Réka Janosik

Optimal Timing of Clean Technology Adoption under Correlated Price Uncertainty

Doctoral School of Economics and Business Informatics

Supervisor: Dávid Zoltán Szabó, PhD

Corvinus University of Budapest

Doctoral School of Economics and Business Informatics

Doctoral Programme in Business Informatics

Optimal Timing of Clean Technology Adoption under Correlated Price Uncertainty

A Two-Factor Optimal Stopping Approach to Irreversible Investment

Doctoral dissertation

Réka Janosik

Budapest, 2026

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

The transition to clean technologies such as electric vehicles (EVs) is a largely irreversible investment decision under multi-factor uncertainty: prospective adopters weigh an immediate upfront capital cost against a stream of future operational savings dictated by volatile energy markets. This dissertation investigates the optimal timing of adoption by developing a continuous-time real options framework that captures both unit cost advantage and fixed adoption cost uncertainties.

The framework is developed in two stages. First, a single-factor analytical model lets the unit cost advantage follow an arithmetic Brownian motion, with fixed adoption costs either constant or subject to discrete Poisson jumps representing breakthroughs or policy shifts. Closed-form solutions for the optimal threshold, expected waiting time, and value function are derived; the jump extension yields the optimal closed-loop rule. Second, the framework is extended to a two-factor optimal stopping model in which the fixed adoption cost follows a mean-reverting Ornstein–Uhlenbeck process correlated with the unit cost advantage. To my knowledge, this is the first application of such a model in the clean-technology adoption literature. Because scale-invariance and separability break down, the two-dimensional free-boundary problem is solved numerically by an implicit finite-difference scheme with projected successive over-relaxation (PSOR), over infinite and finite horizons.

The models are calibrated to a commercial delivery fleet replacement problem using cross-market parameter synthesis, quantifying how cost volatility, expected convergence, the planning horizon, and the operational lifetime interact. A strict hierarchy of adoption barriers emerges: finite operational horizons (structural barriers) dominate expected cost convergence and cost volatility (strategic barriers) by more than an order of magnitude in the calibrated case. Standard net present value calculations and uniform subsidies may therefore miss the binding constraint of marginal adopters. The jump analysis adds a policy timing asymmetry: anticipated cost declines delay adoption while an anticipated subsidy withdrawal accelerates it, so subsidy timing is a fast lever that cuts both ways. Instruments that extend the effective operational lifetime of EVs address the structural barrier that binds, the more cost-effective lever for horizon-constrained adopters. The asymmetry also suggests phasing subsidies out by volume rather than by date, spreading the adjustment across cohorts rather than concentrating it at one deadline.

Chapter 1

MOTIVATION AND OVERVIEW OF THE THESIS

1.1 Research Motivation

The transition from fossil-based technologies to cleaner alternatives presents a decision-making challenge for individuals, firms, and policymakers alike. Clean technologies, from electric vehicles (EVs) in transport to heat pumps in buildings and solar installations in energy generation, typically require a substantial upfront investment but promise long-term reductions in operating costs and greenhouse gas emissions. The economic case for adoption therefore rests on a comparison between a known upfront cost and uncertain future benefits, both of which evolve stochastically over time. Both energy prices and technology costs are subject to persistent volatility, driven by supply dynamics, geopolitical factors, and the pace of technological learning (Regnier, 2007; Efimova and Serletis, 2014; Dutta et al., 2020). At the same time, the purchase prices of clean technologies, while generally declining due to learning effects and economies of scale, are exposed to raw-material price shocks and supply-chain disruptions that can reverse or stall expected cost trajectories (Nykvist and Nilsson, 2015; Kittner et al., 2017; Ziegler and Trancik, 2021). The resulting tension between acting early and waiting constitutes the classic real options trade-off (McDonald and Siegel, 1986; Dixit and Pindyck, 1994). Resolving this tension under multi-factor uncertainty is the central problem addressed by this dissertation.

Despite strong policy support and rapid recent growth in EV sales, the diffusion of clean technologies has remained slower and more uneven than implied by stated decarbonisation targets (International Energy Agency, 2023). Multiple barriers contribute to this gap: high initial capital costs, technological limitations such as range anxiety and charging infrastructure, and behavioural factors (Shen et al., 2019). The economic case for adoption hinges on uncertain future variables. A prospective fleet operator must weigh a known price premium today against fuel savings that depend on volatile energy markets and a technology cost that may fall further, or experience temporary supply-driven spikes, before settling toward a long-run level. An adopter who acts too early foregoes the option value of waiting for more favourable conditions; one who waits too long foregoes years of operating-cost savings. Formally capturing this trade-off requires a dynamic framework that endogenises the timing decision rather than treating adoption as a one-off static comparison.

This dissertation addresses the optimal timing of clean technology adoption using the framework of real options. The adopter is modelled as holding the right, but not the obligation, to commit at any point in time to an investment whose cost is largely sunk, and the objective is to characterise the conditions under which immediate adoption becomes preferable to continued waiting. The option-value argument requires only partial irreversibility, that reversing the decision does not recover the outlay (Dixit and Pindyck, 1994, Ch. 2), a condition met here by transaction costs and steep early depreciation; the model takes the fully irreversible limiting case. The analysis proceeds in two stages, reflecting a deliberate progression in model complexity. The first stage, presented in Chapter 2, derives closed-form solutions for the case in which the unit cost advantage of the clean technology follows a stochastic process but the fixed adoption cost is constant. This single-factor model was originally developed in Szabó et al. (2025) and is presented here, with the permission of the co-authors, in full. It derives closed-form characterisations of the optimal switching rule, the value function, and the expected waiting time. These results are independently applicable to policy design and subsequently serve as the foundation on which the two-factor extension is built. The second stage, presented in Chapter 3, relaxes the constant-cost assumption by allowing the fixed adoption cost itself to evolve as a mean-reverting stochastic process. This two-factor formulation no longer admits a closed-form solution and is therefore solved numerically using finite-difference methods. Together, the two stages yield closed-form adoption rules when costs are fixed and a quantification of the incremental option value generated when costs are themselves stochastic. Finally, Chapter 4 grounds these methodological advancements in an empirical application. By applying the two-factor model to empirical data on vehicle prices, energy costs, and fleet usage patterns, this chapter numerically analyses how finite planning horizons, cost volatility, and expected cost convergence interact to determine the optimal replacement strategy for a commercial delivery fleet.

1.2 Literature review

The economics of clean technology adoption sits at the intersection of three bodies of work that have developed largely in parallel. The first is an empirical literature documenting that the two key financial determinants of adoption timing, the operating cost advantage of the new technology and its upfront purchase price, are both stochastic: energy price spreads are volatile and non-stationary, while clean technology prices exhibit secular decline punctuated by commodity-driven spikes and policy-induced jumps. The second is an applied literature that has modelled fleet electrification and related transitions using dynamic programming, stochastic optimisation, and, most recently, real options; this literature captures the irreversibility and sequential nature of adoption decisions but has generally treated adoption costs as constant or deterministically declining, leaving the interaction between the two sources of uncertainty unmodelled. The third is a formal mathematical literature on optimal stopping and free-boundary problems, which supplies the tools the two-factor problem needs, free-boundary characterisation and finite-difference solvers, where no closed-form solution exists. This review traces each strand and shows that while empirical studies document stochastic dynamics in both operating benefits and technology costs, the applied adoption literature typically models only one of these sources of uncertainty, while the optimal stopping literature provides the analytical and numerical tools required to analyse the resulting multi-factor investment problem. A final subsection surveys adjacent literatures, including end-of-life vehicle management, market entry dynamics, and grid infrastructure, that inform the broader context of clean technology transitions without contributing formal modelling elements. The gaps identified in this review motivate the framework developed in Chapters 2 and 3.

Table 1 organises the surveyed studies along two independent dimensions: uncertainty structure (single-factor, two-factor, jump) and application domain (EV fleet, solar, heat pump, grid/storage), and records whether each study employs analytical, numerical, or empirical methods. The five row groupings correspond to the subsections that follow. Read column-wise, the table documents that single-factor analytical models dominate the early literature; read row-wise, it reveals that applications have broadened from generic irreversible investment toward sector-specific adoption problems. The joint pattern constitutes the defining trajectory of the field: progressive enrichment of the stochastic specification, accompanied by a shift from closed-form to numerical solution as models are brought closer to empirical calibration.

	Model			**Application**				**Method**
Study	**1F**	**2F**	**J**	**EV**	**S**	**HP**	**G**	**A/N**
1.2.1 The Adoption Timing Problem and Its Gaps								
Jovanovic (1982)								Q/E
McDonald and Siegel (1986)	✓	✓						A

Continued on next page...

Table 1 – continued from previous page

	Model			Application				Method
Study	**1F**	**2F**	**J**	**EV**	**S**	**HP**	**G**	**A/N**
Pindyck (1991)	✓							A
Dixit and Pindyck (1994)	✓							A
Abel and Eberly (1996)	✓							A
Caballero (1999)	✓							A
Geroski (2000)								Q/E
Acemoglu et al. (2012)								Q/E
Aghion et al. (2016)				✓				Q/E
Trigeorgis and Tsekrekos (2018)	✓							A
Allcott and Kessler (2019)								Q/E
1.2.2 Fleet Transition and Clean Technology								
Kleindorfer et al. (2012)		✓		✓				A/N
Feng and Figliozzi (2013)				✓				N
Ansaripoor et al. (2016)		✓		✓				N
Pelletier et al. (2019)				✓				N
Shen et al. (2019)								Q/E
Harajli et al. (2020)					✓			E
Knobloch et al. (2020)				✓		✓		E
Falbo et al. (2021)	✓			✓				A/N
Alp et al. (2022)				✓				N
Liu et al. (2023)		✓		✓				A/N
Winkelmann et al. (2024)		✓		✓				N
1.2.3 Stochastic Drivers: Energy Prices, Costs and Policy								
Hirschmann (1964)								E
Beck (1994)								E
Schwartz and Smith (2000)		✓						A/N
Isik (2004)	✓		✓					N
Regnier (2007)								E
Meade (2010)								E
Jacobsen and Zvingilaite (2010)								E
Barradale (2010)					✓			N
Weiss et al. (2012)				✓				E
Efimova and Serletis (2014)								E
Ketterer (2014)								E
Nykvist and Nilsson (2015)				✓				E

Continued on next page...

Table 1 – continued from previous page

	Model			Application				Method
Study	**1F**	**2F**	**J**	**EV**	**S**	**HP**	**G**	**A/N**
Weiss et al. (2015)				✓				E
Adkins and Paxson (2016)		✓	✓		✓			N
Kittner et al. (2017)								E
Safari (2018)				✓				E
Weiss et al. (2019)				✓				E
Babich et al. (2020)		✓			✓			A/N
Dutta et al. (2020)								E
Ziegler and Trancik (2021)				✓				E
Goetzel and Hasanuzzaman (2022)				✓				E
Ramani et al. (2022)				✓				E
Slameršak et al. (2022)				✓				Q/E
Wang et al. (2023)								E
BloombergNEF (2024)				✓				I
International Energy Agency (2023)				✓				I
ING Research (2025)				✓				I
U.S. Geological Survey (2025)				✓				I
1.2.4 Methodological Foundations								
Cryer (1971)								PSOR
Brennan and Schwartz (1978)	✓		✓					FDM
Brennan and Schwartz (1985)		✓						FDM
Longstaff and Schwartz (2001)	✓	✓	✓					LSM
Glasserman (2003)	✓	✓						LSM
Insley (2002)	✓							FDM
Kou (2002)	✓		✓					A
Mordecki (2002a)	✓		✓					A
Dayanik and Karatzas (2003)	✓							A
Kou and Wang (2003)	✓		✓					A
Stentoft (2004)	✓	✓						LSM
Detemple (2006)	✓	✓						N (FDM)
Peskir and Shiryaev (2006)	✓							A
Wilmott (2006)	✓	✓	✓					FDM
Øksendal and Sulem (2007)	✓	✓	✓					A
Szimayer and Maller (2007)	✓		✓					FDM
Pham (2009)	✓							A
Daming et al. (2014)	✓		✓					A

Continued on next page...

Table 1 – continued from previous page

	Model			**Application**				**Method**
Study	**1F**	**2F**	**J**	**EV**	**S**	**HP**	**G**	**A/N**
Borodin and Salminen (2015)	✓							A
Seydel (2017)	✓	✓						FDM
1.2.5 Applied and Peripheral Literature								
Cowan and Hultén (1996)				✓				Q/E
Wesseling et al. (2014)				✓				Q/E
Nijhuis et al. (2015)							✓	Q/E
Kitapbayev et al. (2015)	✓						✓	LSM
Szabó and Martyr (2017)	✓						✓	A
Soo et al. (2017)				✓				Q/E
Abid et al. (2019)								E
Saidani et al. (2019)				✓				Q/E
Thomas and Maine (2019)				✓				Q/E
Karagoz et al. (2020)				✓				Q/E
Soo et al. (2021)				✓				Q/E
Cho and Shin (2022)				✓				Q/E
Demartini et al. (2023)				✓				N
Dagher and Hasanov (2023)								E
Urom et al. (2023)								E
Dagher et al. (2023)								E
Fang et al. (2024)				✓				A/N
Pata et al. (2024)						✓		Q/E
Present Dissertation	✓	✓	✓	✓	~	~	~	**A/FDM**

Notes: **1F**: 1-Factor; **2F**: 2-Factor; **J**: Jump; **EV**: Electric Vehicles & Fleet Transition ; **S**: Solar or other Renewables; **HP**: Heat Pumps; **G**: Grid & Storage; **A**: Analytical; **N**: Numerical (FDM/PSOR/LSM); **E**: Empirical; **Q/E**: Qualitative/Econometric. **I**: Industry/Institutional Report. ✓: applicable model structure or application domain; ~: Theoretically extendable generalised framework.

Table 1: Comprehensive classification of the surveyed studies on clean technology adoption, applied studies, related methodologies and other literature. For formal optimal stopping models, checkmarks indicate the explicitly modelled uncertainty structure. For empirical and policy studies, structural checkmarks are omitted to reflect their qualitative or econometric methodologies.

1.2.1 The Adoption Timing Problem and Its Gaps

While aggregate diffusion models characterise adoption as an epidemic or selection process at the population level (Geroski, 2000; Jovanovic, 1982), they do not endogenise the individual timing decision under price uncertainty. Traditional capital budgeting is equally limited: the Net Present Value criterion prescribes investment when expected benefits equal or exceed the sunk cost, implicitly assuming a now-or-never opportunity (Dixit and Pindyck, 1994) and thereby neglecting the option value of deferral. As established by McDonald and Siegel (1986), surveyed by Pindyck (1991), and formalised for a broader class of processes by Dixit and Pindyck (1994, Ch. 5–6), three features jointly necessitate a departure from the classical NPV rule: irreversibility, stochastic payoffs, and timing flexibility. Trigeorgis and Tsekrekos (2018) provide a comprehensive review of real options methodology in the operations research literature, documenting the range of analytical and numerical approaches available for different process specifications.

The real options approach extends neoclassical investment theory by redefining the user cost of capital. In the framework of Abel and Eberly (1996), the marginal benefit of an investment must exceed the marginal cost by a shadow price equal to the marginal value of preserving the option to delay. This provides a microeconomic explanation for the observed inertia in clean technology adoption: agents do not respond to price signals until the unit cost advantage is substantially larger than what NPV would suggest. When an agent invests, they do not merely exchange capital for an asset; they also extinguish the option to wait for further information. This lost option value represents an opportunity cost that must be included in the full cost of investment. The optimal exercise of a real option is therefore governed by an augmented hurdle rate.

This inertia is reinforced by the presence of legacy capital. An agent considering EV adoption typically owns an existing ICE vehicle whose purchase cost is sunk. However, as discussed in vintage capital models (Caballero, 1999), the replacement decision involves abandoning the residual service life of the current asset. Even under certainty, the effective switching cost therefore exceeds the sticker price of the new technology. When future costs and operating savings are also uncertain, an additional incentive to wait arises: formal optimal-stopping models show that adoption thresholds systematically exceed those predicted by static NPV calculations, because they endogenise *when*, not just whether, to invest. Within the applied fleet-transition literature, real options models have treated fixed adoption costs as constant, modelling operating benefits as the sole stochastic driver.

Adoption costs are themselves subject to uncertainty. As reviewed in Subsection 1.2.3, EV purchase prices have exhibited substantial volatility, driven by battery costs, supply conditions, and policy shifts. When both the unit cost advantage and the fixed adoption cost evolve stochastically, the optimal stopping boundary becomes a surface over two correlated state variables rather than a single threshold. The closest predecessors in the fleet electrification literature, reviewed in Subsection 1.2.2, do not solve this continuous-time two-factor problem.

The urgency of addressing these methodological gaps is reinforced by the broader policy context. Acemoglu et al. (2012) demonstrate that delaying policy interventions in the climate context significantly increases the economic burden of the subsequent technological transition. Analysing patent data in the automotive sector, Aghion et al. (2016) find that while carbon taxes successfully stimulate clean technology development, the strength of this response is heavily dictated by a firm's prior innovation trajectory and historical investments. Allcott and Kessler (2019) show that while behavioural interventions encourage clean technology uptake, conventional welfare estimates often exaggerate these societal gains by omitting the hidden compliance costs borne by consumers.

