Corvinus University of Budapest

Doctoral School of Economics and Business Informatics

Institute of Finance

Optimal Timing of Clean Technology Adoption under Correlated Price Uncertainty

A Two-Factor Optimal Stopping Approach to Irreversible Investment

THESIS BOOKLET

Summary of the PhD Dissertation

Réka Janosik

Supervisor: Dávid Zoltán Szabó, PhD

Budapest, 2026

I. Research Background and Justification for the Selection of the Topic

The transition to clean technologies, of which electric vehicle (EV) adoption is the leading example, confronts every prospective adopter with a largely irreversible investment decision under multi-factor uncertainty. The adopter must weigh a known upfront capital cost against a stream of uncertain future operating savings. Acting too early forfeits the option to wait for better conditions; acting too late forgoes years of savings. Standard net present value (NPV) analysis ignores the option value of waiting entirely. The real options literature since McDonald and Siegel (1986) and Dixit and Pindyck (1994) captures irreversibility, but the workhorse specifications either hold the investment cost fixed or let both the project value and the cost follow geometric Brownian motions.

The empirical record shows that this simplification is untenable for clean technology adoption. Both sides of the adoption ledger are volatile. The operating benefit of an EV relative to an internal combustion vehicle is driven by energy price spreads that are volatile and non-stationary; the upfront cost premium exhibits secular decline through production learning (Nykvist and Nilsson, 2015; Ziegler and Trancik, 2021), punctuated by commodity-driven spikes (the 2022 lithium episode), supply-chain disruptions, and discrete policy shifts such as subsidy introductions and withdrawals. Most recently, the 2025 expiration of the US federal EV tax credit was followed by a pre-deadline surge and post-expiry collapse in US EV sales, consistent with the delay-and-spike cycle analysed in Chapter 5. Moreover, the two sources of uncertainty are correlated: energy market shocks propagate into both the running benefit and the cost of adoption.

Three bodies of literature bear on this problem, but none addresses it in full. An empirical literature documents the stochastic behaviour of energy prices and technology costs. An applied fleet-transition literature models adoption with dynamic programming, stochastic optimisation, and real options (Kleindorfer et al., 2012; Falbo et al., 2021), but treats adoption costs as constant or deterministically declining, leaving the interaction between the two uncertainties unmodelled. A mathematical literature on optimal stopping and free-boundary problems (Peskir and Shiryaev, 2006; Oksendal, 2013; Brennan and Schwartz, 1985; Insley, 2002; Szimayer and Maller, 2007) provides the formal infrastructure, but has not been applied to the correlated clean-technology adoption problem. Within the clean technology adoption literature, no continuous-time real options model jointly captures stochastic operating benefits, mean-reverting fixed adoption costs, correlated two-factor dynamics, and the separation of the planning horizon from the operational lifetime.

The dissertation closes this gap. To the best of my knowledge, it is the first continuous-time two-factor optimal stopping model of clean technology adoption in which the unit cost advantage follows an arithmetic Brownian motion, the fixed adoption cost follows a correlated mean-reverting Ornstein–Uhlenbeck process, and finite planning horizons are treated separately from finite operational lifetimes. The topic is timely: despite strong policy support and rapid growth in EV sales, the diffusion of clean technologies has remained slower and more uneven than stated decarbonisation targets imply, and diagnosing whether adopters are waiting rationally or facing structural barriers determines which policy instruments can accelerate adoption.

The central research question is: *when both the unit cost advantage and the fixed upfront cost of a clean technology are stochastic and correlated, what is the optimal switching rule, and how do finite decision horizons, cost volatility, and expected cost convergence rank as barriers to adoption?* Three subsidiary questions structure the analysis. First, what closed-form guidance does the single-factor benchmark provide, and how do anticipated discrete policy or technology shocks alter the optimal rule? Second, how must the solution concept and the numerical methodology change when the adoption cost becomes a correlated mean-reverting state variable in its own right? Third, which barrier binds for a calibrated, high-utilisation commercial adopter, and what does the answer imply for the design of adoption subsidies?

