital; demand shocks concentrate effects in cyclically sensitive rather than irreversibility-sorted sectors. The irreversibility gradient thus provides a second exclusion restriction.

The dynamic asymmetry. Investment recovery should be discontinuous—accelerating sharply once accumulated successes drive ρ_t above the moderate-credibility threshold of Proposition 3 rather than tracking institutional quality smoothly. Political uncertainty models predict symmetric recovery as uncertainty resolves; structural adjustment models predict recovery tracking completion of adjustment rather than credibility accumulation. The timing and shape of recovery provides a third exclusion restriction.

What would falsify the theory? The theory is falsified if (i) the interaction term β_3 in equation (27) is zero or negative, that is, if high-credibility announcements produce no stronger investment response than low-credibility ones; (ii) the investment gap between high- and low-irreversibility sectors does not widen during reform periods; or (iii) investment recovery tracks institutional quality improvement smoothly rather than accelerating discontinuously at a credibility threshold. Implementing this agenda using cross-country investment data, firm-level datasets, and the Baker et al. [2016] policy uncertainty index represents an important direction for future research.

6 Policy Implications

The empirical predictions of Section 5 follow directly from the model's structure; the policy implications follow just as directly, and turn on a single insight: reducing reform uncertainty may matter more for investment and growth than the reforms themselves.

6.1 Government Credibility and Announcement Effects

The social value of credible institutional reform can be decomposed as:

$$V_{\text{reform}} = \int_{\theta_0}^{\bar{\theta}} V'(\theta)\, d\theta + \left[\Omega(\nu_{\text{high}}) - \Omega(\nu_{\text{low}})\right] + \left[\Psi(\sigma^2_{\alpha,\text{high}}) - \Psi(\sigma^2_{\alpha,\text{low}})\right] \tag{28}$$

The first term captures the direct improvement in the productive environment. The second and third terms, the reductions in the option value of waiting and in the belief uncertainty premium, can easily dominate this direct effect.

The framework reveals why reform announcements so often fail to stimulate investment. Credibility ρ_0 determines how the private sector interprets a given announced reform speed α^{announce}:

$$\hat{\alpha}_0 = \rho_0 \cdot \alpha^{\text{announce}} + (1 - \rho_0) \cdot \alpha^{\text{prior}} \tag{29}$$

with posterior belief uncertainty:

$$\frac{1}{\sigma^2_{\alpha,0}} = \frac{\rho_0^2}{\sigma^2_{\text{announce}}} + \frac{(1-\rho_0)^2}{\sigma^2_{\text{prior}}} \tag{30}$$

A low-credibility government that announces a very ambitious reform speed maximises $\sigma^2_{\alpha,0}$ rather than reducing it. Bold announcements from low-credibility sources are therefore not merely uninformative; they are actively harmful.

6.1.1 The Optimal Announcement Problem

The government maximises expected output:

$$\max_{\alpha^{\text{announce}}} W = \mathbb{E} \int_0^\infty e^{-rt} A(\theta_t) K_t^\eta\, dt \tag{31}$$

The first-order condition balances two competing effects:

$$\frac{dW}{d\alpha^{\text{announce}}} = \rho \cdot \frac{\partial W}{\partial \hat{\alpha}_0} + \frac{\partial W}{\partial \sigma^2_{\alpha,0}} \cdot \frac{d\sigma^2_{\alpha,0}}{d\alpha^{\text{announce}}} = 0 \tag{32}$$

Proposition 3 (Credibility Regions)**.** *Solving the announcement trade-off yields the optimal announcement:*

$$\alpha^{announce*} = \alpha^{true} + \frac{\rho(2-\rho)}{1+\rho^2} \cdot \Gamma \cdot (\alpha^{true} - \alpha^{prior}) \tag{33}$$

where $\Gamma = \frac{\partial^2 W/\partial \sigma^2_{\alpha,0}}{\partial^2 W/\partial \hat{\alpha}_0} \cdot \frac{\sigma^2_{prior}}{\sigma^2_{announce}}$ *measures the relative importance of uncertainty versus expectations for welfare. The coefficient* $\rho(2-\rho)/(1+\rho^2)$ *is increasing in* ρ*, generating three distinct regimes:*