1.2.2 Fleet Transition and Clean Technology

The transition of fleet vehicles to EVs has been studied extensively, with contributions spanning deterministic optimisation, stochastic programming, and dynamic programming before reaching the real options frontier. Early discrete-time formulations, such as the stochastic dynamic programming model developed by Kleindorfer et al. (2012), examine how fuel and battery price uncertainties affect the optimal rollout of electric delivery vehicles for La Poste's mail and parcel distribution. Evaluating the transition strictly on a deterministic basis, Feng and Figliozzi (2013) calculate that battery-electric commercial vehicles reach financial parity with diesel trucks only when annual utilisation exceeds 16,000 miles per vehicle. To manage the financial risks of fleet replacement, Ansaripoor et al. (2016) construct a multi-stage stochastic program that derives the optimal lease volumes and vehicle mix for commercial operations in the UK. They frame the problem as constrained programming rather than optimal stopping, yielding an optimal fleet mix rather than an adoption threshold.

Expanding this operational focus to public transit, Pelletier et al. (2019) apply integer linear programming to coordinate bus procurement and salvage decisions alongside charging infrastructure investments. Their case studies use French data for various EV types with different charging configurations. Similarly, Alp et al. (2022) use integer linear programming to show that firms operating in high-density transport networks can achieve significant emission reductions and operational savings by jointly planning EV adoption and charging station deployments. More recently, Winkelmann et al. (2024) apply approximate dynamic programming to German logistics data, concluding that diesel trucks will remain the standard for long-haul routes, with battery-electric alternatives only becoming competitive for short-distance transport over the medium term. The closest methodological precursor to this dissertation is Falbo et al. (2021), who introduce continuous-time real options to the EV adoption problem with a case study for the Italian market. By specifying the unit cost advantage as a mean-reverting process, they characterise the optimal switching boundary via parabolic cylinder functions, which they subsequently evaluate using numerical methods; their formulation remains single-factor (adoption costs are constant), so the interaction between benefit and cost uncertainty does not arise.

Liu et al. (2023) extend this strand by modelling EV investment under China's dual-credit policy within a real options framework, deriving investment timing thresholds under joint uncertainty in credit prices and fuel-vehicle demand. Both variables follow stochastic diffusion processes, yielding a two-factor investment timing problem with analytical characterisation of the optimal threshold. However, even these extensions maintain a diffusion specification for both stochastic drivers and therefore do not capture the empirically observed mean-reverting dynamics of technology costs. While fleet electrification provides the primary empirical testing ground for the models developed in this dissertation, the underlying optimal stopping mechanics extend directly to other decentralised clean technology adoption decisions. Residential heat pump adoption involves a structurally identical dual-uncertainty profile: the irreversible upfront investment must be weighed against an uncertain spread between electricity and natural gas prices, generating an option value of waiting that static payback calculations cannot capture (Knobloch et al., 2020). The adoption of residential and commercial solar photovoltaic installations presents a parallel structure, combining stochastic electricity prices with declining but volatile panel and installation costs. Harajli et al. (2020) assess the financial viability of commercial solar PV installations across Palestine, Lebanon, and Iraq, finding that adoption economics vary substantially with local energy price levels and policy regimes, confirming that a flexible stochastic framework is necessary to evaluate adoption timing across heterogeneous national contexts.

The diversity of national contexts, data sources, and usage profiles across these studies underscores that adoption economics are local, yet the modelling frameworks employed share a common structural limitation. While a stochastic investment cost has been treated in other contexts, most directly by McDonald and Siegel (1986), the applied literature on fleet electrification and clean-technology adoption has generally treated adoption costs as constant or deterministic.

1.2.3 Stochastic Drivers: Prices, Costs, and Policy

The stochastic environment relevant for clean technology adoption arises from three sources: energy price dynamics, technology cost evolution, and policy uncertainty. The fossil-fuel and electricity prices that jointly determine the unit cost advantage of clean technology adoption are empirically documented stochastic processes subject to persistent volatility. Regnier (2007) demonstrates that crude oil price volatility is a persistent feature of energy markets rather than an anomaly, and that this volatility transmits substantially to end-consumer fuel prices. Jacobsen and Zvingilaite (2010) and Ketterer (2014) demonstrate that the integration of intermittent renewable energy into the grid reduces the wholesale electricity price level through the merit-order effect, but simultaneously increases price volatility. Efimova and Serletis (2014) document that energy market prices exhibit substantial variation at daily and intraday frequencies. Dutta et al. (2020) establish that elevated implied volatility in the energy sector is associated with depressed clean energy asset returns, identifying a transmission channel from fossil-fuel price risk to clean technology investment performance. These empirical properties have direct implications for process specification in adoption timing models. Meade (2010) demonstrates that neither geometric Brownian motion nor mean-reverting processes yield superior long-term oil price forecasts, while cointegration properties documented in commodity markets (Beck, 1994) support arithmetic Brownian motion as the baseline diffusion specification for the unit cost advantage process.

Investment costs for clean technologies exhibit persistent downward trends driven by learning-by-doing and manufacturing scale economies. Supply shocks, material price volatility, and demand fluctuations may prevent these trends from operating predictably, creating persistent uncertainty around expected convergence paths. This distinction matters for adoption timing: deterministic decline creates incentives to wait for lower costs, while stochastic volatility around the trend generates option value that further delays investment. The electric vehicle context exhibits both features. Battery costs have fallen substantially since 2010 (BloombergNEF, 2024), yet the 2021–2023 period saw lithium carbonate prices spike tenfold before partially correcting, demonstrating that volatility persists alongside the secular trend.[1]

The learning curve framework relates unit costs to cumulative production through power laws that capture cost reductions from manufacturing experience. The learning rate (i.e. the percentage cost reduction per doubling of cumulative output) provides a reduced-form measure of technological progress (Hirschmann, 1964). Applied to battery packs, this framework has proved successful: Nykvist and Nilsson (2015) documented learning rates of 6–9% through 2014, while Ziegler and Trancik (2021) found rates of 20–27% when accounting for simultaneous improvements in cell-level energy density. BloombergNEF's widely-cited tracking suggests volume-weighted pack prices fell from over \$1,100/kWh in 2010 to a record low of \$115/kWh by end-2024, a roughly 90% decline consistent with sustained learning (BloombergNEF, 2024).

Learning rates vary substantially across electrified vehicle segments. Weiss et al. (2012) document HEV price-differential learning rates of $23 \pm 5\%$, while Weiss et al. (2015) find a comparatively modest rate of 8% for e-bike prices, reflecting the simpler

[1] U.S. Geological Survey, *Mineral Commodity Summaries 2025*, `https://pubs.usgs.gov/periodicals/mcs2025/mcs2025-lithium.pdf`.

powertrain and more mature bicycle manufacturing base. For BEV passenger cars, Safari (2018) establishes learning rates of $9 \pm 2\%$ for total price and $15 \pm 1\%$ for electrification costs specifically. Weiss et al. (2019) subsequently find a whole-vehicle learning rate of $23 \pm 2\%$ for the German market using updated retail data, with correspondingly faster convergence (half-life: 2.6 years). Applying this framework to segment-level heterogeneity, Goetzel and Hasanuzzaman (2022) document learning rates ranging from 23% (small cars) to 29% (mid-size vehicles), with implied half-lives of roughly 3 to 6 years depending on vehicle class and production growth assumptions. Their cross-sectional analysis reveals that convergence is driven not only by EV cost reductions but also by rising prices of comparable ICE vehicles in the same segments due to emissions compliance mandates.

These learning-curve studies establish that EV–ICE price convergence operates on timescales of roughly 2.5 to 6 years under baseline production growth assumptions, providing empirical grounding for model calibration in adoption timing frameworks. However, the learning-curve methodology measures expected convergence by regressing costs against cumulative production, abstracting from time-varying volatility around the trend. Temporary supply shocks (e.g., lithium price spikes (U.S. Geological Survey, 2025) and semiconductor shortages that constrained vehicle production across major manufacturers in 2021–2022 (Ramani et al., 2022)), policy shifts (subsidy expirations, tariff adjustments), and manufacturer pricing strategies (competitive discounting, model refresh cycles) create deviations from smooth learning curves that can be economically significant for adoption decisions. The coexistence of persistent downward trends and substantial short-term volatility motivates stochastic specifications that combine mean reversion toward declining cost targets with diffusion around the expected path.

Dixit and Pindyck (1994) treat input-cost uncertainty as exogenous to the investing firm and show that it generates an option value of waiting that is qualitatively distinct from benefit-side uncertainty. Schwartz and Smith (2000) decompose a single stochastic commodity price into transitory and persistent components, demonstrating that short-term mean-reverting shocks superimposed on a long-term equilibrium trend require a two-factor specification that single-factor models cannot capture. The lithium carbonate episode documented above is precisely the type of event this framework is designed to accommodate.

Slameršak et al. (2022) document that the energy investment requirements of the low-carbon transition are so substantial and front-loaded that they create a temporary ‘energy trap’, reducing the net energy available to society, which implies that the timing of capital commitment matters not only for individual adopters but also for aggregate transition trajectories. Beyond technology cost dynamics, a distinct stochastic driver operates through the policy environment, compounding this timing sensitivity. Government subsidies such as purchase rebates, tax credits, and feed-in tariffs act as direct reductions in effective adoption costs. Because these instruments are politically contingent, their introduction or removal creates discrete shifts in the adoption cost process that deterministic learning-curve models cannot capture. Isik (2004) models subsidy arrival and withdrawal as a Poisson jump process and shows that the anticipation of a future subsidy increases the option value of waiting: adopters delay investment to avoid the regret of committing capital just before a cost reduction becomes available. Adkins and Paxson (2016) confirm this mechanism in the renewable-energy context, demonstrating that subsidy uncertainty alone can stall the very investment the policy is intended to promote. Babich et al. (2020) extend this analysis to subsidy design, showing that feed-in tariffs and tax rebates generate different adoption dynamics even when their expected present values are equivalent, because the two mechanisms alter upfront costs and ongoing revenue risk differently. The converse effect is equally documented. When an existing subsidy is expected to expire, adopters face a "use-it-or-lose-it" incentive that compresses the optimal adoption window. Barradale (2010) documents the surge-and-stall pattern in US wind investment around the intermittent renewal of the Production Tax Credit, and attributes it not to the absence of the credit in lapsed years but to uncertainty over its return, transmitted through power purchase agreement negotiations. The distinction matters here: it is anticipation of the policy state, rather than the policy state itself, that moves investment timing. ING Research (2025) document the same mechanism in the EV context: US EV market share surged to 13% in September 2025, immediately preceding the expiration of federal tax credits, before dropping sharply thereafter. Wang et al. (2023) confirm the broader pattern: policy uncertainty reduces green investment, consistent with the real options mechanism whereby an unstable policy environment raises the option value of waiting and further delays adoption. Taken together, these findings establish that policy uncertainty introduces a jump component into the adoption cost process that operates independently of, and can dominate, the continuous dynamics driven by market and technology factors. These studies establish the mechanism but stop short of a decision rule: Isik (2004) and Adkins and Paxson (2016) model subsidy arrival as a Poisson event and trace its effect on the option value of waiting, without characterising the adoption boundary a firm facing such risk should follow.

1.2.4 Methodological Foundations

The real options literature has addressed the computational demands of irreversible investment problems through three broad methodological trajectories: analytical closed-form solutions, simulation-based approximations, and partial differential equation (PDE) schemes. The selection among these approaches is governed by a trade-off between economic transparency and the structural complexity of the problem at hand. Single-factor problems with stationary or geometrically growing state variables admit closed-form solutions, yielding explicit investment thresholds whose dependence on volatility and discount rates can be read off directly (McDonald and Siegel, 1986; Dixit and Pindyck, 1994). These analytical results provide powerful economic intuition but require restrictive structural assumptions (constant or deterministically evolving costs, infinite horizons) that limit their applicability when investment costs are themselves stochastic or when decision horizons are finite.

Analytical tractability can be preserved beyond the canonical geometric Brownian motion specification. Modelling the underlying process as an arithmetic Brownian motion (ABM) permits explicit closed-form threshold characterisations by leveraging known Laplace transforms for first-passage times (Borodin and Salminen, 2015; Dayanik and Karatzas, 2003). Analytical tractability is maintained even when incorporating compound Poisson jump processes to capture discontinuous state transitions or structural shocks (Øksendal and Sulem, 2007; Daming et al., 2014). Together these results deliver explicit solutions for the value function, optimal threshold, and expected waiting time under both non-stationary diffusions and discontinuous jumps.

The jump-augmented specification draws on a more specific body of analytical results. When the state process is a Lévy process and the reward is *linear* in the state, the perpetual optimal stopping problem is solved in closed form: Mordecki (2002a) shows, in the Bachelier setting, that the optimal exercise level equals the strike plus the expected supremum of the process over an independent exponentially distributed horizon, and evaluates that expectation explicitly when the positive jumps are mixed exponential. The explicit first-passage transforms for the double-exponential jump diffusion, including the law of the overshoot when a jump carries the process across a level, are due to Kou (2002) and Kou and Wang (2003).

Three structural features of the two-factor model individually and jointly preclude closed-form solutions. Finite planning horizons eliminate the stationarity on which the elliptic ODE reduction depends, requiring time-stepping through a parabolic PDE (Peskir and Shiryaev, 2006; Detemple, 2006). Mean-reverting cost dynamics break the scale-invariance that permits the standard power-function trial solutions employed in canonical real options frameworks (e.g., Dixit and Pindyck, 1994), so the standard guess-and-verify strategy fails. Correlated state variables introduce a mixed partial derivative term that couples the two spatial dimensions and prevents the value function from separating into independent single-factor components (Wilmott, 2006; Seydel, 2017).

Leveraging the mathematical equivalence between optimal stopping problems and American option pricing, Longstaff and Schwartz (2001) demonstrated that optimal exercise strategies can be evaluated by regressing continuation values on basis functions of the state variables along simulated paths (see also Glasserman, 2003), a technique that scales to higher-dimensional problems and accommodates jump-diffusion processes where grid-based methods become prohibitively expensive. However, Stentoft (2004) documents that while Longstaff–Schwartz Monte

Carlo provides accurate option value estimates, the method can exhibit significant errors in characterising optimal exercise boundaries compared to finite-difference methods. Because the framework developed in this dissertation operates in a two-dimensional state space, finite-difference schemes provide the optimal primary solution method, while Monte Carlo implementations remain valuable as independent validation tools.

For state spaces with $d \leq 3$, finite-difference methods exploit the full PDE structure and directly yield the entire value surface and a sharply defined stopping region, making them both faster and more accurate than simulation-based alternatives (Wilmott, 2006; Seydel, 2017). For problems requiring precise boundary characterisation, Brennan and Schwartz (1978) established the finite-difference methodology for contingent claims subject to jump-diffusion processes. Brennan and Schwartz (1985) extended this framework to natural resource investment, with the commodity spot price as the single stochastic factor, establishing the canonical PDE approach to real options on physical assets. Insley (2002) subsequently extended grid-based methods to natural-resource investments involving mean-reverting prices and finite horizons, demonstrating that PDE schemes can resolve free boundaries with high accuracy when the state space remains low-dimensional. For models incorporating jump-diffusions, Szimayer and Maller (2007) establish finite approximation schemes that preserve the properties of the underlying Lévy processes and prove their convergence for optimal stopping problems.

The theory of optimal stopping and free-boundary characterisation underpinning these numerical methods is developed in Peskir and Shiryaev (2006) and Pham (2009); the projected successive over-relaxation (PSOR) algorithm used to solve the resulting linear complementarity problems traces to Cryer (1971). The fleet-transition studies reviewed in Subsection 1.2.2 have employed several of these approaches but none has developed the full two-factor free-boundary solution.

1.2.5 Applied and Peripheral Literature

The preceding subsections focus on the stochastic drivers and methodological tools directly embedded in the dissertation's models. This final subsection surveys adjacent literatures that inform the broader context of clean technology transitions without contributing formal modelling elements: end-of-life vehicle management, market entry dynamics, grid infrastructure constraints, and cross-sectoral heterogeneity.

Energy price fluctuations propagate across borders and interact with broader macroeconomic conditions, shaping the environment in which adoption decisions are made. Abid et al. (2019) empirically document shocks and contagion transmission from U.S. to Middle East and North Africa equity markets, finding that price fluctuations amplify dependence between the two regions. Dagher and Hasanov (2023) find that financial stress in Asian markets is heavily determined by the origin of the oil shock: oil-specific demand shocks reduce financial stress, whereas supply shocks increase it. Urom et al. (2023) find a time-varying feature in the integration between oil price shocks and interest rates, underscoring that monetary policy and energy markets interact in ways that are relevant to investment timing decisions. The consequences of inaction can be severe. In Lebanon, Dagher et al. (2023) document that 78% of households were unable to maintain adequate heating following the 2019 economic collapse, with most forced to rely on costly private diesel generators. Given Lebanon's substantial solar and wind potential, the case illustrates that the barriers to renewable adoption are as much a matter of policy and investment organisation as of resource availability.

The transition to clean technologies requires systemic strategies for managing obsolete legacy assets. End-of-life vehicle (ELV) management has consequently emerged as a critical research domain, reviewed by Karagoz et al. (2020). The regulatory landscape governing ELV treatment varies substantially across jurisdictions. The European Commission imposes quantitative recycling and recovery targets; Soo et al. (2017) demonstrate the efficacy of this approach by comparing Belgium and Australia, finding that Belgium's strict enforcement of the EU ELV Directive yields substantially superior environmental outcomes relative to Australia, where national disposal legislation is absent. Saidani et al. (2019) extend the comparison to the United States, concluding that while the EU is more advanced in its top-down policy framework, U.S. progress toward vehicle circularity relies primarily on voluntary industry initiatives and the advocacy of individual associations rather than federal mandates. Soo et al. (2021) evaluate the impact of these divergent regulatory regimes on material circularity across Europe, Japan, Australia, and the US, finding that regardless of the regional policy environment, the greatest potential for circularity improvement lies in maximising end-of-process scrap utilisation rates. Complementing this, Fang et al. (2024) analyse the entry mode and reverse-channel structure of a new EV manufacturer facing alternative recycling arrangements for used power batteries, showing that when battery treatment costs are low the third-party recycling channel emerges as the equilibrium outcome, while cooperative recycling yields higher consumer surplus and social welfare, suggesting a potential role for policy incentives in shaping channel choice.