The dissertation is organised in five chapters. Chapter 1 motivates the problem and surveys the three literature strands. Chapter 2 develops the single-factor analytical benchmark, including the jump extension and its empirical application across three US metropolitan markets. Chapter 3 introduces deterministic cost decline as an analytical intermediate, then constructs the two-factor model, characterises the value function within the variational-inequality and Snell-envelope framework, and develops the finite-difference solution methodology. Chapter 4 calibrates the framework and conducts the cross-scenario numerical analysis. Chapter 5 concludes and draws out the policy implications, relates the model's predictions to the 2025 US policy cliff episodes, and discusses limitations and future research.

II. The Methods Used

The dissertation develops its framework from closed-form analytics to numerics, with each method chosen to match the structure of the problem it solves.

Closed-form analysis of the single-factor benchmark. The unit cost advantage D_t follows an arithmetic Brownian motion, $dD_t = \alpha\, dt + \sigma\, dW_t$, and the payoff accrues at usage intensity λ. For constant adoption cost, explicit formulas for the value function, the optimal threshold x^{**}, and the expected waiting time are derived using Laplace transforms of first-passage times and smooth-pasting conditions, with optimality established by a verification theorem.

Lévy techniques for the jump extension. Discrete technology and policy shocks are modelled as compound Poisson jumps in the fixed adoption cost. Because the adoption payoff aggregates the two state variables linearly, the two-factor jump problem reduces exactly to one-dimensional optimal stopping of a Lévy process. The optimal *closed-loop* adoption boundary $d^*(C)$ is obtained in closed form by specialising the linear-reward Lévy stopping result of Mordecki (2002); for exponential jumps the threshold is explicit in the roots of the Lévy exponent (Kou, 2002). The result is reported for two calibrated specifications: a Poisson stream of exponentially distributed cost reductions (technological breakthroughs) and a recurring hazard of a fixed cost increase (subsidy withdrawal). All jump results reported below are closed-loop. Chapter 2 additionally analyses a *precommitted* first-passage benchmark, which fixes a trigger on D at time zero and underlies the published article on which the chapter draws; comparing the two rules isolates the economic value of reacting to realised cost jumps.

Deterministic cost decline as an analytical intermediate. Before introducing cost volatility, an exponentially declining cost path is solved in closed form: Lambert-W yields the ex-ante target from any starting state, and the operational boundary is linear in the current cost.

Infinite-horizon two-factor model. The closest antecedent is McDonald and Siegel (1986), in which both the project value and the investment cost follow geometric Brownian motions. Here the cost is an Ornstein–Uhlenbeck process (Uhlenbeck and Ornstein, 1930) rather than geometric, and the benefit is an arithmetic running payoff rather than a geometric present value, so scale-invariance fails and the free boundary is a curve in the two-dimensional state space rather than a scalar threshold. The dynamics of the joint state are

$$dD_t = \alpha\, dt + \sigma\, dW_t^D, \qquad dC_t = \kappa(\theta - C_t)\, dt + \eta\, dW_t^C,$$

where the Brownian motions satisfy $d\langle W^D, W^C\rangle_t = \rho\, dt$. Adoption at a stopping time τ exchanges the cost C_τ for the discounted benefit stream, so the value function is

$$V(c,d) = \sup_{\tau \ge 0} \mathbb{E}_{c,d}\left[-e^{-r\tau} C_\tau + \int_\tau^\infty e^{-rs} \lambda D_s\, ds\right].$$

The value function is stationary and solves the elliptic variational inequality

$$\max\left\{ (\mathcal{L} - r)V(c,d),\ \hat{g}(c,d) - V(c,d) \right\} = 0,$$

where $\hat{g}(c,d) = \lambda d/r + \lambda\alpha/r^2 - c$ is the immediate-adoption payoff and $\mathcal{L} = \kappa(\theta - c)\,\partial_c + \alpha\,\partial_d + \frac{1}{2}\eta^2\,\partial_{cc} + \frac{1}{2}\sigma^2\,\partial_{dd} + \rho\eta\sigma\,\partial_{cd}$ is the generator of the joint process. The continuation equation is homogeneous (benefits accrue only after adoption), and the optimal rule is a time-independent free boundary. The elliptic problem is discretised with finite differences and solved by projected successive over-relaxation (PSOR) on the spatial grid; accuracy is checked by grid and iteration convergence and Richardson extrapolation.