1. ***Low credibility*** *(*$\rho_0 < 0.3$*): the uncertainty effect dominates. The optimal announcement barely deviates from market priors:* $\alpha^{announce*} \approx \alpha^{prior} + 0.2(\alpha^{true} - \alpha^{prior})$*. Truthful announcement is suboptimal not because the government wishes to deceive, but because markets cannot distinguish ambition from incompetence.*

2. ***Moderate credibility*** *(*$0.3 \leq \rho_0 \leq 0.7$*): the two effects approximately balance, making truthful announcement optimal:* $\alpha^{announce*} \approx \alpha^{true}$*.*

3. ***High credibility*** *(*$\rho_0 > 0.7$*): the expectation effect dominates. The government can afford to be aspirational, announcing* $\alpha^{announce*} > \alpha^{true}$ *without generating destabilising uncertainty.*

This non-monotonicity has direct implications for reform design: low-credibility governments should under-promise and over-deliver, building the reputational stock necessary for transformative announcements later.

The formal derivation, including the welfare derivatives and the Gaussian updating structure underlying equation (33), is in Appendix D.

This non-monotonicity is illustrated with particular clarity by the Truss–Kwarteng mini-budget of September 2022: bold fiscal promises made without the institutional backing needed to establish credibility triggered immediate investment paralysis and a sovereign bond market crisis. The announcements were precisely wrong from the perspective of managing belief uncertainty, an ambitious announcement from a low-credibility source maximises $\sigma^2_{\alpha,0}$, deepening paralysis rather than stimulating the investment the budget was intended to generate.

6.1.2 Political Economy Constraints

Our analysis abstracts from several practical constraints on reform credibility. Electoral pressures create perverse incentives: governments approaching elections tend to choose α^{announce} to maximise short-term visibility rather than long-term credibility, interacting destructively with our option value mechanism. Veto players, coalition partners, courts, labour unions, raise implementation uncertainty independently of government competence, linking σ_α^2 directly to institutional architecture. Time inconsistency problems mean investors anticipate potential reneging once investment is sunk, explaining why external commitment devices such as IMF programmes and EU accession processes raise investment by reducing future policy flexibility. A full treatment of these constraints is left for future work.

6.2 Application to Reform Episodes

The following episodes are *illustrative simulations*. Each maps the model's parameters to historically documented values to show the framework can generate quantitative magnitudes consistent with observed investment behaviour. Formal identification of the credibility channel, separating it from credit shocks, demand effects, and political instability, requires the empirical strategy outlined in Section 5 and is left to future work.

6.2.1 Poland's Shock Therapy (1990–1991)

In July 1990, Poland announced the most radical institutional reforms in post-communist Europe. Despite widespread agreement that the reforms were necessary and would ultimately succeed, investment collapsed by 35% in the subsequent year, with firms explicitly citing uncertainty about the new institutional framework as their reason for delay [Balcerowicz, 1995].

While structural factors, credit collapse, the termination of subsidised investment, and the unwinding of central planning relationships—contributed to the decline, our framework identifies a distinct and complementary channel. Mapping to Table 1, the economy passed through the crisis uncertainty regime before credibility could accumulate. By Proposition 3, regime 1 applied by construction: no market economy had previously been built from central planning, providing no track record from which to assess whether the announced reform speed α^{announce} would materialise. This maximised belief uncertainty $\sigma_{\alpha,0}^2$, driving the investment threshold θ^* to crisis-level heights. The reported hurdle rates of 40–60% and real interest rates exceeding 50% observed during this period are consistent with our crisis-regime simulation.

The subsequent recovery from 1992 is consistent with the model's dynamic predictions. As the reform programme delivered consistent early successes, ρ_t rose through the credibility accumulation mechanism of equation (17), reducing $\sigma^2_{\alpha,t}$ and lowering θ^*. Investment accelerated not when structural adjustment was complete, but when reform credibility had accumulated sufficiently to resolve the option value of waiting—consistent with Proposition 1's prediction of a discontinuous threshold crossing.

