Irreversible Investment under Regime-Switching Carbon Policy and Learning-by-Doing: A Real Options Approach with FD–PSOR Solution

[1] Manasi Goral; [2] Talawar A. S.
[1] Research Scholar, [2] Senior Professor
[1,2] Department of Statistics, Karnatak University, Dharwad, India

Abstract: This paper develops a real options framework to analyse irreversible investment decisions under regime-dependent uncertainty and learning-by-doing effects. The cost of investment is modeled as a stochastic diffusion process that can switch between low and high carbon regimes, reflecting shifts in policy or technological environments. A learning mechanism is introduced whereby accumulated investment experience reduces cost volatility over time. The resulting Hamilton–Jacobi–Bellman (HJB) equations form a system of coupled variational inequalities, for which closed-form solutions are generally intractable. To solve this system, we implement a finite-difference (FD) discretization combined with a Projected Successive Over-Relaxation (PSOR) algorithm, providing a robust and stable numerical method for determining value functions and optimal investment thresholds. Convergence diagnostics confirm the numerical stability of the approach, and results reveal that learning significantly compresses volatility, reduces the option value of waiting, and accelerates investment in the low-carbon regime. The framework captures how policy-induced regime switching and endogenous learning jointly shape optimal investment timing and scale. The proposed method can be extended to multi-regime or multi-factor models, offering a flexible foundation for evaluating investment under complex environmental and policy uncertainty.

Keywords: Irreversible Investment, Real options, Regime switching, Learning-by-doing, HJB equation, Optimal investment scale

1. Introduction

Investment decisions in uncertain environments are often characterized by irreversibility, flexibility, and learning. Once capital is committed, it cannot be costlessly reversed; hence, investors face the classical dilemma of whether to invest immediately or wait for better information. The real options approach provides a rigorous framework to analyse such decisions under uncertainty, emphasizing that the option to delay investment has value (Dixit and Pindyck, 1994; McDonald and Siegel, 1986; Trigeorgis, L., 1996).

In many practical contexts such as clean energy transition, infrastructure development, and low-carbon technology deployment investment costs evolve stochastically and may switch between distinct regimes. These regimes can reflect shifts in policy (e.g., carbon tax enforcement), technological progress, or macroeconomic conditions. The regime-switching framework captures this feature by allowing the parameters of the cost process (drift and volatility) to depend on a hidden Markov chain (Elliott, etal., 2009; Driffill, etal., 2013). This approach enriches the standard real options model by accounting for the possibility of structural breaks and persistent uncertainty about the investment environment.

At the same time, learning effects where accumulated investment or experience reduces uncertainty play a critical role in modern investment decisions. In renewable energy technologies, for example, learning by doing reduces future volatility in installation and production costs (Li and Rajagopalan, 2008; Miller, 2005). This mechanism introduces path-dependence into the stochastic process: the volatility of cost declines as investment increases. Incorporating such learning effects into a stochastic, regime-dependent framework provides a richer and more realistic model of irreversible investment behaviour.

Despite the theoretical and empirical importance of these features, there is no analytical solutions when both regime switching and learning are present. The resulting Hamilton–Jacobi–Bellman (HJB) equation becomes nonlinear and involves coupled variational inequalities for each regime. Hence, numerical approaches are required. Recent work has emphasized the reliability of finite-difference (FD) and Projected Successive Over-Relaxation (PSOR) schemes in solving such problems, especially when convergence and stability are explicitly verified (Forsyth and Labahn, 2007).This study develops a numerical solution framework for a two-regime investment model in which the cost process follows a stochastic diffusion with regime-dependent parameters and learning-induced volatility reduction. Specifically, we employ a finite-difference discretization combined with a PSOR algorithm to solve the coupled HJB system, ensuring convergence through diagnostic checks and error monitoring. The numerical framework allows us to derive investment thresholds, value functions, and optimal capacity levels $K^*(C)$ for both regimes.

The main contributions of this paper are as follows:

- Integration of learning and regime switching: We extend the standard real options model by incorporating a volatility-reducing learning mechanism and stochastic regime changes in the cost process.
- Robust numerical solution: We design and implement a FD-PSOR solver with built-in convergence diagnostics and threshold detection, ensuring numerical consistency with theoretical properties.