Finite-horizon two-factor model. A second model imposes a finite planning window T and a finite operational lifetime L. The payoff is a finite annuity rather than a perpetuity, the value function is time-dependent, and the optimal rule is a time-dependent free boundary. The associated parabolic variational inequality is solved by backward induction, nesting PSOR at each time step; the extracted boundary is cross-checked against Longstaff–Schwartz Monte Carlo simulation (Longstaff and Schwartz, 2001).

Empirical calibration. Both two-factor models are calibrated to a high-utilisation delivery fleet's EV replacement decision using cross-market parameter synthesis: benefit dynamics estimated from US energy price data (2015–2023), cost dynamics estimated from UK EV–ICE vehicle price differentials (Nissan Leaf vs Ford Focus, 2018–2024) via maximum likelihood on the interpolated latent differential, and correlation assigned on economic grounds with sensitivity analysis over $\rho \in \{-0.3, 0, +0.3\}$. Eight scenarios (A–H) spanning three US metropolitan markets, jump expectations, deterministic decline, stochastic costs, and finite horizons connect the model hierarchy to observable adoption decisions.

III. Scientific Results of the Dissertation

The scientific results are organised as five theses. Results 1 and 2 correspond to the analytical benchmark of Chapter 2, drawing on the author's co-authored publication listed in Section V; Results 3 and 4 correspond to the two two-factor models of Chapters 3 and 4 (an infinite-horizon problem and a finite-horizon problem, each with its own numerical solution); Result 5 corresponds to the policy analysis of Chapter 5. All quantitative statements refer to the calibrated delivery-fleet case study (usage intensity $\lambda = 480$ hundred-miles per year, equivalently 48,000 miles per year, discount rate $r = 3\%$, current cost premium $C_0 = \$5{,}800$). Table 1 collects the optimal thresholds and decisions across the eight calibrated scenarios, and Figure 1 displays the corresponding boundaries in the state space.

Scenario	**Specification**	**Threshold ($/100mi)**	**Decision**
Constant and jump costs (Chapter 2)			
A	San Francisco, constant cost	6.08	Wait
B	Seattle, constant cost	6.08	Adopt
C	Seattle, anticipated cost declines	11.66	Wait
D	Miami, constant cost	6.08	Wait
E	Miami, anticipated subsidy withdrawal	4.68	Adopt
Stochastic mean-reverting costs (Chapters 3–4)			
F	Deterministic cost decline ($\eta = 0$)	7.31	Wait
G	Stochastic OU cost, infinite horizon	7.93	Wait
H	Stochastic OU cost, finite horizon	26.91	Wait

Table 1: Optimal adoption thresholds and decisions across the eight calibrated scenarios. All scenarios share the benefit process parameters and the current cost premium; Scenarios A–E vary the initial benefit by location (San Francisco 1.45, Seattle 7.94, Miami 4.80 $/100mi). The thresholds in Scenarios C and E are the closed-loop boundaries $d^*(C_0)$ under a Poisson stream of exponentially distributed cost reductions and a recurring hazard of a $7,500 cost increase, respectively. Scenario H imposes a 5-year planning window and a 10-year operating life.

Result 1: Closed-form characterisation of adoption timing under constant costs.

- For the empirically relevant specification (arithmetic Brownian motion benefits, usage-dependent payoffs), the dissertation derives explicit formulas for the optimal threshold x^{**}, the value function, and the expected waiting time, and proves via verification theorem that threshold strategies are optimal.