6.2.2 Post-Brexit Investment Paralysis (2016–2020)

The UK's post-referendum period exemplifies sustained belief uncertainty of a qualitatively different kind. In Poland, ρ_0 was low but $\bar{\theta}$ was well-defined—the destination was a standard market economy whose properties investors could assess. Post-Brexit, both ρ_0 and $\bar{\theta}$ were uncertain: investors could not determine the government's competence to deliver, nor the institutional endpoint toward which reforms were directed. This compound uncertainty maps to the crisis regime of Table 1.

Despite repeated government assurances about "Brexit opportunities," business investment fell in twelve of the fourteen quarters between the referendum and COVID-19. The investment decline concentrated in sectors with high irreversibility—automotive (-30%), aerospace (-25%), chemicals (-20%), exactly as the model predicts given differential option values. Service sectors with lower capital intensity showed considerably more resilience, consistent with the sectoral heterogeneity of Prediction 3. The December 2019 election, by clarifying the political path if not the institutional destination, reduced σ^2_α and coincided with the first investment uptick in eight quarters, consistent with Proposition 3, regime 2. The subsequent COVID-19 shock prevents a clean identification of the counterfactual.

6.2.3 Lessons from the Two Episodes

The two episodes span Proposition 3's credibility spectrum systematically. Poland illustrates regime 1 by construction: the absence of any market-economy precedent made ρ_0 negligible and gradualism impossible. The priority was to accumulate credibility as rapidly as possible through early reform successes. The post-1992 recovery confirms that investment accelerated not when structural adjustment was complete but when credibility had accumulated sufficiently to resolve the option value of waiting.

Post-Brexit illustrates the partial resolution achievable under regime 2 when the institutional destination itself remains uncertain. Unlike Poland, where $\bar{\theta}$ was well-defined and the credibility problem was purely about implementation speed, post-Brexit uncertainty compounded across both ρ_0 and $\bar{\theta}$, generating a more persistent paralysis that moderate credibility improvements could only partially resolve. Without a defined institutional destination, credibility-building through delivery was structurally impeded, which is why external anchors offer a more reliable instrument than domestic reform sequencing alone.

In both episodes the policy implication is consistent: implementation credibility matters more than announcement ambition. The instrument required to build that credibility differs. For Poland it was sequenced early delivery within an unavoidable shock therapy context. For post-Brexit it was, and remains, institutional destination clarity—something only external anchors or negotiated frameworks can reliably provide.

7 Conclusion

This paper began with a puzzle: why do institutional reforms, even those widely expected to succeed, often trigger investment collapse rather than economic renewal? Our answer lies in a fundamental tension between the promise of future improvement and the uncertainty inherent in institutional change. We develop a real options framework in which belief uncertainty about reform credibility amplifies traditional wait-and-see effects through a multiplicative structure at the level of the investment multiplier. Even moderate uncertainty raises investment hurdles by 7–15 percentage points; crisis-level uncertainty raises them by 25 percentage points or more above the no-uncertainty baseline. These effects operate most powerfully at intermediate levels of institutional quality, creating a zone of maximum paralysis (Corollary 1) that can trap economies despite favourable long-run prospects, and they are concentrated in the break-even threshold $\theta_{NPV=0} \propto (r^*)^{1/(1-\eta)}$ rather than in the option multiplier itself.

The policy implications follow directly. Reducing the volatility of the reform path—through gradualism, sequenced delivery, or external anchors—may matter more for investment than the reforms themselves. External commitment devices such as IMF programmes and EU accession processes succeed precisely because they reduce belief uncertainty more than they improve institutional fundamentals, resolving the credibility problem that domestic announcements cannot. Most strikingly, low-credibility governments face a cruel dilemma: bold reform announcements maximise σ_α^2 and thereby deepen the paralysis they are intended to cure.

The framework generates three testable predictions: a credibility-dependent sign reversal in investment responses, an irreversibility gradient across sectors, and a discontinuous pattern of investment recovery, with the interaction term β_3 in equation (27) providing a clean exclusion restriction against alternatives including Bloom-type policy uncertainty, Julio–Yook electoral cycles, and credit constraint models. Implementing this empirical agenda is an important direction for future research.

Our analysis abstracts from the endogenous evolution of credibility through sustained reform delivery, from electoral pressures and veto players that shape announcement incentives, and from the distributional conflicts that determine which reforms are politically feasible. The central message, however, is clear: the binding constraint on investment and growth during institutional transitions may be neither the design of good institutions nor the political will to implement them, but the management of the uncertainty that inevitably surrounds institutional change.