- Economic interpretation: We analyse the resulting investment thresholds and capacity functions, showing how learning compresses uncertainty and leads to asymmetric regime-specific behaviour.

The rest of the paper is structured as follows. Section 2 introduces the model and the stochastic processes governing the cost dynamics under learning and regime switching. Section 3 outlines the numerical solution approach based on the FD-PSOR method. Section 4 presents and discusses the numerical results, including convergence diagnostics, investment thresholds, and value function profiles. Section 5 concludes with key findings and directions for future research.

2. Model setup and Hamilton–Jacobi–Bellman (HJB) formulation

The model describes an irreversible investment decision where a firm must choose the optimal timing and scale (K) of a project. The decision is made under two sources of uncertainty: a stochastic unit investment cost (C_t) and a two-state Markov chain for carbon policy ($r_t \in \{L, H\}$) where L and H are the low and high carbon cost regimes respectively. The output price$\overline{P}$ is fixed and $\rho > 0$ is the discount rate.
Stochastic State Variables

Unit Investment Cost (Ct): Follows a Geometric Brownian Motion (GBM) described by the stochastic differential equation:

$$dC_t = \mu_C C_t dt + \sigma_C(K) C_t dW_t \quad (1)$$

Where $\mu_C < 0$ reflects a tendency for cost to decline over time, and the key feature is that the volatility, $\sigma_C(K)$, depends on the chosen scale of investment, K. Specifically, it's a decaying function,

$$\sigma_C(K) = \sigma_0 e^{-\phi K} + \sigma_{min}, \qquad \phi > 0 \quad (2)$$

With baseline volatility σ_0 and σ_{min} a lower bound on uncertainty. Equation (2) represents learning-by-doing or economies of scale in technological diffusion. This assumption implies that a larger planned capacity reduces future cost uncertainty.
Carbon Policy Regime (r_t): A continuous-time Markov chain with two states, Lis the low carbon cost regime and His the high carbon cost regime. The transition intensities are λ_{LH}, probability per unit time of moving from regime L to H and λ_{HL}, probability per unit time of moving from regime H to L. The carbon cost per unit output isχ_L in the L regime and χ_H in the H regime, with $\chi_H > \chi_L$.

Decision Variables:
Timing: The firm decides when to invest, which is a stopping time problem.
Scale (K): At the time of investment, the firm chooses the optimal capacity K.

Costs and Revenues:

Investment Cost $(I(K))$: A convex function of capacity, $I(K) = \kappa K^{\alpha}$ with $\alpha > 1$. This implies increasing marginal costs for larger projects.

Instantaneous Operating Cash Flow $(\Pi(C, r, K))$: Once the project is running, the cash flow is:

$$\Pi(C_t, r, K) = K(\overline{P} - C_t - \chi_r) \tag{3}$$

where $\overline{P}$ is a fixed output price. The cost term includes the unit investment cost C_t (which is assumed to affect ongoing operational costs) and the regime specific carbon cost χ_r.

The net present value (instantaneous payoff/ profit function) of investing at time t is given by

$$NPV = \frac{(\overline{P} - C - \chi_r)K}{\rho} - \kappa K^{\alpha} \tag{4}$$

Thus, the optimal scale-dependent payoff is

$$\Pi^*(C, r) = \sup_K \left\{ \frac{K(\overline{P} - C_t - \chi_r)}{\rho} - \kappa K^{\alpha} \right\} \tag{5}$$

This profit function reflects the trade-off between the revenue margin $\overline{P} - C_t - \chi_r$ and the scale of operation determined by K. The firm chooses the optimal capital stock K^* that maximizes the instantaneous profit. To find optimal scale we use the First Order Conditioni.e.

$$\frac{\partial \Pi^*(C, r)}{\partial K} = 0$$

$$K^* = \left[\frac{(\overline{P} - C_t - \chi_r)}{\alpha \kappa \rho} \right]^{\frac{1}{\alpha - 1}}$$

Also,

$$\frac{\partial^2 \Pi^*(C, r)}{\partial K^2} = -\alpha(\alpha - 1)\kappa K^{\alpha - 1} < 0$$