- The optimal threshold stands at 4.34 times what a payback rule would require: with calibrated parameters, $x^{**} = 6.08$ \$/100mi against 1.40 \$/100mi for recovering C_0 over a 10-year operating life. Standard NPV reasoning, which ignores the option to wait, therefore overstates the attractiveness of immediate adoption.
- Usage intensity enters multiplicatively: high-utilisation operators face systematically lower adoption hurdles, explaining why commercial fleets adopt before households.

Result 2: The optimal closed-loop adoption rule under jump risk and a fundamental policy timing asymmetry.

- The jump extension reduces exactly to one-dimensional Lévy optimal stopping; the optimal closed-loop boundary $d^*(C)$ is linear in the prevailing cost with slope r/λ for *any* jump specification: jump risk moves only the intercept.
- Anticipated cost declines (a stream of technological breakthroughs) raise the calibrated boundary from 6.08 to 11.66 \$/100mi, a 92% increase that flips Seattle's decision from adopt to wait; the anticipated withdrawal of the \$7,500 federal credit lowers it to 4.68 \$/100mi, a 23% decrease that flips Miami's decision from wait to adopt. The direction of anticipated policy change matters as much as its magnitude.[1]
- A realised cost cut $\Delta C < 0$ lowers the hurdle on D by $(r/\lambda)|\Delta C|$ at the instant it occurs and triggers adoption on the spot whenever the state crosses the boundary, the immediate-response channel that static triggers cannot capture.

Result 3: The infinite-horizon two-factor model and its numerical solution.

[1]The precommitted first-passage rule, which cannot react to realised cost jumps, is a strict lower bound: its trigger ($x^{***} = 10.50$ in the breakthrough scenario) understates the option value of waiting precisely because the closed-loop rule waits *for* the breakthrough and adopts the moment a realised cost cut lowers the hurdle.

- In the infinite-horizon two-factor model, mean-reverting, correlated fixed upfront costs break scale-invariance and separability; the optimal rule is a stationary curved free boundary $d = b(c)$ in the two-dimensional state space that no single scalar threshold can summarise.

- The associated elliptic variational inequality is discretised with finite differences and solved by PSOR. Grid convergence and Richardson extrapolation establish the accuracy of the extracted boundary; the method extends naturally to other correlated two-factor stopping problems.

- In the infinite-horizon calibration, anticipated convergence and volatility contribute unequally: deterministic cost decline raises the threshold from 6.08 to 7.31 \$/100mi, and cost volatility raises it further to 7.93 \$/100mi under the baseline $\rho = -0.3$; expected convergence contributes roughly twice the increment of pure volatility. Across $\rho \in \{-0.3, 0, +0.3\}$ the threshold varies by approximately 9% (from 7.19 under $\rho = +0.3$ to 7.93).

Result 4: The finite-horizon two-factor model: a strict hierarchy of adoption barriers and its policy taxonomy.

- A second, finite-horizon model replaces the perpetuity with a finite annuity and the stationary boundary with a time-dependent free boundary $d^*(c, t)$. With a 5-year planning window T and a 10-year operational lifetime L, the threshold at the decision date is 26.91 \$/100mi. The incremental contribution of the horizon constraint (18.98 \$/100mi) is roughly fifteen times the infinite-horizon convergence increment (1.23) and thirty times the volatility increment (0.62). The two horizon parameters are not interchangeable: L (the amortisation window) dominates T (the decision window) under realistic fleet-replacement assumptions, so operational lifetime, not decision urgency, is the primary driver.

- The associated parabolic variational inequality is solved by backward induction nested with PSOR; the extracted boundary is cross-checked against Longstaff–Schwartz Monte Carlo.