Appendix A: Proofs and Derivations

A.1 Proof of Proposition 1: Investment Threshold with Belief Uncertainty

Step 1: Certainty-Equivalent Representation. By Lemma 1, the joint dynamics of $(\theta_t, \hat{\alpha}_t)$ may be reduced to the one-dimensional state θ, with belief uncertainty entering through the adjusted discount rate r^*. An installed project earning flow output $A(\theta_t)k^\eta$ per unit time grows at the institutional drift rate $g(\hat{\alpha}, \theta) \equiv \hat{\alpha}(\bar{\theta}/\theta - 1)$. Discounting at rate r yields the perpetuity value:

$$V(\theta, \hat{\alpha}) = \frac{A(\theta)k^\eta}{r - g(\hat{\alpha}, \theta)}, \qquad g(\hat{\alpha}, \theta) \equiv \hat{\alpha}\left(\frac{\bar{\theta}}{\theta} - 1\right) \tag{34}$$

Taking a second-order Taylor expansion of $V(\theta, \tilde{\alpha})$ around $\hat{\alpha}$ and using $\mathbb{E}[\tilde{\alpha} - \hat{\alpha}] = 0$:

$$\mathbb{E}_{\tilde{\alpha}}\big[V(\theta, \tilde{\alpha})\big] = V(\theta, \hat{\alpha}) + \tfrac{1}{2}V_{\alpha\alpha}(\theta, \hat{\alpha})\sigma_\alpha^2 \tag{35}$$

Differentiating twice with respect to $\hat{\alpha}$:

$$V_{\alpha\alpha} = -\frac{2A(\theta)k^\eta(\bar{\theta} - \theta)^2/\theta^2}{[r - g(\hat{\alpha}, \theta)]^3} < 0 \tag{36}$$

Since $V_{\alpha\alpha} < 0$, uncertainty about $\tilde{\alpha}$ reduces the expected value of an installed project, creating a precautionary motive to delay investment.

Step 2: Equivalent Discount Rate. We seek r^* satisfying

$$\frac{A(\theta)k^\eta}{r^*} = \mathbb{E}_{\tilde{\alpha}}\left[\frac{A(\theta)k^\eta}{r - g(\tilde{\alpha}, \theta)}\right].$$

For small σ_α^2, a second-order expansion yields $r^* \approx r + \lambda\sigma_\alpha^2$, where $\lambda = (\bar{\theta} - \theta_0)^2/(2\theta_0\bar{\theta})$ satisfies the three natural properties discussed in Section 3.2.

Step 3: Modified Characteristic Equation. In the waiting region ($\theta < \theta^*$), the value of waiting W satisfies the HJB equation:

$$r^* W = \hat{\alpha}(\bar{\theta} - \theta)W_\theta + \tfrac{1}{2}\nu^2\theta(\bar{\theta} - \theta)W_{\theta\theta} \tag{37}$$

We seek $W(\theta) = B\theta^\beta$ in a neighbourhood of $\theta^* \in (0, \bar{\theta})$. Approximating $\theta(\bar{\theta} - \theta) \approx \theta^*(\bar{\theta} - \theta^*) \equiv \Delta^*$ and absorbing constants into rescaled coefficients $\tilde{\alpha} = \hat{\alpha}(\bar{\theta} - \theta^*)/\theta^*$ and $\tilde{\nu}^2 = \nu^2(\bar{\theta} - \theta^*)/\theta^*$, the characteristic equation becomes (13), with positive root (14).