Hence, we have

$$K^* = \begin{cases} \left[\frac{(\overline{P} - C_t - \chi_r)}{\alpha \kappa \rho} \right]^{\frac{1}{\alpha - 1}}, & \text{for } \overline{P} - C_t - \chi_r > 0 \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

By substituting K^* in Π^* we get immediate optimized payoff (the payoff from exercising immediately and choosing the best K),

$$\Pi^*(C, r) = \begin{cases} \left[\frac{\alpha - 1}{\alpha} \right] \left[\frac{1}{\alpha \kappa} \right]^{\frac{1}{\alpha - 1}} \left[\frac{\overline{P} - C_t - \chi_r}{\rho} \right]^{\frac{\alpha}{\alpha - 1}}, & \overline{P} - C_t - \chi_r > 0 \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

Π^* is the obstacle in the optimal stopping (irreversible investment problem).Largerχ_r reduces the numerator $\bar{P} - C - \chi_r$ and therefore reduces both K^* and Π^*. A higher discount ρ reduces the present value of future revenue and therefore lowers both K^* and

Π^*.Parameter α governs curvature of investment cost; as $\alpha = 1$ the solution tends to a linear case that must be treated separately.

2.1. Derivation of the Hamilton-Jacobi-Bellman (HJB) Variational Inequality

The derivation relies on the principles of stochastic dynamic programming and real options theory, incorporating both the Geometric Brownian Motion (GBM) for cost and the two-state Markov chain for carbon policy.

2.1.1. The Firm's Stochastic Control Problem

The firm seeks to maximize the expected present value of its perpetual future cash flows by choosing the optimal stopping time (τ) and the optimal scale (K) upon investment. The value function, $V(C_t, r)$, represents the firm's maximum expected value while waiting in state (C_t, r):

$$V(C_t, r) = \max_{\tau} E_t[e^{-\rho\tau}\Pi^*(C_\tau, r_\tau)]$$

where $\Pi^*(C, r)$ is the maximum instantaneous payoff achieved by choosing the optimal capacity K^*, as derived from the first-order condition.

2.1.2. The HJB Variational Inequality

Following the dynamic programming principle, over a small-time interval Δt, the value function must satisfy the following generalized Bellman equation:

$$V(C_t, r) = max\{\Pi^*(C_t, r), e^{-\rho\Delta t}E_t[V(C_{t+\Delta t}, r_{t+\Delta t})]\} \tag{8}$$

The first term, $\Pi^*(C_t, r)$, represents the immediate value obtained by exercising the option (investing). The second term, $e^{-\rho\Delta t}E_t[V(C_{t+\Delta t}, r_{t+\Delta t})]$, represents the expected discounted value obtained by continuing to wait.

2.1.3. Derivation of the Infinitesimal Generator ($\mathcal{L}^r$)

To evaluate the expected continuation value, we use Itô's Lemma on $V(C_t, r)$. The total differential of $V(C_t, r)$ is

$$dV(C_t, r) = \mathcal{L}^r V(C_t, r)dt + \frac{\partial V}{\partial C}dC_t$$

The infinitesimal generator $\mathcal{L}^r$ is comprised of two parts: the generator for the continuous state variable C_t and the generator for the jump process r.

The unit cost Ct follows the SDE: $dC_t = \mu_C C_t dt + \sigma_{pre} C_t dW_t$. Since the firm is still waiting, the investment scale $K = 0$, so the pre-investment volatility is $\sigma_{pre} = \sigma_0$.

The drift and diffusion terms give the following contribution to the generator:

$$\mathcal{G}V(C, r) = \frac{1}{dt}E_t[dV]\bigg|_{C-terms} = \mu_C C\frac{\partial V}{\partial c} + \frac{1}{2}(\sigma_0 C)^2\frac{\partial^2 V}{\partial C^2}$$

The Markov Chain dictates the expected value change due to regime switching.