Technological transitions disrupt incumbent market structures and create conditions for new entry. Cowan and Hultén (1996) provide a historical baseline, documenting that in the mid-1990s EV production remained below 2,000 units annually, relying on battery technology essentially unchanged since 1910–20, with no mass-market prospect in sight. The subsequent reversal of this trajectory is partly attributed to competitive dynamics: Wesseling et al. (2014), analysing patent activity in low-emission vehicle technologies over the period 1990–2010, find a positive relationship between market rivalry and continued technological development, with new entrant activity amplifying this effect. Thomas and Maine (2019) examine

Tesla's commercialisation strategy and find that it succeeded not through disruptive low-end entry but through architectural innovation, reconfiguring existing product components around a novel value proposition at the premium segment. Cho and Shin (2022), using fuzzy-set qualitative comparative analysis on 14 EV entrants over 2005–2017, identify configurations of technological and market conditions associated with successful entry, including low-cost high-capacity batteries, competitive vehicle pricing, long driving range, and high safety ratings.

Nijhuis et al. (2015) document the grid and distribution network challenges arising from the energy transition, noting that large-scale EV adoption places substantial demands on electricity infrastructure that existing grid architectures were not designed to accommodate. Pata et al. (2024) reinforce this point at the macro level, finding that R&D investments and ICT adoption accelerate the energy transition in Germany while urbanisation exerts a countervailing effect, underscoring that no single cost or policy trajectory applies universally. Demartini et al. (2023) analyse the supply chain, circular economy, and workforce implications of the EV transition, developing a system dynamics model incorporating agent-based elements to examine interactions among stakeholders across the EV value chain, showing that the transition depends on coordinated changes across industrial and institutional systems. The optimal stopping methodology also extends naturally to adjacent energy assets: Kitapbayev et al. (2015) apply stochastic control to the valuation of thermal storage in combined heat-and-power systems under uncertain electricity and gas prices, solving the resulting dynamic programming problem numerically using a Least-Squares Monte Carlo algorithm. Szabó and Martyr (2017) model decremental balancing reserve contracts for electricity storage as a real options problem with stochastic system imbalance modelled as Brownian motion, deriving a closed-form optimal stopping boundary and demonstrating that the free-boundary structure identified in the EV context recurs across decarbonisation technologies.

1.3 Main Contributions and Outline

This dissertation makes four main contributions to the literature on irreversible investment under uncertainty, to its application to clean technology adoption, and to the design of adoption policy:

(1) Analytical characterisation of optimal adoption timing under constant costs. The literature review identifies that the applied fleet-transition literature requires a process specification combining arithmetic Brownian motion, usage-dependent payoffs, and discrete policy jumps, yet no existing single-factor model solves this joint specification in closed form. The foundational results of McDonald and Siegel (1986) and Dixit and Pindyck (1994) derive closed-form thresholds only for geometric Brownian motion, and in the two-factor case of McDonald and Siegel

(1986) only because the scale invariance of the pair permits reduction to a single ratio; Falbo et al. (2021) adopt mean-reverting dynamics but require numerical evaluation of parabolic cylinder functions. Chapter 2 fills this gap by providing a closed-form baseline for the empirically relevant specification before multi-factor complexity is introduced. The chapter proves via a verification theorem that threshold strategies are optimal, derives explicit expressions for the value function and expected waiting time, and characterises the effect of anticipated subsidy changes on the optimal switching rule. Under Poisson jump risk, the chapter obtains the optimal closed-loop adoption boundary by reducing the two-factor problem to a one-dimensional linear-reward stopping problem and specialising the Lévy result of Mordecki (2002a). The jump extension yields a policy timing asymmetry in closed form: an anticipated subsidy raises the adoption threshold while an anticipated expiry lowers it, so the boundary moves on announcement rather than when the cost itself changes. An empirical application to delivery fleet electrification illustrates the model's practical implications. These analytical results serve a dual purpose: they bear directly on the design and timing of adoption incentives, and they provide the single-factor benchmark against which the additional option value generated by cost stochasticity is isolated in the two-factor extension.

(2) Two-factor optimal stopping framework with mean-reverting costs. The central theoretical gap emerges from the intersection of Subsections 1.2.2 and 1.2.3. The empirical evidence surveyed in Subsection 1.2.3 establishes that clean technology prices exhibit both commodity-driven energy price volatility and supply-chain-driven cost volatility alongside secular decline. Yet the applied fleet-transition models reviewed in Subsection 1.2.2 have not incorporated these joint dynamics within a continuous-time optimal stopping framework. Kleindorfer et al. (2012) model both fuel and battery prices as stochastic within a discrete-time dynamic programme; Liu et al. (2023) solve a continuous-time two-factor problem but specify both drivers as diffusions without mean reversion; and Falbo et al. (2021) solve a continuous-time optimal stopping problem with mean-reverting operating benefits while assuming constant adoption costs. Chapter 3 bridges these approaches by developing, to the best of my knowledge within the clean-technology adoption literature, a continuous-time real options model in which the fixed adoption cost follows a mean-reverting process correlated with the unit cost advantage. Modelling the adoption cost as an Ornstein–Uhlenbeck process captures both short-term volatility from supply-chain shocks and long-run convergence toward production-cost parity. Within the same reviewed literature, the correlation parameter between the two factors is a new element, endogenising the channel through which energy-market shocks propagate into technology supply chains. The lithium carbonate episode and semiconductor shortages suggest this interaction is empirically relevant, yet independent single-factor models cannot represent it. The model accommodates both infinite-horizon formulations, which isolate the pure effect of stochastic costs, and finite-horizon specifications that separate the planning horizon from the operational lifetime, reflecting the distinct roles of the decision window and the asset amortisation period in fleet replacement decisions.

(3) Numerical solution methodology and validation. Subsection 1.2.4 establishes that the two-factor model breaks all three structural assumptions on which closed-form solutions in the single-factor literature rely: stationarity (violated by finite horizons), scale-invariance (violated by mean reversion), and separability (violated by correlated multi-factor dynamics). Chapter 3 therefore develops a finite-difference framework that overcomes all three failure modes simultaneously, solving the stationary variational inequality directly for infinite-horizon problems and employing implicit time-stepping with backward induction for finite horizons. This PDE-based approach directly yields the entire value surface and a sharply defined stopping region, overcoming the limitation of Longstaff–Schwartz Monte Carlo. Grid convergence analysis and Richardson extrapolation establish that discretisation errors remain relatively small, and independent validation against Longstaff–Schwartz confirms numerical accuracy, consistent with the findings of Stentoft (2004). The methodology extends naturally to other optimal stopping problems with correlated two-dimensional diffusions.

(4) Hierarchy of adoption barriers and the resulting policy diagnostic. The preceding contributions establish that planning horizon constraints, cost volatility, mean reversion, and benefit-cost correlation all affect optimal adoption timing, but the existing literature provides no quantification of their relative importance for a given adopter. The empirical fleet-transition studies surveyed in Subsection 1.2.2 either calibrate single-factor optimal stopping models (Falbo et al., 2021) or analyse stochastic fleet renewal decisions within discrete-time dynamic programming frameworks (Kleindorfer et al., 2012), leaving unresolved which constraint binds most tightly in a continuous-time two-factor setting: finite decision windows, cost volatility, or the speed of expected convergence. Chapter 4 addresses this through a case study of delivery fleet electrification, calibrating the two-factor model using cross-market parameter synthesis: UK vehicle price data for fixed adoption cost dynamics and US energy price data for unit cost advantage dynamics, with the correlation structure assigned on economic grounds where joint observations are unavailable. Sensitivity analysis across model specifications establishes a hierarchy of adoption barriers: finite planning horizons dominate anticipated cost convergence, which dominates cost volatility. This motivates a distinction between the *strategic* barrier, where the firm optimally waits because of volatility and expected convergence, and the *structural* barrier, where the remaining lifetime cannot amortise the fixed cost regardless of expectations; in the calibrated case the structural barrier exceeds the volatility effect by more than an order of magnitude. The practical conclusion is a policy diagnostic: the appropriate instrument depends on which barrier binds for the target population. Instruments extending the effective payoff horizon, such as vehicle-to-grid revenues or battery second-life value, target the binding structural barrier, whereas uniform purchase subsidies address the smaller, strategic one.

The remainder of the dissertation is organised as follows. Chapter 2 derives closed-form solutions for the single-factor model under constant costs and discrete jumps, with an empirical application to delivery fleet electrification. Chapter 3 introduces deterministic cost decline as an analytical intermediate, then develops the full two-factor stochastic framework and the finite-difference solution methodology for both infinite and finite horizons. Chapter 4 calibrates the two-factor model to a specific vehicle pair using UK price data and US energy cost data, presents numerical results across model specifications, and establishes the hierarchy of adoption barriers. Chapter 5 synthesises the policy implications of the model hierarchy, situates the results within the broader energy transition literature, and identifies directions for future research.

Chapter 2

BENCHMARK MODELS: ANALYTICAL SOLUTIONS

This chapter derives closed-form solutions for the optimal adoption timing problem when the benefit differential follows arithmetic Brownian motion and costs are either constant or subject to discrete jumps. Section 2.2 establishes the value function, optimal threshold, and waiting time for constant costs. Section 2.3 extends the analysis to incorporate jump risk from technological breakthroughs or policy shocks. Section 2.4 calibrates the model using historical energy price data and analyses adoption decisions for individuals and commercial fleets across different geographic markets. Section 2.5 interprets the results and draws policy implications.

2.1 Model Specification

This section establishes the common elements that underpin all adoption models analysed in this dissertation. The model parameters, the unit cost advantage process (identical across all model variants), and the payoff structure used for infinite horizon problems are specified.

2.1.1 Model parameters

The model employs six core parameters that characterise the adoption decision. While the subsequent presentation uses electric vehicle adoption as the primary illustrative example, the framework applies to any technology transition where there is a trade-off between an initial investment and a potential gain in operating cost

This chapter is based on Szabó et al. (2025), joint work with Dávid Zoltán Szabó and Péter Csóka published in Technological Forecasting & Social Change. *The material is adapted to fit the dissertation framework and used with permission in accordance with Elsevier's author rights policy for doctoral dissertations. Subsection 2.3.1 extends the published analysis and is not part of the article.*

advantage driven by stochastic price dynamics. The model parameters and their interpretations are as follows.

Unit cost advantage (D). Cost advantage of using the new technology compared to the old technology for one unit of consumption. It is important to ensure that, when modelling the unit cost advantage, two technologies are made comparable in terms of monetary costs. For the purposes of the car example, the additional cost or savings in dollars from travelling 100 miles by car is suggested as the unit of measurement. For specific examples, it requires an estimation or concrete knowledge of the efficiency of the two cars (usually measured in "kWh of energy" for EVs and "gallons of gasoline per mile" for petrol cars), as well as the consumer price of the respective energy source. Beyond the cost difference arising from the energy sources, it can also cover usage-dependent maintenance costs. Furthermore, the cost-benefit analysis can incorporate both monetary and ethical costs and benefits arising from the disparity in CO_2 emissions between the two technologies. Monetary costs can also arise from a carbon tax, which represents future, hypothetical, or shadow costs of emissions.

Usage intensity (λ). Average (annual) intensity of using the technology. The unit should be consistent with the unit of D, i.e. if the unit of D is the cost per 100 miles, λ should be the average annual miles travelled by car, divided by 100. The product of the usage intensity and the unit cost advantage is referred to as the variable cost advantage per year.

Fixed cost of adopting the new technology (C). It can be either a lump sum paid when switching, which is considered an investment in the adoption of the new technology, or the discounted value of the monthly lease payment difference. The former represents the price difference of a new car in the event that an individual user or a company is considering purchasing new vehicles (or replacing old ones). The latter represents the discounted additional leasing costs for an electric car.[1] C can also include the potential tax benefits, like the $7,500 federal tax credit then offered for the purchase of certain EVs, although it was only available for some models in the US. The subjective costs can also potentially be offset by individual motivations to willingly pay extra for sustainability considerations and to enjoy the benefits of the new technology. For some, owning a new EV (or any new technology) might even be a direct or indirect factor of joy (e.g. by virtue signalling or genuinely

[1]The two can be considered equivalent since the fixed cost difference can be interpreted as the discounted payment difference between leasing the new technology minus leasing the old technology up to infinity. For types of vehicles lasting for decades and having reasonably high discount rates, a good approximation is the price difference between the new and the old technology. For types of vehicles lasting for years and having lower discount rates, on top of a one-time price difference, one has to discount the expected price difference after each round of amortisation up to infinity.

appreciating). To include subjective factors that might make the usage of EVs less comfortable, either D and/or C can be adjusted. Range anxiety, long charging times and inconvenient and insufficient charging infrastructure are significant barriers that are impeding the widespread adoption of EVs, as reported by Shen et al. (2019). Today, the charging time for EVs at home is around 10 hours,[2] allowing drivers to travel around 155 miles. People are willing to pay between $425 and $3,250 to decrease the charging time by one hour for 50 miles (Hidrue et al., 2011). In general, advancements in technology will decrease charging time in the long term, thereby influencing these costs.

Unit cost advantage improvement pace (α). The expected forward-looking improvement speed in technology efficiency, expressed in terms of an annualised (positive) cost difference increase, driven by technology improvements, the relative scarcity of resources and taxes, among other things. This represents the drift of the underlying stochastic process and will be referred to as drift. It can be estimated from the (inflation-adjusted) historical values of D, assuming that the trend will continue at the same pace.

Volatility of the unit cost advantage (σ). The expected volatility of the forward-looking unit cost advantage. It can be estimated from the (inflation-adjusted) historical values of D, assuming that future volatility will be at the same level as historical volatility.

Interest rate (r). The real (continuously compounded, annualised) discounting rate, assumed to be constant in the model. It can also incorporate the time value of money based on individual preferences through the application of a subjective discount rate. Haq and Weiss (2018) found that consumers tend to apply a subjective discount rate of $19 \pm 14\%$ when discounting initial investments against later savings for efficient energy and transport technologies. The analysis adheres to real discount rates.

2.1.2 The problem as a stochastic optimisation

The decision-maker is an individual, a company, or any organisation (hereafter company) that currently owns or leases a fleet of vehicles that run on conventional fuel. The company plans to gradually replace its fleet with EVs, but it is undecided about the timing of each transaction. The timing will depend on the specific usage profile and value trade-off of each vehicle. The electric car is henceforth referred to as *new technology* and the conventional car as *old technology*. The company

[2] https://www.mobilityhouse.com/int_en/knowledge-center/charging-time-summary

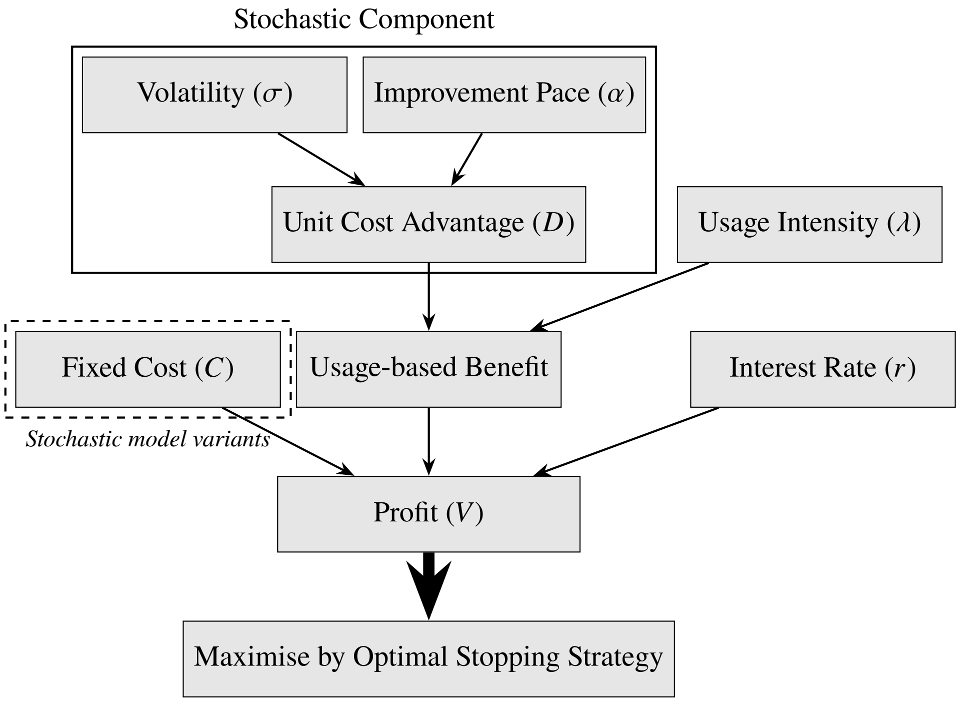

Figure 1: Factors influencing the decision to transit from the old technology to the new one. Different model specifications treat Fixed Cost differently: Chapter 2 introduces discrete jumps, while Chapter 3 models it as a mean reverting, stochastic process.

constantly tracks the cost-benefit of using the new technology compared to the old technology and wants to determine the optimal time to maximise their financial profit or subjective utility. This comprises various factors that contribute to savings, such as travel expenses related to different energy sources, maintenance costs and both monetary and ethical considerations.