- This motivates a taxonomy of adoption barriers. The infinite-horizon threshold of 7.93 \$/100mi represents the *strategic barrier*: the firm has perpetual time to act and chooses to wait, leveraging EV–ICE price volatility and mean reversion to optimise its entry point. The finite-horizon threshold of 26.91 \$/100mi represents the *structural barrier*: the annuity factors available over a ten-year operating life capture only a fraction of the infinite-horizon value, so the threshold explodes irrespective of option value considerations. The taxonomy suggests a hypothesis about the longstanding energy efficiency gap, that its origin is structural rather than strategic, with firms lacking the operational horizon rather than miscalculating option value.

- The policy implication is diagnostic: for short-horizon adopters a uniform purchase subsidy lowers the premium without extending the window over which it is amortised, so it is necessary and insufficient when the annuity penalty is the binding constraint, whereas lifetime-extending instruments (transferable battery warranties, second-life value guarantees, vehicle-to-grid revenue access) act on that constraint directly.

Result 5: Policy timing asymmetry and the architecture of adoption subsidies.

- A subsidy with a firm expiry date is economically analogous to an anticipated positive jump in the adoption cost: the model predicts that announcement shifts thresholds immediately, pulling adoption forward into the pre-deadline window and collapsing it thereafter. The correspondence is not exact: the stationary Poisson specification places the jump at an exponentially distributed time rather than at a known date.

- Market outcomes around the 2025 US federal credit expiry and Germany's 2023 Umweltbonus termination are consistent with this prediction. US EV sales reached 458,298 units in Q3 2025 ahead of the September deadline, 22% above a year earlier, then fell 43%

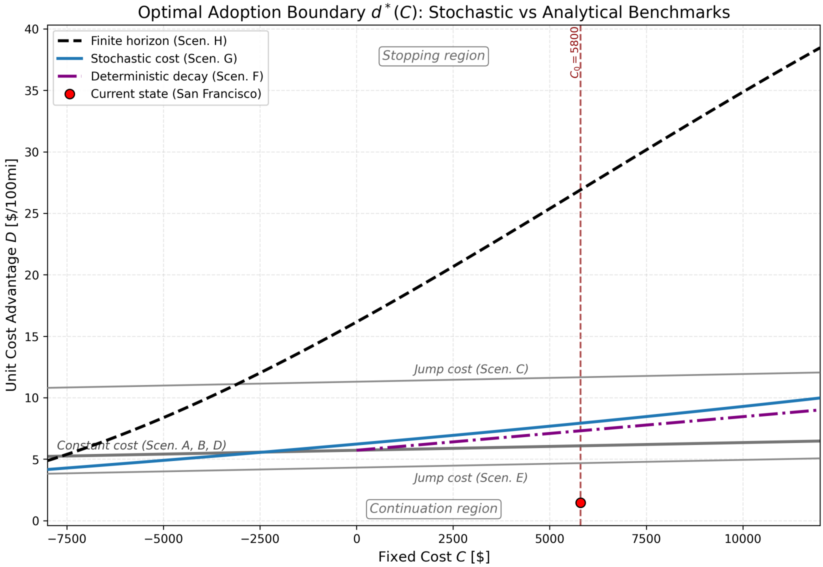

Figure 1: Optimal adoption boundaries across model specifications in the (C, D) state space. The constant-cost benchmark (Scenarios A, B and D) is flanked by the closed-loop jump boundaries: anticipated cost declines (Scenario C) raise the boundary while the threat of subsidy withdrawal (Scenario E) lowers it. Deterministic cost decline (Scenario F) steepens the boundary, cost volatility (Scenario G) shifts it up further, and the finite horizon (Scenario H) steepens it most, the annuity penalty amplifying the sensitivity to the prevailing cost. The red dot marks the calibrated current state $(C_0, D_0) = (5{,}800,\ 1.45)$.

below the prior year in Q4; German battery-electric registrations fell 27% in 2024. Both series and their sources are documented in Chapter 5. The episodes are consistent with the mechanism rather than identifying it, since model launches, inventory positioning and the general 2025 sales cycle moved in the same direction over the same window. Recent Hungarian and EU purchase-grant schemes follow the same time-limited, budget-capped design. Volume-based or algorithmic phase-outs would distribute the adjustment across adopter cohorts.