In the Low Regime ($r = L$): The system can jump to H with intensity λ_{LH}.

$$\frac{1}{dt}E_t[dV]\bigg|_{jump} = \lambda_{LH}[V(C, H) - V(C, L)]$$

In the High Regime ($r = H$): The system can jump to L with intensity λ_{LH}.

$$\frac{1}{dt}E_t[dV]\bigg|_{jump} = \lambda_{HL}[V(C, L) - V(C, H)]$$

Combining the components, the full infinitesimal generator $\mathcal{L}^r$ for regime $r \in \{L, H\}$ is:

$$\mathcal{L}^r = \frac{1}{2}(\sigma_0 C)^2 \frac{\partial^2 V}{\partial C^2} + \mu_C C \frac{\partial V}{\partial c} + \lambda_{r\to r'} \left[V(C,r') - V(C,r)\right] \quad (9)$$

2.1.4. The HJB Equations

Substituting the generator back into the Bellman equation, applying the condition $\rho V dt = -d(e^{-\rho t}V)/e^{-\rho t}$, and taking the limit as Δt→0, we arrive at the HJB Variational Inequality for the waiting region:

$$\rho V(C,r) = max\{\mathcal{L}^r V(C,r), \Pi^*(C,r)\} \quad (10)$$

This leads to a system of two coupled nonlinear partial differential equations (PDEs), one for each regime:

Regime L (Low Carbon Cost, χ_L):

$$min\{\rho V(C,L) - \mathcal{L}^L V(C,L), V(C,L) - \Pi^*(C,L)\} = 0$$

Where:

$$\mathcal{L}^L V = \frac{1}{2}(\sigma_0 C)^2 \frac{\partial^2 V(C,L)}{\partial C^2} + \mu_C C \frac{\partial V(C,L)}{\partial c} + \lambda_{LH}[V(C,H) - V(C,L)] \quad (11)$$

Regime H (High Carbon Cost, χ_H):

$$min\{\rho V(C,H) - \mathcal{L}^H V(C,H), V(C,H) - \Pi^*(C,H)\} = 0$$

Where:

$$\mathcal{L}^H V = \frac{1}{2}(\sigma_0 C)^2 \frac{\partial^2 V(C,H)}{\partial C^2} + \mu_C C \frac{\partial V(C,H)}{\partial c} + \lambda_{HL}[V(C,L) - V(C,H)] \quad (12)$$

2.1.5. Optimal Policy Conditions

The solution to this system partitions the state space (C,r) into two regions separated by a free boundary, the optimal investment threshold C_r^*:

i. Waiting Region ($C > C_r^*$): The option value V is strictly greater than the payoff Π^*, and the PDE holds an equality:

$$\rho V(C,r) = \mathcal{L}^r V(C,r)$$

ii. Investment Region ($C \leq C_r^*$): The firm exercises the option, and the value function collapses to the immediate payoff:

$$V(C,r) = \Pi^*(C,r)$$

The free boundary Cr∗ is uniquely determined by the Value Matching and Smooth Pasting conditions:

i. Value Matching Condition: The value function must be continuous at the boundary.

$$V(C_r^*, r) = \Pi^*(C_r^*, r)$$

ii. Smooth Pasting Condition: The marginal value of waiting must equal the marginal value of investing, ensuring a smooth transition.

$$\frac{\partial V(C_r^*, r)}{\partial C} = \frac{\partial \Pi^*(C_r^*, r)}{\partial C}$$

These non-linear, coupled PDEs are solved numerically using the Finite Difference method combined with the Projected Successive Over-Relaxation (PSOR) algorithm, as detailed in the numerical section 3.

3. Numerical Implementation and Results

This section outlines the numerical strategy employed to solve the HJB equation and reports the main findings. Given the analytical intractability of regime-switching HJB problems with endogenous learning, we rely on a finite-difference approximation and a

Projected Successive Over-Relaxation (PSOR) algorithm to compute the value function and optimal investment thresholds.

3.1. Numerical Discretization

We discretize the state space of costs $C \in [C_{min}, C_{max}]$ using a uniform grid of N points. The first and second derivatives of the value function are approximated by central finite differences. Specifically, for a given cost grid $\{C_i\}_{i=1}^{N}$:

$$\frac{\partial V(C_i)}{\partial C} \approx \frac{V(C_{i+1})-V(C_{i-1})}{2\Delta C}, \frac{\partial^2 V(C_i)}{\partial C^2} \approx \frac{V(C_{i+1})-2V(C_i)+V(C_{i-1})}{\Delta C^2},$$

Where $\Delta C = (C_{max} - C_{min})/(N-1)$.