A stochastic differential equation is assumed to describe the unit cost advantage of the new technology compared to the old one, denoted by $D(t)$ and specified as

$$dD(t) = \alpha\, dt + \sigma\, dW(t), \quad D(0) = D_0, \tag{2.1}$$

where α and σ are constants that represent the drift and volatility of the process, respectively. $W(t)$ represents a standard Brownian motion, defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P}_{D_0})$, where $\mathbb{P}_{D_0}$ denotes the probability measure under which $D(0) = D_0$. Let $\mathbb{E}_{D_0}$ denote expectations with respect to $\mathbb{P}_{D_0}$.

Arithmetic Brownian motion for fuel prices is well-supported empirically. The detailed justification with supporting citations and statistical evidence is presented in Subsection 2.2.1. The arithmetic Brownian motion process has several properties that are important for the optimal stopping analysis:

- **State space:** $D(t) \in \mathbb{R}$ (entire real line), allowing for both positive (cost advantage) and negative (cost disadvantage) values

- **No finite natural boundaries:** Requires computational domain truncation for numerical methods (Subsection 3.4.1)
- **Infinitesimal generator:** $\mathcal{L}_D = \frac{\sigma^2}{2}\frac{\partial^2}{\partial D^2} + \alpha\frac{\partial}{\partial D}$, which appears in the HJB variational inequality (Subsection 3.1.2)
- **Explicit solution:** $D(t) = D_0 + \alpha t + \sigma W(t)$, a Gaussian process with mean $D_0 + \alpha t$ and variance $\sigma^2 t$
- **First-passage times:** Passage times to thresholds admit known Laplace transforms, enabling analytical solutions for restricted optimal stopping problems

This unit cost advantage process $D(t)$ is common to *all* models in Chapters 2 and 3. What distinguishes the models is the specification of cost dynamics $C(t)$, the time horizon, and whether costs are subject to jump risk.

The canonical real options result under geometric Brownian motion establishes the benchmark investment threshold from which the arithmetic Brownian motion specialisation of this chapter departs.

Definition 2.1 (Investment Opportunity as a Call Option)**.** An investment opportunity is a claim on an underlying asset with stochastic value V_t whose cost of exercise is I. The value of this opportunity, $F(V)$, is characterised by the optimal stopping problem

$$F(V) = \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[(V_\tau - I)e^{-r\tau}\right], \tag{2.2}$$

where $\mathcal{T}$ denotes the set of all stopping times with respect to the filtration $\mathcal{F}_t$.

Theorem 2.2 (Fundamental Real Options Inequality)**.** *With irreversibility and the ability to delay, the optimal investment threshold V^* strictly exceeds the direct cost I. Let V follow a geometric Brownian motion with drift α and volatility σ. The optimal investment rule is to invest when $V \geq V^*$, where*

$$V^* = \frac{\beta_1}{\beta_1 - 1} I, \tag{2.3}$$

with $\beta_1 > 1$ the positive root of the characteristic quadratic $\frac{1}{2}\sigma^2\beta(\beta-1) + \alpha\beta - r = 0$. The term $\frac{\beta_1}{\beta_1-1}$ represents the option value multiplier.[3]

Proof. The general solution to the homogeneous ODE $\frac{1}{2}\sigma^2 V^2 F''(V) + \alpha V F'(V) - rF(V) = 0$ is $F(V) = AV^{\beta_1} + BV^{\beta_2}$, where $\beta_1 > 1$ and $\beta_2 < 0$ are the roots of the characteristic equation. The boundary condition $F(0) = 0$ eliminates the second term ($B = 0$). Applying value matching $F(V^*) = V^* - I$ and smooth pasting $F'(V^*) = 1$ to $F(V) = AV^{\beta_1}$ yields the threshold V^* defined in (2.3). Since $\beta_1 > 1$, it follows that $V^* > I$ (Dixit and Pindyck, 1994, pp. 142–152). □

[3] This canonical result uses geometric Brownian motion. The models in Chapters 2 and 3 employ arithmetic Brownian motion for the benefit differential, yielding analogous threshold structures with different functional forms.

Remark 2.3. The GBM derivation above proceeds via ODE methods: the value function solves a second-order homogeneous ODE, and the optimal threshold follows from value matching and smooth pasting conditions. The ABM optimal stopping problems in this chapter admit an analogous threshold structure but require different analytical tools. Because arithmetic Brownian motion lacks the scale invariance of GBM, the value function does not take a power-law form, and the characteristic ODE has non-constant coefficients. The threshold characterisation instead exploits known Laplace transforms for ABM first-passage times (Borodin and Salminen, 2015), with optimality established via the verification theorem of Dayanik and Karatzas (2003). Both approaches recover the fundamental real options principle: the optimal threshold strictly exceeds the intrinsic value of immediate adoption by an option value premium that increases in volatility.

2.1.3 Value of Adoption and Payoff Structures

The value of switching from old to new technology at any point in time depends on the expected discounted profit stream from adoption minus the upfront cost. The expected discounted profit of the company is calculated by comparing two strategies:

Scenario 1 (Threshold Strategy): The company transitions to the new technology when the unit cost advantage process $D(t)$ first crosses a threshold level x^*.

Scenario 2 (Never Switch): The company maintains the old technology throughout its entire lifetime.

By comparing these scenarios, the company incurs the same cost over the time period before adoption. At the moment of switching, there is an additional cost C (negative profit) from the upfront investment. However, over the entire remaining time after adoption, the company realises profit linked with the evolution of the unit cost advantage $D(t)$.

Infinite Horizon (Perpetual Option)

One would naturally expect α to be positive due to continuous technological improvements, making it intuitive to wait until $D(t)$ reaches a sufficiently high threshold before switching. The first passage time of $D(t)$ to threshold x^* is:

$$D_{x^*} := \inf_{t \geq 0} \{D(t) \geq x^*\}. \tag{2.4}$$

The threshold is written as a passage rather than a hitting condition so that the same definition serves both the diffusion of this section and the jump-diffusion of Section 2.3, where the state can cross a level by overshooting it. For a continuous D started below x^* the two coincide, and for $x^* \leq D_0$ the definition returns $D_{x^*} = 0$, which is the immediate-adoption convention used throughout.

The related profit stream from Scenario 1 reads as

$$V_{D_{x^*}} := \mathbb{E}_{D_0}\left[-e^{-rD_{x^*}}C + \int_{D_{x^*}}^{\infty} e^{-rt}D(t)\lambda\, dt\right]. \tag{2.5}$$

The terms in Equation (2.5) are interpreted as follows. At time D_{x^*}, there is an additional cost C (negative profit) from the adoption investment. However, over the entire remaining time interval (D_{x^*}, ∞), the company realises an everlasting profit linked with the trajectory of D. This is calculated with the help of an integral on (D_{x^*}, ∞), equipped with the constant positive intensity parameter λ, which represents the usage intensity of the technology. Therefore, at a given future time point t, the integrand is $e^{-rt}D(t)\lambda$.

After some algebraic manipulation and taking expectations (detailed derivations are provided in Subsection B.8.3), this can be rewritten as:

$$\Pi_\infty(D, C) = \frac{\lambda D}{r} + \frac{\lambda\alpha}{r^2} - C. \tag{2.6}$$

Interpretation:

- $\frac{\lambda D}{r}$: Capitalised (perpetual annuity) value of the current benefit flow λD
- $\frac{\lambda\alpha}{r^2}$: Additional value from the expected drift α in the unit cost advantage process
- C: One-time adoption cost paid at the switching time

2.2 Analytical Results for the Constant-Cost Model

This section derives the analytical characterisation of the optimal stopping problem when the adoption cost is constant. The optimal stopping formulation is established, the value function is derived in closed form, optimal threshold is determined, and the expected waiting time is characterised.

2.2.1 Empirical Motivation for the Stochastic Model

The model addresses a decision faced by individuals or companies considering replacing current vehicles or expanding their fleet: whether to choose an electric vehicle (EV) instead of a conventional combustion engine vehicle. Both options present distinct trade-offs in financial terms and personal preferences. EVs typically involve higher upfront costs, which may be offset by lower operating costs due to cheaper electricity compared to fossil fuels. Additional cost considerations include insurance, maintenance, and repairs. Beyond financial factors, motivations for EV adoption include contributing to the green transition and reducing local air pollution. However, adoption faces significant barriers: limited charging infrastructure compared to gas stations, substantially longer refueling times, and range anxiety (Shen et al., 2019).

Accurate evaluation of these trade-offs requires comprehensive analysis that accounts for the stochastic nature of energy markets. The retail prices of gasoline and electricity have a significant impact on the decision of whether it is worth switching to EVs or waiting for potentially higher expected discounted profits. In the end, we have to pay for each mile driven, whether it is for charging our EVs or refilling our tanks at a petrol station. If the prices are volatile, our natural intuition is to wait until the net benefit becomes stable and sufficiently high to offset the initial investment. To emphasise the significant volatility in energy markets, we analyse the prices of gasoline and energy in the US and European regions.

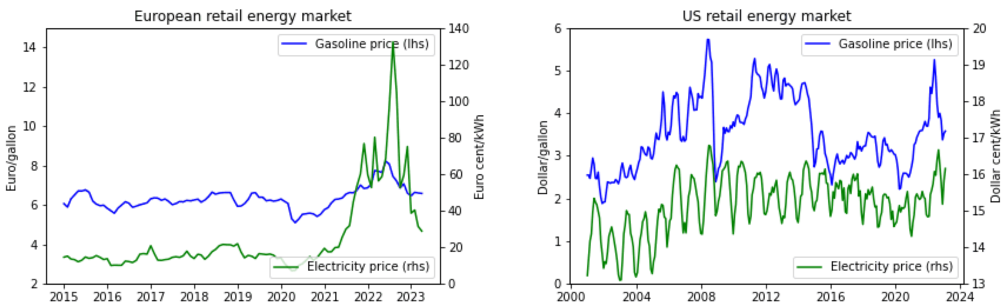

Figure 2: Historical Retail Prices in the Energy Markets in the US and Europe. Historical average monthly retail prices are adjusted for inflation.[4]

Figure 2 shows the average retail price of one gallon of gasoline and one kWh of electricity, as reported monthly by the U.S. Energy Information Administration, European Commission and Ember. Historical prices are adjusted for inflation and reported in dollars for the US and euros for Europe. Both markets show characteristics of volatility and significant price movements. In the US, the energy price features some seasonality. Nonetheless, the volatility of the gasoline price vastly outstrips that of the energy price. Therefore, we can ignore the seasonal effect when comparing the price differences. In Europe, price levels are generally higher for energy products, but electricity has always been relatively cheaper. The electricity price was already volatile before 2021. Additionally, there was a strong spurt in prices during the European Energy Crisis period, and there has been no return to stable low prices since then. In the examined period, the standard deviation of the gasoline price was 23.96% for the US and 8.7% for Europe. The standard deviation of the electricity price was 5.1% for the US and 90.71% for Europe.

Although it is beyond the scope of this analysis to forecast the volatility of energy markets, there is no reason to assume that the market structure would change and markets would stabilise drastically without targeted policy intervention. Moreover, the potential green transition, which would require heavy investments in green technology and electricity networks and would entail a structural change in the

[4]Further details on historical data sources are available in Appendix A.

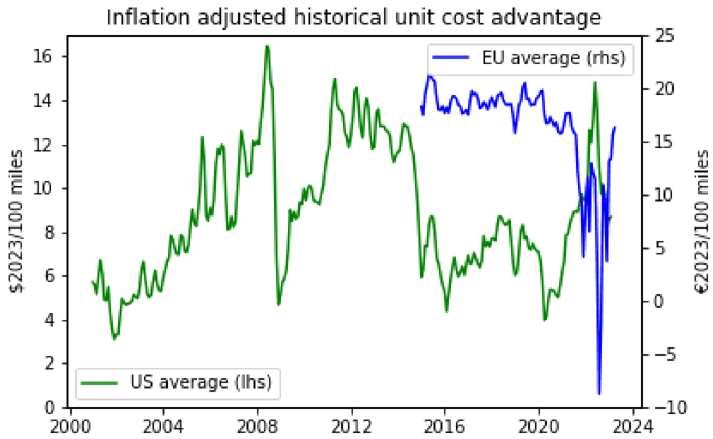

Figure 3: Historical unit cost advantage of travelling 100 miles by an electric car compared to a conventional car.[5]

energy markets, might prove to be disruptive and elevate energy price volatility during the transition period. This risk is especially pronounced in divergent scenarios and poses a threat to consumers who rely on electricity as a substitute for fossil-fuel-intensive options. From this, we conclude that there is value in investigating how a stochastic modelling approach could enhance the optimisation of the timing of the technology switch.

Arithmetic Brownian motion for fuel price has been assumed in Kleindorfer et al. (2012), and this modelling choice is supported by Meade (2010), who argue that for long-term oil prices, neither geometric Brownian motion nor mean-reverting stochastic processes yield superior forecasts, as surveyed in Subsection 1.2.3. Moreover, cointegration and non-stationary features have been generally observed in commodity markets (Beck, 1994). This supports a modelling framework in which the two prices representing gasoline and electricity uncertainty contain trends. Finally, one can also notice that the difference between two arithmetic Brownian motions is an arithmetic Brownian motion itself.

Our suggested modelling choice also aligns with the historical unit cost advantage time series, calculated from electricity and fuel prices from Europe and the US, as presented in Figure 3. Although external shocks, such as the European Energy Crisis in 2022, can cause market dislocations, the trend is observable over several decades.

Statistical tests on the most recent five years of data support the arithmetic Brownian motion specification over mean-reverting or stationary alternatives. Table 2

[5]Note that the unit cost advantage is not a simple difference between electricity and gasoline prices. To calculate the historical unit cost advantage, the same assumptions on fuel efficiency are applied as described in Section 2.4.2, where other attributes of this historical time series are discussed in detail.

presents results from the Augmented Dickey-Fuller (ADF) test and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test. Both tests include a constant but no trend term. For both regions, the ADF test fails to reject the presence of a unit root at standard confidence levels, while the KPSS test rejects stationarity at 95% confidence, consistent with non-stationary dynamics.

Test	Region	Test Statistic	p-value	Null Hypothesis
ADF	EUR	-2.17	0.219	Unit root present
KPSS	EUR	0.867	0.01	Stationary
ADF	US	-1.93	0.319	Unit root present
KPSS	US	0.497	0.042	Stationary

Table 2: Statistical tests for unit root and stationarity applied to five-year historical unit cost advantage data in Europe and the United States.

These empirical patterns motivate the arithmetic Brownian motion specification used throughout this chapter. The following subsection formalises the optimal stopping problem under this stochastic benefit process.

2.2.2 Optimal Stopping and Closed-Form Value Function

The unit cost advantage $D(t)$ follows the arithmetic Brownian motion process specified in Subsection 2.1.2, equation (2.1). Recall from Subsection 2.1.3 that the company seeks to choose a threshold x^* to maximise the expected discounted profit $V_{D_{x^*}}$ defined in (2.5), where D_{x^*} denotes the first passage time (2.4). The firm's objective is to solve:

$$V^* := \sup_{x^*} V_{D_{x^*}} = \sup_{x^*} \mathbb{E}_{D_0}\Big(-e^{-rD_{x^*}}C + \int_{D_{x^*}}^{\infty} e^{-rt} D(t)\lambda\, dt\Big). \tag{2.7}$$

To analyse the company's profit stream, an analytical formula for $V_{D_{x^*}}$ is derived below.

Theorem 2.4. *Assuming that $r > 0$, for a given level x^*, $V_{D_{x^*}}$ is given as*

$$V_{D_{x^*}} = \Big(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C\Big) e^{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma}(x^*-D_0) - |x^*-D_0|\sqrt{\frac{\alpha^2}{\sigma^2}+2r}\right)}. \tag{2.8}$$

Proof. Equation (2.7) is rearranged as

$$V_{D_{x^*}} = \mathbb{E}_{D_0}\Big(-e^{-rD_{x^*}}C + \int_{D_{x^*}}^{\infty} e^{-rt}(D(t))\lambda dt\Big) = \tag{2.9}$$

$$\mathbb{E}_{D_0}\Big(\mathbb{E}_{D_0}\Big(-e^{-rD_{x^*}}C + \int_{D_{x^*}}^{\infty} e^{-rt}(D(t))\lambda dt | D_{x^*} = s\Big)\Big) =$$

$$\int_0^{\infty} f_{D_{x^*}}(s)\mathbb{E}_{D_0}\Big(-e^{-rs}C + \int_s^{\infty} e^{-rt}(x^* + \alpha(t-s) + \sigma\hat{W}(t-s))\lambda dt\Big) ds,$$

where $f_{D_{x^*}}$ is the density function of D_{x^*}, and, by the strong Markov property, $\{\hat{W}(u)\}_{u\geq 0}$ is a standard Wiener process started at zero and independent of $\mathcal{F}_{D_{x^*}}$. Note that $D(D_{x^*}) = x^*$, as $D(t)$ is a continuous process.

Equation (2.9) is expanded into four terms as

$$V_{D_{x^*}} = -\int_0^\infty f_{D_{x^*}}(s)e^{-rs}Cds + \int_0^\infty f_{D_{x^*}}(s)\int_s^\infty e^{-rt}(x^* - \alpha s)\lambda dtds + \quad (2.10)$$
$$\int_0^\infty f_{D_{x^*}}(s)\int_s^\infty e^{-rt}\alpha\lambda t dtds + \int_0^\infty f_{D_{x^*}}(s)\mathbb{E}_{D_0}\int_s^\infty e^{-rt}(\sigma\hat{W}(t-s))\lambda dtds.$$

The last three terms in Equation (2.10) are calculated separately.