Step 4: Investment Threshold. Applying value-matching and smooth-pasting at the threshold gives equation (11), where the net effect on θ^* is positive because the rise in $\theta_{NPV=0} \propto (r^*)^{1/(1-\eta)}$ dominates the fall in the multiplier $\beta^*/(\beta^* - 1)$. $\square$

A.2 Parameter Calibration

Table 2: Baseline Parameter Values and Sources

Parameter	Symbol	Value	Source/Justification
Risk-free rate	r	0.10	Real rate in emerging markets (2000–2020 avg); IMF IFS
Institutional volatility	$\tilde{\nu}$	0.15	Locally approx. diffusion; WB Governance Indicators
Expected reform speed	$\hat{\alpha}$	0.20	Panel estimation, transition economies (Poland, Czech Republic, Estonia), 1990–2010
Belief uncertainty	σ_α	0.20	CV$=1.0$; forecast dispersion, Consensus Economics
Capital share	η	0.33	Standard value in growth literature
Institutional gap	$(\bar{\theta} - \theta_0)/\theta_0$	0.67	Crisis calibration ($\theta_0 = 0.60\bar{\theta}$); Poland 1990, Egypt 2004–2010 (WGI)
Gap parameter	λ	0.133	Derived: $(0.40)^2/(2 \times 0.60) = 0.133$
Prior credibility	ρ_0	0.40	Posterior prob. of competence at announcement; low-to-moderate baseline

Notes: Within-country s.d. of institutional quality ranges from 0.10 to 0.25 (median 0.15). CV $\sigma_\alpha/\hat{\alpha}$ typically 0.5–2.0 during major reforms. Credibility regions in Proposition 3: low ($\rho_0 < 0.3$), moderate ($0.3 \leq \rho_0 \leq 0.7$), high ($\rho_0 > 0.7$).

A.3 The Feynman–Kac Representation and Generator Decomposition

From the Feynman–Kac theorem, the value function satisfies $rV = A(\theta)k^{\eta} + \mathcal{L}V$, where the infinitesimal generator decomposes as:

$$\mathcal{L}V = \underbrace{\hat{\alpha}(\bar{\theta}-\theta)V_\theta}_{\text{institutional drift}} + \underbrace{\tfrac{1}{2}\nu^2\theta(\bar{\theta}-\theta)V_{\theta\theta}}_{\text{institutional convexity}} + \underbrace{\tfrac{1}{2}\sigma^2_{\alpha,t}(\bar{\theta}-\theta)^2 V_{\alpha\alpha}}_{\text{belief convexity}} \tag{38}$$

The complete generator also contains a cross-derivative term $\sigma^2_\alpha(\bar{\theta}-\theta)V_{\theta\hat{\alpha}}$; Appendix B establishes that this affects the investment threshold only at order $O(\sigma^4_\alpha)$.

The option value resides in the second and third terms. By a second-order Taylor expansion: $\mathbb{E}[V(\theta+\Delta\theta)] = V(\theta) + V_\theta\mathbb{E}[\Delta\theta] + \frac{1}{2}V_{\theta\theta}\mathrm{Var}(\Delta\theta)$. The last term is positive when $V_{\theta\theta} > 0$. The same logic applies to $V_{\alpha\alpha}$: convexity in beliefs means that waiting to learn about $\tilde{\alpha}$ has positive value, with the effect scaling as $(\bar{\theta}-\theta)^2$, strongest when institutions are far from steady state and vanishing as $\theta \to \bar{\theta}$.

Appendix B: Validity of the Certainty-Equivalent Approximation

B.1 The Full Two-Dimensional HJB Equation

Institutional quality evolves as in equation (6). From the Kalman filter (7):

$$d\hat{\alpha}_t = \frac{\sigma^2_{\alpha,t}}{\nu\sqrt{\theta_t(\bar{\theta}-\theta_t)}}\,dB_t \tag{39}$$

Since both state variables are driven by the same Brownian motion B_t:

$$d\langle\theta,\hat{\alpha}\rangle_t = \sigma^2_{\alpha,t}\,dt \tag{40}$$

The complete infinitesimal generator acting on a smooth function $W(\theta,\hat{\alpha})$ is:

$$\mathcal{L}W = \hat{\alpha}(\bar{\theta}-\theta)W_\theta + \tfrac{1}{2}\nu^2\theta(\bar{\theta}-\theta)W_{\theta\theta} + \frac{\sigma^4_\alpha}{2\nu^2\theta(\bar{\theta}-\theta)}W_{\hat{\alpha}\hat{\alpha}} + \sigma^2_\alpha W_{\theta\hat{\alpha}} \tag{41}$$

The full HJB in the waiting region is $rW = \mathcal{L}W$.