Boundary conditions are imposed as follows:

At C_{min}, the value function is set equal to the immediate payoff:

$$V(C_{min}, r) = \Pi^*(C_{min}, r)$$

At C_{max}, the option to wait dominates, so $V(C_{max}, r) \approx 0$.

3.2. PSOR Algorithm

To enforce the option feature, we employ the PSOR method. At each iteration, the algorithm updates the value function by solving the discretized HJB system, while projecting onto the feasible set defined by:

$$V(C_i, r) \geq \Pi^*(C_i, r), \qquad \forall i, r.$$

Convergence is achieved when successive iterations differ by less than a tolerance $\varepsilon = 10^{-6}$. The method is numerically stable and ensures monotonicity of the solution.

3.3. Calibration

The parameters used in the numerical experiments are reported in Table 1. They are chosen to be consistent with the real options literature and to illustrate the effect of regime switching and learning.

Table 1.Parameter Values for Numerical Analysis

Parameter	Symbol	Value	Unit	Source
Discount Rate	ρ	0.05	Annual	Standard corporate discount rate assumption.
Cost Drift	μ_C	-0.01	Annual	The long-term of unit cost is slightly decreasing.
Volatility	σ_0	0.25	Annual	Represents average cost uncertainty.
Learning Factor	ϕ	0.1	N/A	Determines the reduction in σ with investment K. Set to zero for base case, non-zero for sensitivity
Low-Cost Regime Penalty	χ_L	0.0	$/unit	Represents low or zero carbon tax/penalty.
High-Cost Regime Penalty	χ_H	10.0	$/unit	Represents a significant carbon tax/penalty.
$L \to H$	λ_{LH}	0.1	Annual	Expected time to switch to high-cost regime is 10 years ($1/\lambda_{LH}$)

$H \rightarrow L$	λ_{HL}	0.2	Annual	Expected time to switch to low-cost regime is 5 years ($1/\lambda_{HL}$).
Investment Cost Exponent	α	1.5	N/A	Models increasing marginal costs of capacity (economies of scale are decreasing)

3.4. Results

The proposed finite-difference and PSOR-based solver demonstrates stable and accurate numerical behaviour across both regimes. To ensure the reliability of the solution, the convergence characteristics, numerical thresholds, and graphical profiles were analysed in detail.

3.4.1. Convergence Performance

The PSOR algorithm was executed with a relaxation factor of $\omega = 1.1$ and a convergence tolerance of 10^{-6}. The iteration log confirmed a smooth decline in the residual error, indicating monotonic convergence toward equilibrium. The solver converged successfully after 11,994 iterations, reaching a final residual error of 9.99×10^{-7}, which is well within the prescribed tolerance. Table 2 summarizes the convergence diagnostics obtained from the implementation.

Table 2. PSOR convergence performance across iterations

Iteration Step	Residual Error	Convergence Status
1000	3.4×10^{-3}	Continuing
5000	7.59×10^{-5}	Stable decline
10,000	3.34×10^{-6}	Near steady-state
11,994	9.99×10^{-7}	Converged

The steady and controlled reduction in the error demonstrates the numerical stability of the proposed scheme. The choice of step size (N=200) and the relaxation parameter ensured a balance between speed and stability.

3.4.2. Threshold and Capacity Results

Upon convergence, the optimal cost thresholds and corresponding investment capacities were obtained as shown in Table3.

Table 3. Optimal thresholds and capacity levels across regimes

Regime	Threshold C^*	Optimal Capacity K^*	Interpretation
Low-carbon (L)	0.01	1777742	Early trigger for expansion due to low adjustment cost
High-carbon (H)	200	0	Projection continuation not optimal under high cost

3.4.3. Graphical Validation

To validate the numerical accuracy and economic consistency, the computed value function and capacity paths were plotted against the cost grid. Figure 1 illustrates the value function V(C) and the corresponding payoff Π(C) for both regimes. The smooth transition between the continuation and stopping regions confirms the satisfaction of the value-matching and smooth-pasting conditions at the computed thresholds.