The second term reads as

$$\int_0^\infty f_{D_{x^*}}(s)\int_s^\infty e^{-rt}(x^* - \alpha s)\lambda dtds = \quad (2.11)$$
$$\int_0^\infty f_{D_{x^*}}(s)(x^* - \alpha s)\lambda\frac{e^{-rs}}{r}ds.$$

The third term reads as

$$\int_0^\infty f_{D_{x^*}}(s)\int_s^\infty e^{-rt}\alpha\lambda t dtds = \quad (2.12)$$
$$\int_0^\infty f_{D_{x^*}}(s)\lambda\alpha\Big(\frac{e^{-rs}s}{r} + \frac{e^{-rs}}{r^2}\Big)ds.$$

By applying Fubini's theorem and noticing that $\mathbb{E}\,\hat{W}(u) = 0$ for every $u \geq 0$, the fourth term reads as

$$\int_0^\infty f_{D_{x^*}}(s)\mathbb{E}_{D_0}\int_s^\infty e^{-rt}(\sigma\hat{W}(t-s))\lambda dtds = \quad (2.13)$$
$$\int_0^\infty f_{D_{x^*}}(s)\int_s^\infty e^{-rt}\sigma\,\mathbb{E}_{D_0}\big(\hat{W}(t-s)\big)\lambda dtds = 0.$$

Having simplified each term, the sum is expressed as

$$V_{D_{x^*}} = \int_0^\infty f_{D_{x^*}}(s)e^{-rs}\Big(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C\Big)ds = \quad (2.14)$$
$$\Big(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C\Big)\mathbb{E}_{D_0}(e^{-rD_{x^*}}).$$

If T_a is the first hitting time of a standard Wiener process with drift ($X(t) = ct + B(t)$, $B(0) = D_0$) at level a, then according to (Borodin and Salminen (2015) p.295):

$$\mathbb{E}_{D_0}(e^{-xT_a}) = e^{c(a-D_0)-|a-D_0|\sqrt{c^2+2x}}. \quad (2.15)$$

Tailoring Equation (2.15) to the present problem yields

$$V_{D_{x^*}} = \Big(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C\Big)e^{\frac{1}{\sigma}(\frac{\alpha}{\sigma}(x^*-D_0)-|x^*-D_0|\sqrt{\frac{\alpha^2}{\sigma^2}+2r})}. \quad (2.16)$$

□

The value function in (2.16) has a natural economic reading. The bracketed term is the payoff collected at the moment of switching, net of the fixed cost C paid upfront: $\frac{\lambda x^*}{r}$ is the discounted profit of the annuity realised upon adoption if the unit cost advantage stayed fixed at its switching-time level, while $\frac{\lambda\alpha}{r^2}$ adds the value of its further expected growth at rate α. The exponential factor multiplying the bracket is the expected discount factor of the switching time, so the value of the threshold strategy is simply the adoption payoff at the threshold, discounted back over the random waiting time.

Corollary 2.5. *The profit the company obtains if they decide to switch to the new technology immediately ($x^* = D_0$ and $D_{x^*} = 0$) is*

$$\frac{\lambda D_0}{r} + \frac{\lambda\alpha}{r^2} - C. \tag{2.17}$$

2.2.3 Optimal Threshold Characterisation

The level x^* that maximises $V_{D_{x^*}}$ is determined next.

Theorem 2.6. *For a given set of parameters, $V_{D_{x^*}}$ reaches its supremum at a specific point*

$$x_{max} = \begin{cases} x^{**}, & x^{**} > D_0. \\ D_0, & otherwise, \end{cases} \tag{2.18}$$

where

$$x^{**} := \frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\alpha}{r} + \frac{Cr}{\lambda}. \tag{2.19}$$

The corresponding maximum value can be characterised as

$$V_{max} := \sup_{x^* \in \mathbb{R}} V_{D_{x^*}} = \begin{cases} \left(\frac{\lambda x^{**}}{r} + \frac{\lambda\alpha}{r^2} - C\right) e^{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma} - \sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right)(x^{**} - D_0)}, & x^{**} > D_0. \\ \frac{\lambda D_0}{r} + \frac{\lambda\alpha}{r^2} - C, & \text{otherwise.} \end{cases} \tag{2.20}$$

Proof. First, it follows from Equation (2.16) that $V_{max} > 0$ for any given set of parameters. This follows from choosing a sufficiently large value for x^* to make $\left(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C\right)$ positive while keeping all other parameters fixed.

Second, $x_{max} \geq D_0$ for any given set of parameters. Suppose $x_{max} < D_0$ for a fixed set of parameters. If $V_{max} \leq 0$, this contradicts the first claim. If $V_{max} > 0$, then $V_{max} < V_{D_0}$ as $\frac{\lambda D_0}{r} + \frac{\lambda\alpha}{r^2} - C > \left(\frac{\lambda x_{max}}{r} + \frac{\lambda\alpha}{r^2} - C\right) e^{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma} + \sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right)(x_{max} - D_0)}$, yielding a contradiction.

It therefore suffices to consider the case $x^* \geq D_0$ for the maximisation problem. In this case, $V_{D_{x^*}}$ reads as

$$\begin{aligned}
\hat{V}_{D_{x^*}} =& \left(Ax^* + B\right)e^{G(x^*+H)}, \text{ where} \\
A =& \frac{\lambda}{r}, \\
B =& \frac{\lambda\alpha}{r^2} - C, \\
G =& \frac{\frac{\alpha}{\sigma} - \sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}{\sigma}, \\
H =& - D_0.
\end{aligned} \tag{2.21}$$

The first derivative of $\hat{V}_{D_{x^*}}$ with respect to x^* is set to zero to find x^{**} that maximises $\hat{V}_{D_{x^*}}$ as a function of x^*, with all other parameters fixed, provided the second-order derivative is negative. It follows that

$$\frac{d\hat{V}_{D_{x^*}}}{dx^*} = \left(AGx^* + A + BG\right)e^{G(x^*+H)}. \tag{2.22}$$

Because of the positive feature of the exponential function, $\frac{d\hat{V}_{D_{x^*}}}{dx^*}$ is zero exactly when the coefficient $(AGx^* + A + BG)$ is zero.

$$\begin{aligned}
\left(AGx^* + A + BG\right) &= 0, \text{ from which} \\
x^{**} &= \frac{-1}{G} - \frac{B}{A} \\
x^{**} &= \frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\alpha}{r} + \frac{Cr}{\lambda}.
\end{aligned} \tag{2.23}$$

The second-order derivative should be negative at x^{**}, which reads as

$$\begin{aligned}
\frac{d^2\hat{V}_{D_{x^*}}}{d^2x^*}(x^{**}) &< 0, \\
G(A(Gx^{**} + 2) + BG)e^{G(x^{**}+H)} &< 0, \text{ from which} \\
G((-A - BG + 2A) + BG) = GA &< 0.
\end{aligned} \tag{2.24}$$

Note that for $r > 0$, $G < 0$ and $A > 0$ hold for any parameter combination. Thus, $GA < 0$, and as a result, both $(AGx^* + A + BG)$ and $(e^{G(x^*+H)})$ are decreasing functions of x^*. Therefore, $(AGx^* + A + BG)$ is positive for $x^* < x^{**}$ and negative for $x^* > x^{**}$. Along with the positive property of $e^{G(x^{**}+H)}$, it follows that $\frac{d\hat{V}_{D_{x^*}}}{dx^*}$ changes from positive to negative at x^{**}, meaning that x^{**} is a local maximum of $\hat{V}_{D_{x^*}}$. As $\hat{V}_{D_{x^*}}$ has a unique stationary point, x^{**} is also a global maximum.

Note that x^{**} does not depend on D_0. Thus, for a given set of parameters, if $x^{**} \geq D_0$, then $x_{max} = x^{**}$.

On the other hand, $\hat{V}_{D_{x^*}}$ is decreasing as a function of x^* for $x^* > x^{**}$. Therefore, for the given set of parameters, when $D_0 > x^{**}$, $x_{max} \leq D_0$. Since $x_{max} \geq D_0$ for any set of parameters, it follows that $x_{max} = D_0$. □

The optimal threshold (2.19) reveals several properties. The threshold x^{**} is independent of the current state D_0, reflecting translation invariance of arithmetic Brownian motion. It increases with volatility σ (higher uncertainty delays adoption) and decreases with drift α (stronger expected improvements encourage earlier switching). The threshold depends on fixed cost C and usage intensity λ only through their ratio, so proportional scaling of both parameters leaves x^{**} unchanged. The effect of the discount rate r is non-monotonic, as it discounts both future benefits and the switching cost.

2.2.4 Expected Waiting Time

Beyond determining the optimal threshold, a complete characterisation requires analysing the expected time until the threshold is reached and adoption occurs.

Theorem 2.7. *For $D_0 \geq x^{**}$, it holds that*

$$\mathbb{E}_{D_0}(D_{x_{max}}) = 0. \tag{2.25}$$

For $D_0 < x^{**}$ and $\alpha \leq 0$, it holds that

$$\mathbb{E}_{D_0}(D_{x_{max}}) = \infty. \tag{2.26}$$

For $D_0 < x^{**}$ and $\alpha > 0$, it holds that

$$\mathbb{E}_{D_0}(D_{x_{max}}) = \frac{1}{\alpha}\Big(\frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\alpha}{r} + \frac{Cr}{\lambda} - D_0\Big). \tag{2.27}$$

Proof. See Appendix B.1. □

If the current unit cost benefit D_0 already exceeds the optimal threshold, adoption should occur immediately. With non-positive drift and $D_0 < x^{**}$, the benefit process is not expected to reach the threshold in finite time, so the firm waits indefinitely. With positive drift, expected waiting time is inversely proportional to α and linearly increasing in the gap $(x^{**} - D_0)$. Notably, expected waiting time does not depend on volatility σ, though σ affects the optimal threshold level itself.

2.2.5 General optimal stopping problems in one dimension

The previous results are extended to general optimal stopping problems to prove that V_{max} from Theorem 2.6 is also a solution for the problem, which allows arbitrary stopping times. This means that V_{max} (and x_{max} from 2.6) is not only optimal for a specific strategy where the trigger is defined by the first passage time, but it is also the optimal value regardless of the stopping strategy chosen. In other words, any stopping strategy, whether triggered by the first passage time or any other criterion, will not yield a higher expected profit than V_{max}.

As discussed earlier, the analysis considers a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P}_{D_0})$. $\mathbb{E}_{D_0}$ denotes expectations with respect to $\mathbb{P}_{D_0}$, and $\mathcal{T}$ represents the set of stopping times with respect to the filtration $\{\mathcal{F}_t\}_{t\geq 0}$. First, the objective of the company formulated mathematically as a generic optimal stopping problem reads as

$$V(D_0) = \sup_{\tau\in\mathcal{T}} \mathbb{E}_{D_0}\Big(-e^{-r\tau}C + \int_\tau^\infty e^{-rt}D(t)\lambda dt\Big). \tag{2.28}$$

This can be rearranged as

$$\sup_{\tau\in\mathcal{T}} \mathbb{E}_{D_0}\Big(-e^{-r\tau}C + \int_\tau^\infty e^{-rt}D(t)\lambda dt\Big) = \tag{2.29}$$

$$\sup_{\tau\in\mathcal{T}} \mathbb{E}_{D_0}\Big(-e^{-r\tau}C + e^{-r\tau}\mathbb{E}\Big(\int_\tau^\infty e^{-r(t-\tau)}D(t)\lambda dt|\mathcal{F}_\tau\Big)\Big). \tag{2.30}$$

The following notation is introduced:

$$g(D(\tau)) = \mathbb{E}\Big(\int_\tau^\infty e^{-r(t-\tau)}D(t)\lambda dt|\mathcal{F}_\tau\Big), \quad \hat{g}(D(\tau)) = g(D(\tau)) - C,$$

to finally write

$$V(D_0) = \sup_{\tau\in\mathcal{T}} \mathbb{E}_{D_0}\{e^{-r\tau}\hat{g}(D(\tau))\}. \tag{2.31}$$

Equation (2.31) is related to the general *optimal stopping problem*, which is to find $\Phi(y)$ and $\mathcal{T}^* \in \mathcal{T}$ such that

$$\Phi(y) = \sup_{\tau\in\mathcal{T}} \mathbb{E}_y\{e^{-r\tau}g(D(\tau))\} \stackrel{\Delta}{=} \sup_{\tau\in\mathcal{T}} J^\tau(y) = J^{\mathcal{T}^*}(y), \quad y \in \mathbb{R}. \tag{2.32}$$

The function $\Phi(y)$ is called the *value function*, and the stopping time $\mathcal{T}^*$ (if it exists) is called an *optimal stopping time*.

The results of Dayanik and Karatzas (2003) are utilised next. To this end, a few important concepts are introduced. The state space of the diffusion process D is $\mathbb{R}$, and D is regular in $\mathbb{R}$. Thus, $-\infty$ and ∞ are natural boundaries for process D. The *reward function* $\hat{g}()$ is given as $\hat{g}(x) = \frac{x\lambda}{r} + \frac{\lambda\alpha}{r^2} - C$, which is a bounded function on every compact subset of $\mathbb{R}$. With $\mathcal{L}_D$ the infinitesimal generator of D introduced in Section 2.1, the equation $\mathcal{L}_D u = ru$ has two linearly independent, positive solutions. The increasing solution is denoted $\psi()$ and the decreasing solution $\phi()$.

Further, the following notations are introduced:

$$\ell_a \stackrel{\Delta}{=} \limsup_{x\to-\infty} \frac{\hat{g}^+(x)}{\phi(x)} \quad \text{and} \quad \ell_b \stackrel{\Delta}{=} \limsup_{x\to\infty} \frac{\hat{g}^+(x)}{\psi(x)}, \tag{2.33}$$

where $\hat{g}^+() \stackrel{\Delta}{=} \max\{0, \hat{g}()\}$.

Several new results are presented next.

Theorem 2.8. *$V(x) < \infty$ for all $x \in \mathbb{R}$, corresponding to the problem of Equation* (2.31)*. Moreover, ℓ_a and ℓ_b are both finite.*

Proof. According to Proposition 5.10 of Dayanik and Karatzas (2003), either $V \equiv \infty$ in $\mathbb{R}$ or $V(x) < \infty$ for all $x \in \mathbb{R}$.

Note that the reward function $\hat{g}$ of the optimal stopping problem is bounded on every compact interval and is only unbounded at its ∞ limit, meaning that $\lim_{x\to\infty} \hat{g}(x) = \infty$ and $\hat{g} < \infty$ otherwise. Therefore, the only way that $V \equiv \infty$ in $\mathbb{R}$ would hold is if the value function corresponding to the problem associated with reaching $D(\infty)$ exists. In other words, this would be the case if the optimal stopping time could be characterised as $\mathcal{T}^* = \lim_{x\to\infty} D_x$ for all D_0 initial points, and besides, $\mathbb{E}_{D_0}\{e^{-r\mathcal{T}^*}\hat{g}(D(\mathcal{T}^*))\} = \infty$ would hold for each D_0. Nonetheless, based on Equation (2.8), $\lim_{x^*\to\infty} V_{D_{x^*}}(D_0) = 0$ for all D_0. Hence, $V(x) < \infty$ for all $x \in \mathbb{R}$.

Proposition 5.10 of Dayanik and Karatzas (2003) also implies that ℓ_a and ℓ_b are both finite. □

The following notations are introduced: $\Gamma \triangleq \{x \in \mathbb{R} : V(x) = \hat{g}(x)\}$, $C \triangleq \mathbb{R} \setminus \Gamma$ and $\tau^* = \inf\{t \geq 0 : D(t) \in \Gamma\}$.

Theorem 2.9. *τ^* is an optimal stopping time.*[6]

Proof. If both $\ell_a = 0$ and $\ell_b = 0$, then Proposition 5.13 of Dayanik and Karatzas (2003) immediately yields the theorem. It is straightforward that only $\ell_b > 0$ is possible, so assume $\ell_b > 0$. By Proposition 5.14 of Dayanik and Karatzas (2003) it then suffices to show that there is no $\ell \in \mathbb{R}$ such that $(\ell, \infty) \subset C$, that is, that Γ is unbounded above. Appendix B.2 establishes this by constructing an r-superharmonic majorant of $\hat{g}$ that is tangent to it at x^{**}, which gives $V = \hat{g}$ on $[x^{**}, \infty)$. Hence τ^* is an optimal stopping time. □

Theorem 2.10. *The stopping region for the optimal stopping problem is $\Gamma = [x^{**}, \infty)$, where x^{**} is determined by Equation* (2.23).

Proof. The argument of Appendix B.2 establishes $V(D_0) = \hat{g}(D_0)$ for all $D_0 \geq x^{**}$, so $[x^{**}, \infty) \subset \Gamma$; the interval is closed at x^{**} because V and $\hat{g}$ are continuous and value matching holds at the threshold itself. To complete the proof by contradiction, assume there is a point $D_0 < x^{**}$ in Γ. However, Theorem 2.6 showed that $\hat{g}(D_0) < V_{D_{x^{**}}}$ for all $D_0 < x^{**}$, leading to a contradiction. □

Economically, the reduction of the integral objective to the state-dependent reward $\hat{g}(D(\tau))$ separates the two ingredients of the firm's problem: $\hat{g}(x)$ is the intrinsic value, the net present value realised upon adoption when the state equals x,

[6] The argument differs from the published version (Szabó et al., 2025), which reaches the same conclusion by contradiction using an inequality between $\hat{g}(\gamma^*)$ and V that is not derived there. The verification argument used here is self-contained, and it additionally yields that the stopping region is closed at x^{**}.

while the supremum over stopping times turns $V(D_0)$ into the value of the option to invest. The results above establish that the first-passage strategy τ^* is globally optimal and not merely a plausible heuristic: no alternative strategy, however complex, yields higher expected returns than the threshold rule defined by x^{**}, so the adoption decision is fully characterised by a single implementable threshold.