B.2 Perturbation Expansion in σ_α^2

Let $\varepsilon \equiv \sigma_\alpha^2$ and expand:

$$W(\theta, \hat{\alpha}; \varepsilon) = W^{(0)}(\theta) + \varepsilon W^{(1)}(\theta, \hat{\alpha}) + \varepsilon^2 W^{(2)}(\theta, \hat{\alpha}) + O(\varepsilon^3) \tag{42}$$

Order ε^0. Positing $W^{(0)} = B_0 \theta^{\beta_0}$ and applying the local approximation of Appendix A.1, Step 3, the characteristic equation is:

$$\tfrac{1}{2}\tilde{\nu}^2 \beta_0(\beta_0 - 1) + \tilde{\alpha}\beta_0 - r = 0 \tag{43}$$

with positive root β_0, yielding the standard investment multiplier $\beta_0/(\beta_0 - 1)$.

Order ε^1. Since $W^{(0)}$ depends only on θ, $W^{(0)}_{\theta\hat{\alpha}} = 0$ and $W^{(0)}_{\hat{\alpha}\hat{\alpha}} = 0$. The first-order correction $W^{(1)}$ satisfies:

$$rW^{(1)} = \hat{\alpha}(\bar{\theta} - \theta)W^{(1)}_\theta + \tfrac{1}{2}\nu^2 \theta(\bar{\theta} - \theta)W^{(1)}_{\theta\theta} + R^{(1)}(\theta, \hat{\alpha}) \tag{44}$$

Solving yields $W^{(1)} = B_1(\hat{\alpha})\theta^{\beta^*}$ where β^* solves:

$$\tfrac{1}{2}\tilde{\nu}^2 \beta(\beta - 1) + \tilde{\alpha}\beta - r^* = 0, \qquad r^* = r + \lambda\sigma_\alpha^2 \tag{45}$$

This is precisely the characteristic equation (13) of Proposition 1. The cross-derivative and $W_{\hat{\alpha}\hat{\alpha}}$ terms contribute to the investment threshold only at order $O(\varepsilon^2) = O(\sigma_\alpha^4)$.

B.3 Implication for the Investment Threshold

Applying value-matching and smooth-pasting to the perturbed value function:

$$\theta^* = \theta_{NPV=0} \cdot \frac{\beta^*}{\beta^* - 1} + O(\sigma_\alpha^4) \tag{46}$$

The certainty-equivalent representation in Proposition 1 is exact to first order in σ_α^2, and the omitted terms from the generator (41) are negligible at this order. This completes the proof of Lemma 1.

Appendix C: Notation

Institutional dynamics	
$\theta_t \in [0, \bar{\theta}]$	Institutional quality; $\bar{\theta}$ is the maximum attainable level
$\tilde{\alpha}$, $\hat{\alpha}_t$, $\sigma^2_{\alpha,t}$	True (unobserved) reform speed; posterior mean and variance at time t
ν, B_t	Diffusion coefficient; standard Brownian motion
$g(\hat{\alpha}, \theta)$	Institutional drift, $\hat{\alpha}(\bar{\theta}/\theta - 1)$
Production and investment	
$A(\theta)$, k, $\eta \in (0, 1)$	Productivity of installed capital; capital stock; capital share
I, r	Sunk investment cost; risk-free discount rate
$V(\theta, \hat{\alpha})$	Value of installed project, $A(\theta)k^{\eta}/[r - g(\hat{\alpha}, \theta)]$
Investment thresholds and option values	
θ^*, $\theta_{NPV=0}$	Optimal threshold; break-even threshold under certainty
β^*, β_0	Characteristic roots with and without belief uncertainty
$r^* = r + \lambda\sigma^2_{\alpha}$	Uncertainty-adjusted discount rate
$\lambda = (\bar{\theta} - \theta_0)^2/(2\theta_0\bar{\theta})$	Institutional gap parameter
Cost of capital and credibility	
$\pi(\theta)$	Risk premium from current institutional volatility
$\Omega(\theta, \nu)$, $\Psi(\sigma^2_{\alpha,t})$	Option value of waiting; belief uncertainty premium
$\rho_t \in [0, 1]$, ρ_0	Posterior probability government is competent; prior credibility at announcement
α^{announce}, α^{prior}	Announced and prior reform speeds
$\sigma^2_{\text{announce}}$, σ^2_{prior}	Variance of announcement signal; variance of prior belief
Aggregate growth	
$g^*(\bar{\theta})$, Λ, $\Phi(\sigma^2_{\alpha,t})$	Potential growth rate; aggregate option value of waiting; aggregate belief uncertainty penalty