IV. Main References

Merrill Jones Barradale. Impact of public policy uncertainty on renewable energy investment: Wind power and the production tax credit. *Energy Policy*, 38(12):7698–7709, 2010. doi: 10.1016/j.enpol.2010.08.021. URL `https://www.sciencedirect.com/science/article/pii/S0301421510006361`.

Michael J Brennan and Eduardo S Schwartz. Evaluating natural resource investments. *Journal of Business*, pages 135–157, 1985. doi: 10.1086/296288.

Avinash K Dixit and Robert S Pindyck. *Investment under uncertainty*. Princeton University Press, 1994.

Nicole El Karoui. Les aspects probabilistes du contrôle stochastique. In *Ecole d'Eté de Probabilités de Saint-Flour IX-1979*, volume 876 of *Lecture Notes in Mathematics*, pages 73–238. Springer, 1981.

Paolo Falbo, Giorgio Ferrari, Giorgio Rizzini, and Maren Diane Schmeck. Optimal switch from a fossil-fueled to an electric vehicle. *Decisions in Economics and Finance*, 44:1147–1178, 2021. doi: 10.1007/s10203-021-00359-2.

Margaret Insley. A real options approach to the valuation of a forestry investment. *Journal of environmental economics and management*, 44 (3):471–492, 2002. doi: 10.1006/jeem.2001.1209.

Murat Isik. Incentives for technology adoption under environmental policy uncertainty: Implications for green payment programs. *Environmental and Resource Economics*, 27(3):247–263, 2004. doi: 10.1023/B:EARE.0000017624.07757.3f. URL `https://link.springer.com/content/pdf/10.1023/B:EARE.0000017624.07757.3f.pdf`.

Paul R Kleindorfer, Andrei Neboian, Alain Roset, and Stefan Spinler. Fleet renewal with electric vehicles at la poste. *Interfaces*, 42(5): 465–477, 2012. doi: 10.1287/inte.1120.0640.

Steven G. Kou. A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101, 2002. doi: 10.1287/mnsc.48.8.1086.166.

Francis A Longstaff and Eduardo S Schwartz. Valuing american options by simulation: a simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001. doi: 10.1093/rfs/14.1.113.

Robert McDonald and Daniel Siegel. The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727, 1986. doi: 10.2307/1884175.

Ernesto Mordecki. Perpetual options for Lévy processes in the Bachelier model. *Proceedings of the Steklov Institute of Mathematics*, 237: 247–255, 2002. URL `https://www.cmat.edu.uy/~mordecki/articles/bachelier/`. Russian original: Trudy Mat. Inst. Steklova 237, 256–264.

Björn Nykvist and Mans Nilsson. Rapidly falling costs of battery packs for electric vehicles. *Nature Climate Change*, 5(4):329–332, 2015. doi: 10.1038/nclimate2564.

Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.

Huyên Pham. *Continuous-time stochastic control and optimization with financial applications*, volume 61. Springer Science & Business Media, 2009.

Alexander Szimayer and Ross A. Maller. Finite approximation schemes for Lévy processes, and their application to optimal stopping problems. *Stochastic Processes and their Applications*, 117(10):1422–1447, 2007. doi: 10.1016/j.spa.2007.01.012.

George E Uhlenbeck and Leonard S Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930. doi: 10.1103/PhysRev.36.823.

Micah S. Ziegler and Jessika E. Trancik. Re-examining rates of lithium-ion battery technology improvement and cost decline. *Nature Energy*, 6(9):871–878, 2021. doi: 10.1038/s41560-021-00864-5.

V. List of Own Publications on the Topic

- Szabó, D. Z., Csóka, P., & Janosik, R. (2025): The optimal timing of clean technology adoption: A stochastic cost–benefit analysis. *Technological Forecasting & Social Change*, 219, 124276. DOI: 10.1016/j.techfore.2025.124276