2.3 Analytical Results under Jump Risk

Recall that in our benchmark model, the unit cost advantage of the new technology compared to the old one is a Brownian motion with drift, specified as

$$dD(t) = \alpha\, dt + \sigma\, dW(t), \quad D(0) = D_0, \tag{2.34}$$

where α is the drift, σ is the volatility, and $W(t)$ is a standard Brownian motion.

In this section, the fixed cost of adopting the new technology $C(t)$ is assumed to accommodate abrupt jumps. These jumps can be interpreted as an instantaneous reduction in the upfront cost of EVs, reflecting sudden price drops within the EV industry. Such reductions can be attributed to various factors, including advancements in battery technology, the entry of new manufacturers into the market, the introduction of incentives for EV purchases (such as tax credits, rebates, and grants), and heightened regulatory pressures. Given the necessity of such incentives to achieve the Sustainable Development Goals (SDGs), these positive abrupt changes are anticipated to occur with a certain frequency in the foreseeable future. However, these jumps can also represent policy changes in the opposite direction, where the removal of subsidies or tax credits leads to sudden increases in the effective cost of an EV.

Mathematically, this is represented as

$$C(t) = C_0 + Y(t), \tag{2.35}$$

where $Y(t)$ denotes a compound Poisson process that captures the cumulative effect of these technological shifts. Poisson processes have been widely utilised in the literature to model the impact of radical technological innovations (Daming et al., 2014). For instance, automotive manufacturers have historically introduced new vehicle models at regular intervals, and it is expected that, due to the factors mentioned above, these new models will either enhance overall quality or reduce manufacturing costs. These sporadic, abrupt innovations are likely to contribute to a decline in the fixed cost $C(t)$, and such dynamics are effectively modelled using Poisson processes. The aggregation of these discrete changes results in a compound Poisson process allowing for positive and negative jumps as well, which provides a robust framework for analysing the cost evolution in the EV industry. Throughout, the jump sizes are i.i.d. and independent of the arrival process, and Y is independent of the Brownian motion W driving D. In economic terms, the events behind an abrupt price cut, such as a battery breakthrough, a new entrant, or a subsidy introduced or withdrawn, are assumed not to be systematically timed with the fuel-price shocks that move the unit cost advantage, and the magnitude of one such event carries no information about when the next one arrives. The following additional jump parameters are introduced:

Intensity of jump (ν). Represents the expected number of jump events occurring within a fixed time interval. It quantifies how frequently abrupt changes, such as technological advancements or policy shifts, impact the cost structure. Since it is annualised, ν denotes the expected number of jumps per year.

Average jump size in the fixed cost (d_e). Captures the magnitude of cost changes triggered by jumps. A positive d_e indicates that EVs become more expensive, while a negative d_e reflects a reduction in the fixed cost $C(t)$ due to factors such as policy incentives or technological breakthroughs. This parameter characterises the expected directional shift in EV prices over time, driven by discrete events rather than gradual trends.

Two solution concepts are developed for this extension, and it is instructive to present them together. Subsection 2.3.1 obtains the *optimal (closed-loop) adoption rule*: because the adoption payoff aggregates the two state variables linearly, the two-factor problem reduces *exactly* to one-dimensional optimal stopping of a Lévy process with linear reward, for which Mordecki (2002a) supplies the threshold in closed form. Subsection 2.3.2 then analyses a *precommitted first-passage benchmark* that fixes a trigger on D at time zero. Subsection 2.3.3 closes the section with the economic implications of the two rules, and quantifies the cost of ignoring realised cost jumps.

2.3.1 Closed-loop solution of the jump risk model

Under jump risk the state of the adoption problem is the pair (C, D), and the optimal rule is a boundary in this plane: adoption is triggered by any combination of a sufficiently high unit cost advantage and a sufficiently low fixed cost. This subsection derives the boundary in closed form. In brief, the solution is that adoption is optimal the first time the immediate-adoption payoff, a single Lévy process Z aggregating the two state variables, reaches a level z^* determined explicitly by the two positive roots of a cubic; Theorem 2.11 states the rule and the boundary it induces in the (C, D) plane.

Assume throughout the closed-loop analysis that $r > 0$ and that the fixed cost follows $C(t) = C_0 + Y(t)$ with exponential jump sizes: the i.i.d. cost reductions $\xi \sim \mathrm{Exp}(\gamma)$, $\gamma > 0$, arrive at rate $\nu > 0$, so that $d_e = -1/\gamma$, and Y remains independent of W. The specialisation is one-sided, only cost reductions being modelled here; cost-increase risk is treated separately in Proposition B.9, which needs no distributional assumption on the jump size.

The company chooses an adoption time τ, a stopping time of the filtration generated by (C, D), so as to maximise the profit stream

$$V_\tau = \mathbb{E}_{D_0,C_0}\left[-e^{-r\tau}\, C(\tau) + \int_\tau^\infty e^{-rt} \lambda\, D(t)\, dt\right], \tag{2.36}$$

where $\mathbb{E}_{D_0,C_0}$ denotes expectation under the measure with $D(0) = D_0$ and $C(0) = C_0$.

Let

$$Z_t := \frac{\lambda D_t}{r} + \frac{\lambda\alpha}{r^2} - C_t \tag{2.37}$$

denote the value of adopting at time t, that is the immediate-adoption payoff of Corollary 2.5 evaluated at the prevailing state rather than at $t = 0$. Because Y is independent of W, the superposition of the Brownian and the jump component again has stationary and independent increments, so Z is a Lévy process, with drift $\mu_Z = \lambda\alpha/r$, volatility $\sigma_Z = \lambda\sigma/r$, and upward jumps of size $\xi \sim \mathrm{Exp}(\gamma)$ arriving at rate ν, the sign reversal reflecting that a cost reduction raises the adoption payoff. Stationary increments make Z *spatially homogeneous*.[7]

The reduction to this single state variable is a chain of equalities. Fix a stopping time τ of the filtration generated by (C, D). Then

$$\begin{aligned}
V_\tau &= \mathbb{E}\Big[-e^{-r\tau} C_\tau + \int_\tau^\infty e^{-rs} \lambda D_s\, ds\Big] \\
&= \mathbb{E}\Big[\,\mathbb{E}\Big[-e^{-r\tau} C_\tau + \int_\tau^\infty e^{-rs} \lambda D_s\, ds \,\Big|\, \mathcal{F}_\tau\Big]\Big] \\
&= \mathbb{E}\Big[-e^{-r\tau} C_\tau + e^{-r\tau}\Big(\frac{\lambda D_\tau}{r} + \frac{\lambda\alpha}{r^2}\Big)\Big] \\
&= \mathbb{E}\big[e^{-r\tau} Z_\tau\big].
\end{aligned} \tag{2.38}$$

The second equality is the tower property, the third the strong Markov property, which restarts D from D_τ and so repeats the integration of Corollary 2.5 at the state prevailing at τ, and the fourth the definition (2.37) of Z. The two-factor problem is therefore equivalent to optimal stopping of the one-dimensional process Z with linear reward, and because (2.38) holds for every such τ, the equivalence costs no generality. The precommitted rules of Subsection 2.3.2 are the special case in which τ is a first-passage time of D alone.

In this form the problem is one the literature has already solved. Stopping a Lévy process with a linear reward is the Bachelier-model perpetual call solved by Mordecki (2002a), whose Theorem 1 identifies the optimal level with an expected supremum and whose Theorem 2 evaluates it in closed form when the positive jumps are mixed exponential. The present model is the case of a single exponential

[7] Started from z it evolves as $Z_t = z + \hat{Z}_t$, where $\hat{Z}$ is the same process started at zero. Nothing about the dynamics depends on the level, only on distances travelled, and this translation invariance will make the optimal trigger independent of the initial state.

component with zero strike. The hypotheses of that theorem are verified below, and its threshold is then translated back into the adoption boundary.

Both the threshold and the value function are expressed there in terms of the roots of a characteristic equation and the mixture weights those roots generate; the roots and the weights are computed for the present model below. The Lévy exponent of Z, defined by $\psi_Z(\beta) := \log \mathbb{E}[e^{\beta(Z_1-Z_0)}]$, splits over the Brownian and jump components, which are independent. The Gaussian factor is finite for every β; the compound Poisson factor is $\exp\{\nu(\mathbb{E}[e^{\beta\xi}]-1)\}$, and $\mathbb{E}[e^{\beta\xi}] = \gamma/(\gamma-\beta)$ is finite if and only if $\beta < \gamma$, since for $\beta \geq \gamma$ a single jump already has infinite exponential moment. The exponent is therefore finite exactly for $\beta < \gamma$, and given there by

$$\psi_Z(\beta) = \mu_Z\beta + \tfrac{1}{2}\sigma_Z^2\beta^2 + \nu\Big(\frac{\gamma}{\gamma-\beta} - 1\Big). \tag{2.39}$$

The right-hand side is a rational function of β and extends ψ_Z to every $\beta \neq \gamma$; roots of $\psi_Z(\beta) = r$ are understood throughout in the sense of this extension.[8] With this convention $\psi_Z(\beta) = r$ has exactly two positive roots $0 < \beta_1 < \gamma < \beta_2$, the one-sided case of the interlacing stated there and proved for the present cubic in Lemma B.5. Write

$$p = \frac{\beta_2(\gamma-\beta_1)}{\gamma(\beta_2-\beta_1)}, \qquad 1-p = \frac{\beta_1(\beta_2-\gamma)}{\gamma(\beta_2-\beta_1)}, \tag{2.40}$$

both of which lie in $(0,1)$ by the interlacing. These are the mixture weights A_1, A_2 of that paper's equation (13), evaluated at a single exponential rate $\alpha_1 = \gamma$.[9]

Three hypotheses have to hold: the discount rate is positive, $r > 0$; the Gaussian component is non-degenerate, $\sigma_Z = \lambda\sigma/r > 0$, as Theorem 2 there requires; and the expected supremum is finite, which for a compound Poisson jump component is exactly the finite-mean condition $\nu\,\mathbb{E}|\xi| < \infty$ assumed throughout. All three hold for every admissible calibration of the model.

Theorem 2.11. *Let* $0 < \beta_1 < \gamma < \beta_2$ *be the two positive roots of* $\psi_Z(\beta) = r$ *for the exponent* (2.39) *and let* p *be as in* (2.40). *Then adoption is optimal at*

$$z_{max} = \begin{cases} z^*, & z^* > z, \\ z, & \textit{otherwise}, \end{cases} \qquad \textit{where} \qquad z^* = \frac{p}{\beta_1} + \frac{1-p}{\beta_2}, \tag{2.41}$$

[8] The extension is harmless because the expectation is only ever evaluated where it is finite, that is at $\beta < \gamma$. The larger root $\beta_2 > \gamma$ never enters an expectation: it is a root of the cubic obtained by clearing the denominator, and appears afterwards only in finite expressions such as (2.41).

[9] The weights carry a probabilistic meaning. The supremum M of the process killed at rate r has density $p\,\beta_1 e^{-\beta_1 y} + (1-p)\,\beta_2 e^{-\beta_2 y}$: with probability p the highest level ever reached behaves as $\mathrm{Exp}(\beta_1)$, the diffusion-dominated component, and with probability $1-p$ as $\mathrm{Exp}(\beta_2)$, the overshoot component created by jumps that carry the process strictly above its previous record. The threshold (2.41) is the mean of this mixture, $z^* = \mathbb{E}M = p/\beta_1 + (1-p)/\beta_2$, and the interlacing guarantees that both weights are positive, so the decomposition is a proper probability density. As $\nu \to 0$ the jump component vanishes, $p \to 1$, and the mixture collapses to the classical exponential law of the supremum of a killed Brownian motion, the limit made precise in Proposition 2.12.

so that $\tau^* = \inf\{t \geq 0 : Z_t \geq z^*\}$ *and, in the original coordinates,* $\tau^* = \inf\{t \geq 0 : D_t \geq d^*(C_t)\}$ *with*

$$d^*(c) = \frac{r}{\lambda}\, c + \frac{r}{\lambda} z^* - \frac{\alpha}{r}. \tag{2.42}$$

The corresponding value is

$$V(z) = \begin{cases} \dfrac{p}{\beta_1}\, e^{-\beta_1(z^*-z)} + \dfrac{1-p}{\beta_2}\, e^{-\beta_2(z^*-z)}, & z^* > z, \\ z, & \textit{otherwise.} \end{cases} \tag{2.43}$$

Proof. Equations (2.41) and (2.43) are Theorems 1 and 2 of Mordecki (2002a) for a single exponential jump component and zero strike, under the hypotheses verified above, with the mixture weights (2.40) substituted for the A_k of that paper. It remains only to change coordinates: substituting $Z_t = \lambda D_t/r + \lambda\alpha/r^2 - C_t$ into the rule $Z_t \geq z^*$ and solving for D_t,

$$\frac{\lambda}{r} D_t \geq z^* + C_t - \frac{\lambda\alpha}{r^2} \qquad \Longleftrightarrow \qquad D_t \geq \frac{r}{\lambda} C_t + \frac{r}{\lambda} z^* - \frac{\alpha}{r},$$

the right-hand side being $d^*(C_t)$ of (2.42). □

The level z^* admits an equivalent form that is often more convenient numerically. Substituting the weights (2.40) into (2.41) and factoring the numerator over the common denominator $\gamma\beta_1\beta_2(\beta_2 - \beta_1)$,

$$\begin{aligned} z^* &= \frac{p}{\beta_1} + \frac{1-p}{\beta_2} \\ &= \frac{(\beta_2 - \beta_1)\big(\gamma(\beta_1 + \beta_2) - \beta_1\beta_2\big)}{\gamma\beta_1\beta_2(\beta_2 - \beta_1)} \\ &= \frac{1}{\beta_1} + \frac{1}{\beta_2} - \frac{1}{\gamma}. \end{aligned} \tag{2.44}$$

This is also what the pasting conditions deliver in Appendix B.9.1.

Unlike the constant-cost threshold (2.19), which is an explicit formula in α, σ, r, C and λ, the boundary depends on the parameters through the two roots. These are explicit in the following sense. Multiplying $\psi_Z(\beta) = r$ by $\gamma - \beta$ turns the root condition into the cubic of Lemma B.5,

$$(\gamma - \beta)\Big(\tfrac{1}{2}\sigma_Z^2\beta^2 + \mu_Z\beta - (\nu + r)\Big) + \nu\gamma = 0, \tag{2.45}$$

whose three real roots are one negative root and the pair $0 < \beta_1 < \gamma < \beta_2$. A cubic is solvable in radicals, so β_1 and β_2, and with them p, z^* and d^*, are closed-form algebraic functions of the primitives, although the radical expressions for the roots are too unwieldy to be informative and are not reproduced here. In practice the interlacing supplies guaranteed brackets instead: ψ_Z increases strictly on $(0, \gamma)$ from $\psi_Z(0) = 0 < r$ to $+\infty$ at the pole, and on (γ, ∞) from $-\infty$ to $+\infty$, so each root is the unique sign change of $\psi_Z(\beta) - r$ in its interval and bisection, or safeguarded

Newton, converges unconditionally. Equivalently, the three roots of (2.45) can be obtained from any polynomial root finder and identified by the interlacing. Once the roots are in hand, p, z^* and d^* follow from (2.40), (2.41) and (2.42) by arithmetic alone; adoption occurs the first time the state reaches the boundary $d^*(C_t)$. The rule is stationary in calendar time and linear in the prevailing cost, the hurdle on D falling at rate r/λ as C declines.

Setting $\nu = 0$ switches the jump off, so the constant-cost model of Section 2.2 should reappear as a limiting case. This cannot be read off the value function directly. The exponents β_1 and β_2 and the mixture weight p are defined implicitly by the cubic (2.45), so the limit has to be taken on the roots themselves before the resulting threshold can be compared with Theorem 2.6. The following proposition carries this out and establishes the nesting.