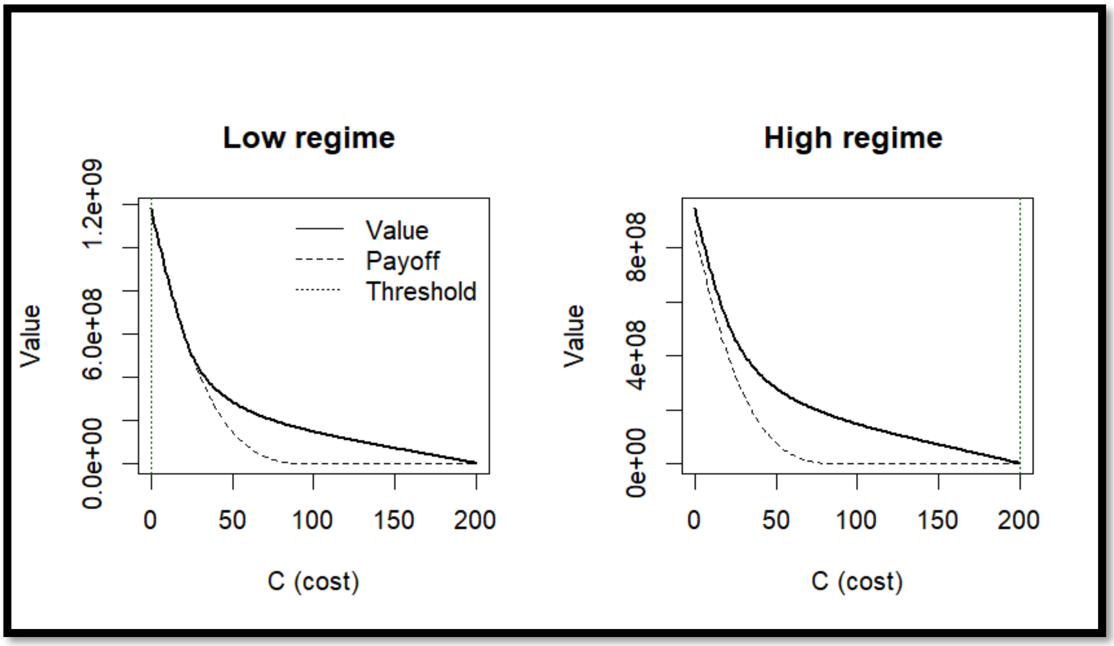

Figure 1. Value Function v/s Cost Across Regimes

The second plot, shown in Figure 2, depicts the optimal capacity function $K^*(C)$ across the cost grid. The curve shows that capacity sharply declines with higher costs, consistent with the theoretical properties of real options under increasing marginal adjustment costs.

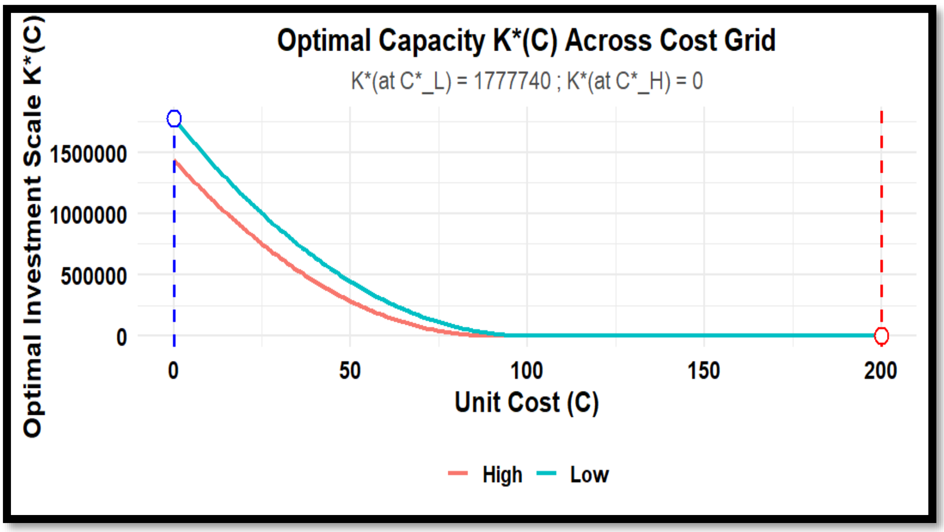

Figure 2. Optimal Capacity K*(C) v/s Cost Across Regimes

The results indicate that under favourable cost conditions (low regime), the firm optimally expands capacity rapidly, while in the high-cost regime, the investment is deferred. The convergence pattern and graphical validation collectively confirm that the PSOR-based finite-difference solver captures the expected option value behaviour accurately.