Appendix D: Derivation of Proposition 3

Standard Gaussian updating of the prior $\tilde{\alpha} \sim \mathcal{N}(\alpha_{\text{prior}}, \sigma^2_{\text{prior}})$ on the signal $\alpha^{\text{announce}} \sim \mathcal{N}(\tilde{\alpha}, \sigma^2_{\text{announce}})$ yields equations (28)–(29), with the weight $\rho_0 = \sigma^2_{\text{prior}}/(\sigma^2_{\text{prior}} + \sigma^2_{\text{announce}})$ identified with credibility.

By Proposition 1 and threshold dominance (Section 4.5), $\partial W/\partial\hat{\alpha}_0 > 0$ and $\partial W/\partial\sigma^2_{\alpha,0} < 0$. The government maximises $W(\hat{\alpha}_0(\alpha^{\text{announce}}), \sigma^2_{\alpha,0}(\alpha^{\text{announce}}))$; the first-order condition (31) then yields (32) on substituting (28)–(29) and defining Γ as the ratio of the two welfare derivatives weighted by $\sigma^2_{\text{prior}}/\sigma^2_{\text{announce}}$. The coefficient $\rho_0(2-\rho_0)/(1+\rho_0^2)$ is strictly increasing in ρ_0 on $[0, 1]$, generating the three regimes of Proposition 3. $\square$

Note: During the preparation of this work the author used Claude (Anthropic) to assist with editing, cross-reference checking, and Latex corrections.

References

Abel, A. B., Dixit, A. K., Eberly, J. C., and Pindyck, R. S. (1996). Options, the value of capital, and investment. *Quarterly Journal of Economics*, 111(3), 753–777.

Abel, A. B. and Eberly, J. C. (1994). A unified model of investment under uncertainty. *American Economic Review*, 84(5), 1369–1384.

Acemoglu, D., Johnson, S., and Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5), 1369–1401.

Acemoglu, D. and Robinson, J. A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown Publishers.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4), 1593–1636.

Balcerowicz, L. (1995). *Socialism, Capitalism, Transformation*. Central European University Press.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics*, 98(1), 85–106.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623–685.

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

Décamps, J. P. and Mariotti, T. (2004). Investment timing and learning externalities. *Journal of Economic Theory*, 118(1), 80–102.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press.

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

Gennotte, G. (1986). Optimal portfolio choice under incomplete information. *Journal of Finance*, 41(3), 733–746.

Handley, K. and Limaõ, N. (2015). Trade and investment under policy uncertainty: Theory and firm evidence. *American Economic Journal: Economic Policy*, 7(4), 189–222.

Julio, B. and Yook, Y. (2012). Political uncertainty and corporate investment cycles. *Journal of Finance*, 67(1), 45–83.

Karlin, S. and Taylor, H. M. (1981). *A Second Course in Stochastic Processes*. Academic Press, New York.

Liptser, R. S. and Shiryaev, A. N. (2001). *Statistics of Random Processes I: General Theory* (2nd ed.). Springer-Verlag.

McDonald, R. and Siegel, D. (1986). The value of waiting to invest. *Quarterly Journal of Economics*, 101(4), 707–727.

North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.

Oksendal, B. (2003). *Stochastic Differential Equations: An Introduction with Applications* (6th ed.). Springer-Verlag.

Pastor, L. and Veronesi, P. (2012). Uncertainty about government policy and stock prices. *Journal of Finance*, 67(4), 1219–1264.

Pastor, L. and Veronesi, P. (2013). Political uncertainty and risk premia. *Journal of Financial Economics*, 110(3), 520–545.

Rodrik, D. (1991). Policy uncertainty and private investment in developing countries. *Journal of Development Economics*, 36(2), 229–242.

Stokey, N. L. (2016). Technology adoption under uncertainty: A real options approach. *Journal of Economic Dynamics and Control*, 63, 1–24.