Proposition 2.12 (The constant-cost model is the case $\nu = 0$)**.** *Assume $\beta_0 < \gamma$, where β_0 is the positive root of the no-jump exponent. As $\nu \downarrow 0$ the roots satisfy $\beta_1 \to \beta_0$ and $\beta_2 \to \gamma$, and the mixture weight satisfies $p \to 1$. Consequently the optimal level becomes*

$$z_0^* = \frac{1}{\beta_0} = \frac{\lambda}{r}\,\frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha}, \tag{2.46}$$

and the boundary evaluated at the initial cost is the constant-cost threshold, $d^*(C_0) = x^{**}$. The value (2.43)*, evaluated at the initial state, reduces to the V_{max} of Theorem 2.6, which is the trigger formula* (2.16) *at the optimal trigger. Theorem 2.6 is therefore Theorem 2.11 at $\nu = 0$, evaluated at the initial cost.*

Proof. The jump term drops out of (2.39), so $\psi_Z(\beta) = r$ becomes the quadratic $\frac{1}{2}\sigma_Z^2\beta^2 + \mu_Z\beta - r = 0$, whose positive root is

$$\beta_0 = \frac{\sqrt{\mu_Z^2 + 2r\sigma_Z^2} - \mu_Z}{\sigma_Z^2}. \tag{2.47}$$

The coefficients of the cubic (2.45) are affine in ν, and the roots of a polynomial depend continuously on its coefficients, so as $\nu \downarrow 0$ the two positive roots converge to the positive roots of the limit cubic $(\gamma - \beta)(\frac{1}{2}\sigma_Z^2\beta^2 + \mu_Z\beta - r)$, namely β_0 and γ; under the standing assumption $\beta_0 < \gamma$ the interlacing $\beta_1 < \gamma < \beta_2$ forces $\beta_1 \to \beta_0$ and $\beta_2 \to \gamma$.[10] The mixture weight (2.40) then satisfies

$$p = \frac{\beta_2(\gamma - \beta_1)}{\gamma(\beta_2 - \beta_1)} \longrightarrow \frac{\gamma(\gamma - \beta_0)}{\gamma(\gamma - \beta_0)} = 1,$$

so the two-term mixture in (2.43) collapses to the single exponential $(1/\beta_0)\, e^{-\beta_0(z^* - z)}$, and by (2.44) the level satisfies $z^* = 1/\beta_1 + 1/\beta_2 - 1/\gamma \to 1/\beta_0$,

[10] In the complementary case $\beta_0 \geq \gamma$ the labels interchange and $p \to 0$; the threshold and the collapsed value are unchanged, but a general proof has to track this relabelling and is not given here.

which is (2.46). Substituting $\mu_Z = \lambda\alpha/r$ and $\sigma_Z = \lambda\sigma/r$ into (2.47),

$$\beta_0 = \frac{(\lambda/r)\sqrt{\alpha^2 + 2r\sigma^2} - \lambda\alpha/r}{\lambda^2\sigma^2/r^2} = \frac{r}{\lambda} \cdot \frac{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha}{\sigma^2} = -\frac{r}{\lambda} G$$

for the constant G of (2.21), so the collapsed value function is exactly the Brownian transform (2.15) written in the Z coordinate, the factor r/λ converting distances in D into distances in Z; evaluated at the initial state it is the trigger formula (2.16) at the optimal trigger, that is the V_{max} of Theorem 2.6; the prefactors agree because the state (C_0, x^{**}) lies on the boundary, where the adoption payoff equals $z_0^* = 1/\beta_0$ by the definition of Z. Finally, substituting z_0^* into (2.42) turns the intercept $(r/\lambda)z_0^*$ into $\sigma^2/(\sqrt{\alpha^2 + 2r\sigma^2} - \alpha)$, so that $d^*(C_0)$ reproduces (2.19). □

Theorem 2.11 is stated for stopping times of Z, whereas the firm in (2.36) may condition on the whole history of (C, D), which is a priori more information. That extra information turns out to be worthless, because Z remains a Lévy process in the larger filtration.

Proposition 2.13 (Global optimality)**.** *The rule τ^* of Theorem 2.11 is optimal among all stopping times of the filtration generated by (C, D), and the stopping region is the half-line $[z^*, \infty)$.*

Proof. Both W and Y are adapted to the filtration generated by (C, D) and retain stationary independent increments with respect to it, so Z is a Lévy process in that larger filtration and not merely in its own. Theorem 1 of Mordecki (2002a) applies verbatim with the larger filtration in place of $\mathcal{F}^Z$, and by (2.38) the reward of every (C, D)-stopping time is $\mathbb{E}[e^{-r\tau} Z_\tau]$. The supremum over the larger class is therefore attained at the same first-passage rule, and observing C and D separately confers no advantage over observing the payoff Z alone. □

The analysis so far has been phrased as a stopping rule: a level to wait for, and a boundary in the state plane. Chapter 3 approaches the two-factor problem in the complementary language of free-boundary partial differential equations solved on a grid, because once the fixed cost is itself stochastic, the problem no longer collapses to a single payoff level. The two descriptions meet as follows, which allows the results of the two chapters to be compared.

Equivalently, the value function solves a quasi-variational inequality on the two-dimensional state. Write

$$\mathcal{V}(c, d) := V\big(\hat{g}(c, d)\big) \tag{2.48}$$

for the value expressed in the original coordinates, where V is the one-dimensional value (2.43) of the payoff level and $\hat{g}(c, d)$ is the immediate-adoption payoff of Corollary 2.5 evaluated at the prevailing state. The two agree by construction, since

$\hat{g}(C_t, D_t) = Z_t$ is exactly the payoff process of (2.37); the calligraphic letter only distinguishes the surface $\mathcal{V}(c, d)$ from the function $V(z)$ of a single argument. Then

$$\max\{(\mathcal{L} - r)\mathcal{V},\ \hat{g} - \mathcal{V}\} = 0, \tag{2.49}$$

where $\mathcal{L}$ is the infinitesimal generator of the jump diffusion (C, D),

$$\mathcal{L}\mathcal{V} = \alpha\,\mathcal{V}_d + \tfrac{1}{2}\sigma^2\mathcal{V}_{dd} + \nu\,\mathbb{E}\big[\mathcal{V}(c - \xi, d) - \mathcal{V}(c, d)\big], \tag{2.50}$$

where subscripts denote partial derivatives. The first two terms of (2.50) are the diffusion generator $\mathcal{L}_D$ of Section 2.1. The integral term is the standard compensator of a compound Poisson component, weighting the change in value induced by a jump of size ξ by the arrival rate ν (see, e.g., Øksendal and Sulem, 2007; Kou, 2002). This is the one-factor, jump-augmented instance of the free-boundary problem solved numerically in Chapter 3: supplementing $\mathcal{L}_D$ with the jump operator turns the variational inequality into a quasi-variational one, and the benefit stream enters only through the obstacle $\hat{g}(c, d) = \lambda d/r + \lambda\alpha/r^2 - c$.

The closed form (2.43), extended to the state plane by (2.48), therefore offers a benchmark for the numerical scheme of Chapter 3 in this instance, on the entire value surface and not only on the location of the free boundary.

The characterisation of z^* itself requires no exponential assumption. Theorem 1 of Mordecki (2002a) holds for any jump law with $\mathbb{E}M < \infty$, that is whenever the upward jumps have finite mean, and gives

$$z^* = \mathbb{E}\Big[\sup_{t \le e_r}(Z_t - Z_0)\Big], \tag{2.51}$$

where e_r is an exponential time of rate r independent of Z. The clock carries the discounting and adds no assumption to the model: a payoff at time t is discounted by e^{-rt}, which is precisely the probability $\mathbb{P}(e_r > t)$ that such a clock has not yet rung by t, so a discounted payoff and a payoff collected only when it arrives before the clock have the same expected value. The window has mean length $1/r$, and (2.51) then says that the firm holds out until the advantage in hand matches the largest gain it could still expect to see within that window. Read economically, the clock is a constant chance per unit of time that the opportunity ceases to be available. The linear boundary and its economic reading therefore survive well beyond the exponential case, and only the closed form (2.41) is specific to it. Under heavier upward tails the expectation diverges and adoption is never optimal, the prospect of arbitrarily large cost reductions making waiting worth more than any finite advantage.[11]

A constructive verification of z^* for the exponential specification, through value matching, smooth pasting, and the integro-differential pasting condition governing

[11] The representation of the optimal level as an expected supremum goes back to Darling et al. (1972) for random walks; for general Lévy processes it appears in Mordecki (2002b) and Novikov and Shiryaev (2004), in the geometric and power-reward cases respectively.

jumps across the boundary, is carried out in Appendix B.9.1.[12] The representation (2.51) itself yields two facts on which the rest of this subsection rests: a sign for z^*, and a change of variables that turns the average over paths into an integral over distances.

The sign of z^* follows from the representation directly: the supremum in (2.51) runs over an interval containing $t = 0$, where the increment $Z_t - Z_0$ is zero, so the supremum is nonnegative along every path and so is its expectation, giving $z^* \geq 0$ for every admissible jump distribution.

The change of variables is the layer-cake formula. Write

$$\tau_h := \inf\{t \geq 0 : \ Z_t - Z_0 \geq h\}, \qquad \Lambda(h) := \mathbb{E}\big[e^{-r\tau_h}\big]$$

for the first time Z gains a distance h over its starting point and for the expected discount factor of that passage, and let e_r be the exponential clock of rate r in (2.51), independent of Z. Then

$$\begin{aligned} z^* &= \mathbb{E}\Big[\sup_{t \leq e_r}(Z_t - Z_0)\Big] \\ &= \int_0^\infty \mathbb{P}\Big(\sup_{t \leq e_r}(Z_t - Z_0) > h\Big)\,\mathrm{d}h \\ &= \int_0^\infty \mathbb{P}(\tau_h < e_r)\,\mathrm{d}h \\ &= \int_0^\infty \mathbb{E}\big[e^{-r\tau_h}\big]\,\mathrm{d}h = \int_0^\infty \Lambda(h)\,\mathrm{d}h. \end{aligned} \tag{2.52}$$

The second equality is the layer-cake formula $\mathbb{E}X = \int_0^\infty \mathbb{P}(X > h)\,\mathrm{d}h$ for a nonnegative random variable, the third the identity of events $\{\sup_{t \leq e_r}(Z_t - Z_0) > h\} = \{\tau_h < e_r\}$, and the fourth conditioning on τ_h and using the independence of the clock, $\mathbb{P}(e_r > \tau_h \mid \tau_h) = e^{-r\tau_h}$.

The optimal level thus aggregates the discount factors of reaching every distance. Each layer h contributes what it is worth to the firm to still be waiting when the advantage has grown by h, and the hurdle is the sum of those contributions over all sizes of advantage. Since the law of Z enters (2.52) only through Λ, the dynamics move z^* through a single channel, the speed at which Z covers ground: a specification that reaches a given distance sooner raises $\Lambda(h)$ at every h, and raises z^* with it. Three properties of the adoption boundary follow from this: its slope, its position relative to the no-jump case, and its response to a jump that has already occurred.

The slope is fixed by the payoff alone: because the stopping region is the half-line $[z^*, \infty)$ in the payoff variable, the boundary in the (c, d)-plane is the level set

[12] Smooth pasting is legitimate here because the trigger is *regular* for the stopping region: started exactly at the boundary, Z enters (z^*, ∞) immediately, almost surely. Since Z is spatially homogeneous, this is equivalent to the origin being regular for $(0, \infty)$, and it holds because $\sigma > 0$ gives Z a Brownian component, which oscillates across any level instantly. Were the boundary reachable only by jumps, regularity could fail and only continuous fit would be guaranteed (Alili and Kyprianou, 2005). The third pasting condition is integro-differential because upward jumps overshoot the boundary rather than crossing it continuously.

$\hat{g}(c,d) = z^*$, and its slope is therefore the slope of an iso-payoff line. With $\hat{g}_c = -1$ and $\hat{g}_d = \lambda/r$, the implicit function theorem gives $-\hat{g}_c/\hat{g}_d = r/\lambda$, a calculation into which no feature of the dynamics enters. For any specification under which Z remains a Lévy process with upward jumps of finite mean, the boundary is a line of this same slope, and the dynamics act on the intercept alone, through (2.51). Curvature could arise only if spatial homogeneity failed, for instance under state-dependent jump intensity such as learning-by-doing, a case outside the present framework. The sign obtained above then locates the whole line: since $z^* \geq 0$, the boundary never falls below the zero-NPV line along which immediate adoption breaks even.

Jump risk acts on that intercept, and because a change in the jump specification reaches z^* only through Λ, its direction can be read off (2.52) directly. Cost-reducing jumps raise the threshold, because upward jumps carry Z to any given level sooner and so raise $\Lambda(h)$ at every distance h, and with it the integral: $z^* > z_0^*$. Since the slope is unchanged, the boundary lies uniformly above its no-jump counterpart, shifted upward in parallel by $(r/\lambda)(z^* - z_0^*)$.

Cost-increase risk works the other way round: when the jumps of the fixed cost are upward, Z has downward jumps and therefore crosses every level continuously, so the overshoot that lifted Λ in the previous case is absent. The threshold falls to $z^* = 1/\Phi_Z(r) < z_0^*$, where $\Phi_Z(r)$ is the unique positive root of the corresponding exponent equation. Both statements are proved in Appendix B.9.2, and their economic content is discussed in Subsection 2.3.3.

A realised jump, finally, requires no separate argument, since the boundary (2.42) is linear in the prevailing cost: substituting $c + \Delta C$ shows that a cost cut $\Delta C < 0$ lowers the hurdle on D by $(r/\lambda)|\Delta C|$ at the instant it occurs.

2.3.2 A precommitted first-passage benchmark

The closed-loop rule of Theorem 2.11 responds to both state variables. It is nevertheless useful to analyse the best rule within the restricted class of *precommitted* first-passage triggers on D alone. Three reasons motivate the detour: the benchmark yields fully explicit formulas valid for *any* jump-size distribution with finite mean; it is the specification estimated in the published article underlying this chapter; and the comparison isolates, both analytically and numerically, the value of being able to react to realised cost jumps.

Restricting τ in (2.36) to a first-passage time of D, the profit stream from switching at the first time $D(t)$ reaches a threshold x^* is given by

$$V_{D_{x^*},C} = \mathbb{E}_{D_0,C_0}\left[-e^{-r\tau_{x^*}}\, C(\tau_{x^*}) + \int_{\tau_{x^*}}^{\infty} e^{-rt}\lambda\, D(t)\, dt\right], \tag{2.53}$$

where

$$\tau_{x^*} = D_{x^*} = \inf\{t \geq 0 : D(t) \geq x^*\} \tag{2.54}$$

is the first hitting time of the level x^* by the process $D(t)$ and $\mathbb{E}_{D_0,C_0}$ denotes expectation with respect to the probability measure under which $D(0) = D_0$ and $C(0) = C_0$.

The threshold x^* is fixed at time zero, so the rule cannot respond to realised jumps in the cost process C: a sudden cost reduction cannot trigger adoption before D reaches x^*. The results below characterise the best rule within this restricted class, to be compared with the unrestricted optimum of Theorem 2.11.

For $s \geq 0$

$$\mathbb{E}_{D_0,C_0}[Y(s)] = \nu\, d_e\, s, \tag{2.55}$$

where ν is the intensity of the Poisson process, and d_e is the average jump size.

Theorem 2.14. *Assume that the jump process $Y(t)$ is independent of the diffusion $D(t)$, that $\nu\,\mathbb{E}|\xi| < \infty$ for the jump sizes ξ, and that $x^* \geq D_0$. For the precommitted first-passage rule $\tau_{x^*} = \inf\{t \geq 0 : D(t) = x^*\}$, the profit stream $V_{D_{x^*},C}$ is given by*

$$V_{D_{x^*},C} = \left[\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C_0 - \frac{\nu\, d_e\,(x^* - D_0)}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}\right] \times \exp\left\{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma}(x^* - D_0) - (x^* - D_0)\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right)\right\}. \tag{2.56}$$

Proof. See Appendix B.3. □

Equation (2.56) modifies the no-jump value in a structured way. The baseline term $\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C_0$ is the value in the absence of jumps. The adjustment $\frac{\nu d_e(x^*-D_0)}{\sigma\sqrt{\alpha^2/\sigma^2+2r}}$ equals $\nu d_e\,\mathbb{E}[\tau_{x^*}e^{-r\tau_{x^*}}]/\mathbb{E}[e^{-r\tau_{x^*}}]$, the expected cumulative jump drift in the cost at the discount-weighted adoption date. For cost-reducing jumps ($d_e < 0$) the adjustment is negative and, entering with a minus sign, *raises* the value of every rule with $x^* > D_0$: expected costs at adoption are lower.

The optimal precommitted threshold is characterised as follows.

Theorem 2.15. *For a given set of parameters, assuming that $\tilde{A} := \frac{\lambda}{r} - \frac{\nu\, d_e}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2}+2r}} > 0$, $V_{D_{x^*},C}$ reaches its supremum over the precommitted rules $x^* \geq D_0$ at a specific point*

$$x_{max} = \begin{cases} x^{***}, & x^{***} > D_0. \\ D_0, & \textit{otherwise}, \end{cases} \tag{2.57}$$

where

$$x^{***} := \frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\dfrac{\lambda\alpha}{r^2} - C_0 + \dfrac{\nu\, d_e\, D_0}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}}{\dfrac{\lambda}{r} - \dfrac{\nu\, d_e}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}}. \tag{2.58}$$

The corresponding maximum value can be characterised as

$$V_{max,C} := \sup_{x^*\geq D_0} V_{D_{x^*},C} = \begin{cases} V_{D_{x^{***}},C} & x^{***} > D_0. \\ \frac{\lambda D_0}{r} + \frac{\lambda\alpha}{r^2} - C_0, & \textit{otherwise.} \end{cases} \tag{2.59}$$

where

$$V_{D_{x^{***}},C} = \left[\frac{\lambda x^{***}}{r} + \frac{\lambda\alpha}{r^2} - C_0 - \frac{\nu\, d_e\,(x^{***} - D_0)}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}\right] \times \exp\left\{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma}(x^{***} - D_0) - (x^{***} - D_0)\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right)\right\}. \tag{2.60}$$

Proof. See Appendix B.4. □

Remark 2.16. Theorem 2.15 optimises within the class of precommitted first-passage rules on D alone. Such rules cannot respond to realised cost jumps, and the resulting threshold x^{***} depends on the initial state D_0, which a dynamic-programming optimum of a Markov stopping problem cannot. Subsection 2.3.1 solves the unrestricted problem. The threshold x^{***} coincides with the closed-loop solution when $\nu d_e = 0$.

The hypothesis $\tilde{A} > 0$ is mild: it holds unconditionally for cost-reducing jumps ($d_e < 0$), and for cost-increasing jumps it holds whenever usage intensity is moderate to high and the remaining parameters take standard values. In the degenerate case $\nu d_e = 0$ the threshold collapses to $x^{***} = x^{**}$, and for $d_e < 0$ every precommitted rule with $x^* \geq D_0$ is worth more than its no-jump counterpart, $V_{D_{x^*},C} \geq V_{D_{x^*}}$.