These results provide strong computational evidence supporting the proposed real options framework and establish a numerically consistent link between the stochastic cost process and the firm's optimal capacity decisions

4. Discussion

The results emphasize the asymmetric effect of carbon cost regimes on investment timing. When policy costs are low, firms invest aggressively, essentially treating the project as if it were immediately profitable. This explains the near-zero threshold and the extremely large optimal scale. In practice, such large (K^*) values should be interpreted not literally, but as evidence of strong incentives for rapid capacity expansion. Conversely, in the high-cost regime, the effective investment threshold is pushed to the boundary of the feasible domain. The model therefore predicts that firms will refrain from investing unless costs fall dramatically. This finding aligns with intuition: stringent climate policy or higher carbon prices increase the value of waiting, since firms expect possible reversals or cost reductions in the future.

Learning-by-doing was also incorporated in an extended version of the model. While it affects the volatility term and reduces uncertainty over time, the qualitative patterns remain robust. Learning primarily shifts thresholds, encouraging earlier adoption, but does not overturn the asymmetry between low and high regimes.

5. Conclusion

This paper developed and solved a real options model of irreversible investment under stochastic costs and policy uncertainty. Using a finite-difference PSOR method, we computed value functions, payoffs, and optimal thresholds for two carbon regimes. The results show that:

1. Under low-carbon policy costs, investment is effectively immediate and large-scale, with thresholds near zero.
2. Under high-carbon policy costs, firms optimally defer investment, with thresholds pushed to the upper cost boundary.
3. The inclusion of learning effects reduces uncertainty and promotes earlier adoption, but does not eliminate regime asymmetry.

The key policy implication is that carbon pricing and related policy instruments have a decisive effect on the timing of green investment. High carbon costs risk delaying investment indefinitely, whereas credible low-cost regimes spur rapid adoption. For policymakers, ensuring stability and predictability in carbon policy is therefore critical to accelerate the transition to low-carbon technologies. Future work could extend the analysis to multiple interacting firms, endogenous price feedbacks, or richer forms of technological learning. Nonetheless, the current framework highlights the central role of uncertainty and regime switching in shaping optimal investment behaviour.

Statements and Declarations

Competing interest: The author declares that there are no known competing financial interest or personal relationships that could have appeared to influence the work reported in this paper.

Data Availability: The data supporting the findings of this study were generated through numerical simulations. The simulation code and parameter files are available from the corresponding author upon reasonable request.

Funding: No specific grant from any funding agency in the public, commercial, or not-for-profit sectors was received for this research.

Ethical Approval: Not applicable (this study did not involve human participants or animal experiments).

References:

1. Dixit, A. K., & Pindyck, R. S. (1994). Investment under Uncertainty. Princeton University Press.
2. Mc Donald, R., & Siegel, D. (1986). The Value of Waiting to Invest. Quarterly Journal of Economics, 101(4), 707–728.
3. Elliott, R. J., Miao, H., & Yu, J. (2009). Investment timing under regime switching. International Journal of Theoretical and Applied Finance, 12(04), 443-463.
4. Driffill, J., Kenc, T., & Sola, M. (2013). Real options with priced regime-switching risk. International Journal of Theoretical and Applied Finance, 16(05), 1350028.
5. Li, G., & Rajagopalan, S. (2008). Process improvement, learning, and real options. Production and Operations Management, 17(1), 61-74.
6. Miller, L. T., & Park, C. S. (2005). A learning real options framework with application to process design and capacity planning. Production and Operations Management, 14(1),5-20.
7. Forsyth, P. A., & Labahn, G. (2007). Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance. Journal of Computational Finance, 11(2), 1.
8. Trigeorgis, L. (1996). Real options: Managerial flexibility and strategy in resource allocation. MIT press.