The effect on the precommitted threshold is signed by the net present value of immediate adoption at the initial state:

Proposition 2.17. *Under the assumptions of Theorem 2.15,*

$$\frac{\partial x^{***}}{\partial(\nu d_e)} = -\frac{\Pi_\infty(D_0, C_0)}{\tilde{A}^2\,\sigma\sqrt{\alpha^2/\sigma^2 + 2r}}, \qquad \Pi_\infty(D_0, C_0) = \frac{\lambda D_0}{r} + \frac{\lambda\alpha}{r^2} - C_0.$$

In particular, if immediate adoption has negative net present value at (D_0, C_0), cost-reducing jumps lower x^{***} and cost-increasing jumps raise it.

Proof. Write $J = \nu d_e/(\sigma\sqrt{\alpha^2/\sigma^2 + 2r})$ for the jump adjustment, $A_0 = \lambda/r$ and $B_0 = \lambda\alpha/r^2 - C_0$, so that $\tilde{A} = A_0 - J$ and (2.58) reads

$$x^{***} = -\frac{1}{G} - \frac{B_0 + JD_0}{A_0 - J},$$

with G the constant of (2.21), which involves only α, σ and r. The first term is therefore free of J. Differentiating the second by the quotient rule,

$$\frac{\partial}{\partial J}\frac{B_0 + JD_0}{A_0 - J} = \frac{D_0(A_0 - J) + (B_0 + JD_0)}{(A_0 - J)^2} = \frac{A_0 D_0 + B_0}{\tilde{A}^2},$$

the two terms in J cancelling in the numerator. Since $A_0 D_0 + B_0 = \lambda D_0/r + \lambda\alpha/r^2 - C_0 = \Pi_\infty(D_0, C_0)$, this gives

$$\frac{\partial x^{***}}{\partial J} = -\frac{\Pi_\infty(D_0, C_0)}{\tilde{A}^2},$$

and the chain rule with $\partial J/\partial(\nu d_e) = 1/(\sigma\sqrt{\alpha^2/\sigma^2 + 2r})$ yields the stated derivative. The sign claim follows because $\tilde{A}^2$ and $\sigma\sqrt{\alpha^2/\sigma^2 + 2r}$ are both positive, so the sign of $\partial x^{***}/\partial(\nu d_e)$ is that of $-\Pi_\infty(D_0, C_0)$. □

Two points follow. First, for the calibration of this chapter (and, more generally, throughout the region where immediate adoption has positive net present value) $\Pi_\infty(D_0, C_0) > 0$, so the precommitted and closed-loop rules share the same qualitative comparative statics: anticipated cost reductions raise the adoption boundary and anticipated cost increases lower it. The realistic parameter range is comfortably inside this region: $\Pi_\infty < 0$ would require an implausibly negative initial advantage.[13] Second, the agreement is not guaranteed in general: the precommitted sign is $-\operatorname{sign}\Pi_\infty$, so it can *reverse* when immediate adoption has negative net present value, whereas the closed-loop rule delivers the standard real-options sign unconditionally. Subsection 2.3.3 quantifies the difference in magnitude and the value of reacting to *realised* jumps.

For a fixed combined jump drift νd_e, the jump adjustment term scales with σ^{-1}, so higher volatility dampens the impact of jumps on the precommitted threshold. In highly volatile environments, discrete innovations matter less than continuous uncertainty. Higher drift α affects both the base threshold and jump adjustment, amplifying or attenuating responses depending on parameter values.

2.3.3 Comparative Analysis of Economic Implications

With both solution concepts in hand, the economic content of the jump extension can be read off. The closed-loop solution delivers the comparative statics of adoption timing in a transparent geometric form. The adoption boundary $d^*(c)$ is linear in the prevailing cost with slope r/λ for any Lévy jump specification with integrable upward jumps, so jump risk moves only the intercept, through the optimal level z^*. Anticipated cost reductions raise z^* (Proposition B.8) and shift the boundary up uniformly: a strictly better state is required to justify committing, because part of the value of waiting now consists of the chance that a breakthrough makes adoption cheaper. The effect on expected adoption *timing* is nevertheless ambiguous. The same jumps that raise the hurdle also raise the drift of the payoff process Z from μ_Z to $\mu_Z + \nu/\gamma$, so the state approaches the boundary faster even as the boundary recedes. Which force dominates depends on the parameters, so a higher hurdle does not by itself imply later adoption.

[13] For the delivery-fleet calibration of Section 2.4, $\Pi_\infty < 0$ requires $D_0 < -11$, while the estimated values lie between 1.45 and 7.94.

The threat of cost *increases* works in the opposite direction. When the only jumps are upward movements in the fixed cost, for example, the anticipated withdrawal of a purchase subsidy, Proposition B.9 shows that the boundary shifts downward and adoption accelerates: the firm commits at operating advantages it would otherwise consider insufficient, because waiting now carries the risk of losing the current favourable cost conditions.

Realised jumps operate through a separate channel: the instantaneous hurdle reduction of $(r/\lambda)|\Delta C|$ recorded in Subsection 2.3.1 triggers adoption on the spot whenever it carries the state across the boundary. Adoption therefore occurs through two distinguishable routes. Under diffusive adoption, Z reaches z^* continuously and the realised surplus $Z_{\tau^*} - z^*$ is zero. Under jump adoption, a cost reduction carries Z strictly across the boundary and, by the memorylessness of the exponential jump size, the realised surplus is $\mathrm{Exp}(\gamma)$. The model thus predicts a mixed distribution of adoption surplus, a mass point at zero generated by gradual growth of the operating advantage plus an exponential tail generated by discrete cost announcements. The precommitted benchmark of Subsection 2.3.2 lacks this immediate-response channel. In the empirically relevant region the benchmark agrees with the closed-loop rule on the *direction* of the response to anticipated jumps, but the closed-loop rule improves on it in the magnitude of that response and, decisively, in the ability to act on jumps once they occur.

These results have policy implications: credible announcements of future cost reductions delay near-term adoption, creating a trade-off between signalling innovation and encouraging immediate deployment, while anticipated subsidy removals accelerate adoption, creating "adoption cliffs" around policy expiration dates. Preannouncing technological breakthroughs can slow current adoption as firms wait for superior options. Gradual, predictable subsidy phase-outs may therefore support smoother transitions than abrupt policy changes.

2.4 Empirical Analysis

In this section, the theoretical model is applied to an empirical example, demonstrating its relevance and practical utility in guiding electric vehicle (EV) adoption strategies. First, in Section 2.4.1, a detailed case study is presented by specifying all model parameters. This case study focuses on a delivery company considering the replacement of its internal combustion engine (ICE) vehicle fleet with EV alternatives of comparable types. The company operates in three cities (San Francisco, Seattle, and Miami) each featuring regional variations in the baseline unit cost advantage. Through the examination of twin scenarios, the impact of positive and negative jumps is further elucidated by modelling the effects of either the withdrawal of tax incentives or the anticipation of technological advancements that reduce the fixed cost of adopting the EV. By selecting a representative example and analysing varying conditions, the analysis explores where and when the circumstances are most favourable for EV adoption, taking individual parameters into account.

These insights reinforce the model's policy implications, emphasising the need for stable economic environments to support timely and efficient technology transitions. Additionally, the results suggest that users may delay technology adoption if they anticipate future advancements leading to more favourable EV pricing, thereby emphasising the importance of strategically planning policy initiatives supported by an in-depth understanding of the factors affecting technology adoption timing and decisions.

Section 2.4.2 provides a detailed analysis of model calibration using historical data on EV unit cost advantages and volatility. This examination highlights how fixed cost, usage intensity, drift, and volatility influence adoption timing, demonstrating the model's applicability across diverse scenarios. It is important to emphasise the inherent complexity and critical importance of accurately estimating forward-looking parameters, given the substantial uncertainty surrounding future technology costs and market developments. Consequently, incorporating individual views and expert judgment becomes essential to effectively capture scenario-specific expectations and enhance decision-making accuracy.

2.4.1 Case Study: Delivery Company's Fleet Replacement

The analysis is refined by examining a case study in which a delivery company evaluates the decision to replace its fleet of vehicles. This involves the precise choice and calibration of all the model parameters, including scenarios in which the company anticipates jumps in the fixed cost process $C(t)$. To ensure a comprehensive assessment, five distinct scenarios are explained, each representing different types of economic conditions. These scenarios are designed to provide a nuanced understanding of optimal switching decisions under varying circumstances for the company.

In this context, a delivery company is considered operating within an industry characterised by an average annual mileage of 48,000 miles per vehicle.[14] Given the nature of its business, the company requires compact, cost-efficient vehicles that support rapid urban deliveries. The decision involves replacing the current ICE model with a new ICE model or a new EV model, necessitating a choice between the following two models presented in Table 3.

The Manufacturer's Suggested Retail Price (MSRP) is used as a reference, although actual purchase prices may vary. The Commercial Clean Vehicle Credit (Internal Revenue Code Section 45W) provided a 30% federal tax credit up to

[14] `https://www.caltrux.org/driver-faqs/`

Vehicle Type	EV	Price [$]	Subsidy [$]	Efficiency	Size
Nissan Sentra	No	25,000	No	2.9 gal/100 mi	Mid-size sedan
Nissan Leaf	Yes	38,300	7,500	31 kWh/100 mi	Compact

Table 3: Comparison of ICE and EV alternatives considered by the delivery company. Pricing and specifications were retrieved from Nissan's official inventory page as of February 2025. Source: `https://www.nissanusa.com/shopping-tools/search-inventory`.

$7,500 for qualifying EVs at the time of calibration, effectively lowering the Nissan Leaf's acquisition cost.[15]

The company adopts a moderately optimistic yet empirically substantiated outlook on future energy market conditions, contingent on the assumption that the long-term trends identified in the U.S. energy market since 2015 will remain consistent. In line with this perspective, the firm systematically integrates historical drift and volatility parameters into its decision-making framework. More details about the calibration of volatility and drift can be found in Section 2.4.2. The unit cost advantage (D_0) is calculated using regional gasoline and electricity prices as of December 2024, based on data from the U.S. Bureau of Labor Statistics.[16]

For instance, in San Francisco, the unit cost advantage is

$$\begin{aligned} D_0 &= \big(2.9\,\text{gal}/100\,\text{mi}\big) \cdot 4.456\,\frac{\$}{\text{gal}} \;-\; \big(31\,\text{kWh}/100\,\text{mi}\big) \cdot 0.37\,\frac{\$}{\text{kWh}} \\ &= 12.92 \;-\; 11.47 \;=\; 1.45\,\$/100mi. \end{aligned}$$

Thus, under current price assumptions, the Leaf is $1.45 cheaper per 100 miles than the Sentra in San Francisco. The corresponding parameters without jump components are reported in Table 4.

Five distinct cases are examined. The company operates in three different locations, with two scenarios explicitly considering jumps in the fixed cost process. Table 5 summarises the adjusted parameters along with the optimal levels for each scenario: the no-jump threshold x^{**} where applicable, and the closed-loop adoption boundary $d^*(C_0)$ for the jump scenarios, evaluated from Theorem 2.11 in Scenario C and from Proposition B.9 in Scenario E, where the cost jumps upward instead. Scenario A uses the benchmark parameters of Table 4. Scenarios B and D differ from

[15] Additional California-specific benefits, such as High-Occupancy Vehicle (HOV) lane access, may provide operational advantages but are not explicitly monetised in this analysis. Moreover, it is assumed that the company has sufficient tax liability to fully utilise the federal tax credit, making the reduction applicable. The Section 45W credit was subsequently repealed by the One Big Beautiful Bill Act (Pub. L. 119-21, 4 July 2025, §70503), which terminated it for vehicles acquired after 30 September 2025. The parameters of this chapter are retained as originally published in Szabó et al. (2025) and reflect the policy environment of February 2025. Scenario E, which treats credit withdrawal as an anticipated positive jump in $C(t)$, is therefore an ex-ante analysis of an event that has since been realised. The ex-post evidence is discussed in Section 5.3.

[16] `https://www.bls.gov/regions/midwest/data/averageenergyprices_selectedareas_table.htm`

Variable	Parameter value [unit]	Source
α	0.34	Historical drift of $D(t)$
σ	2.42	Historical volatility of $D(t)$
λ	480 [x 100 miles per year]	Average annual mileage for delivery
C_0	5,800 [$]	Fixed cost
r	3%	Discount rate
D_0	$1.45	Automobile specific unit cost advantage

Table 4: Analysis Parameters for the Delivery Company. The company estimated a historical volatility and drift from the US oil and electricity prices for the period from 2015 to 2023. More details about the calibration of volatility and drift can be found in Section 2.4.2. Fixed cost is estimated from actual listed inventory prices for the corresponding models minus the Federal Tax Credit. The unit cost advantage (D_0) is calculated from the two automobile model-specific efficiency numbers as reported by the manufacturer and summarised in Table 3, and San Francisco-specific average energy prices in December 2024 from the U.S. Bureau of Labor Statistics.

Scenario A only in the company's location, and hence in the current unit cost advantage D_0. Scenarios C and E add jump risk to B and D respectively. The economic story behind each jump determines both its sign and the Lévy specification used to price it:

- Scenario A: Fleet replacement in San Francisco under current conditions.
- Scenario B: Fleet replacement in Seattle, where lower electricity prices raise the unit cost advantage relative to California.
- Scenario C: Seattle, with the firm now *anticipating a stream of technological breakthroughs* (falling battery prices, new entrants) that gradually reduce the EV premium. This is a cost-*reducing* risk, modelled as a Poisson process of exponentially distributed downward jumps in $C(t)$ at rate $\nu = 1$ with mean size $|d_e| = \$3{,}325$. The calibration is set out in the notes to Table 5.[17]
- Scenario D: Fleet replacement in Miami under current conditions.
- Scenario E: Miami, with the firm now *facing the risk of losing the federal EV tax credit*. This is a cost-*increasing* risk, modelled as a Poisson process of fixed upward jumps in $C(t)$ of size $d_e = \$7{,}500$ (the credit amount) arriving at hazard rate $\nu = \frac{1}{8}$. The calibration is set out in the notes to Table 5.

In Scenario A, even a simplified linear calculation, which neglects discounting and stochastic effects, indicates an approximate annual saving of $700 from an immediate EV transition. However, this saving alone does not offset the higher initial

[17]The industry projection of EV–ICE price parity around 2029 is used to anchor the expected annual cost decline. Source: `https://insideevs.com/news/709746/lower-battery-prices-to-make-ev-costs-equal-gas-cars/`.

Variable	Scenario A	Scenario B	Scenario C	Scenario D	Scenario E
Location	San Francisco	Seattle	Seattle	Miami	Miami
C_0	5,800	5,800	5,800	5,800	5,800
D_0	1.45	7.94	7.94	4.8	4.8
ν	0	0	1	0	$\frac{1}{8}$
d_e	-	-	-3,325	-	7,500
x^{**} or $d^*(C_0)$	6.08	6.08	11.66	6.08	4.68
Optimal Decision	ICE	EV	ICE	ICE	EV

Table 5: Key variables for the delivery company under different scenarios. The fixed cost of EV adoption (C_0) is the after-credit price premium of the EV model. The unit cost advantage (D_0) is calculated from regional gasoline and electricity prices as of December 2024 (U.S. Bureau of Labor Statistics). The jump intensity ν is the expected annual frequency of cost jumps, and d_e their signed size ($d_e < 0$ for cost-reducing breakthroughs, $d_e > 0$ for cost-increasing policy shocks): a mean size in Scenario C, where the jumps are exponentially distributed, and an exact size in Scenario E, where the jump is deterministic. *Scenario C calibration:* industry projections of EV–ICE price parity imply that the pre-incentive price gap of roughly \$13,300 closes over about four years, an expected decline of \$13,300/4 $\approx$ \$3,325 per year. This is matched with a Poisson stream of exponential cost cuts at rate $\nu = 1$ and mean $|d_e| = \$3{,}325$ (so $\nu|d_e| = \$3{,}325/\text{yr}$). *Scenario E calibration:* a 50% chance that the \$7,500 federal credit is revoked over four years is spread evenly across the horizon, which gives an arrival rate of $\nu = \frac{1}{4} \cdot 50\% = \frac{1}{8}$ per year for a fixed upward jump of $d_e = \$7{,}500$. This is the conservative reading of the risk: a constant hazard matched exactly to the four-year probability, $\nu = \ln 2/4 \approx 0.17$, would lower the boundary further, to $d^*(C_0) = 4.15$, and so strengthen the case for immediate adoption. The row x^{**} or $d^*(C_0)$ reports the critical unit cost advantage above which immediate transition is optimal: the threshold x^{**} of Theorem 2.6 for the no-jump scenarios, and the closed-loop adoption boundary $d^*(C_0)$ evaluated at C_0 for the jump scenarios, from Theorem 2.11 in Scenario C and from Proposition B.9 in Scenario E.

cost of an electric vehicle in San Francisco, California, resulting in a break-even period of approximately eight years. Moreover, when incorporating the stochastic nature of unit cost advantages, the model suggests delaying adoption until these advantages stabilise at a higher and more favourable level.

Electricity in California is notably expensive due to several interrelated factors. A significant contributor is the state's ambitious environmental policy, which mandates substantial investments in renewable energy sources and infrastructure upgrades to support a greener grid. These initiatives, while environmentally beneficial, often come with high costs that are passed on to consumers.[18] In effect, though with intentions to help the economic transformation by greening the electric grid, high electricity prices delay the adoption of EVs according to the model. This highlights the importance of comprehensive policy planning. To ensure the feasibility

[18] `https://www.wsj.com/business/energy-oil/why-californians-have-some-of-the-highest-power-bills-in-the-u-s-a831b60e`

of widespread electrification despite high electricity costs, it is essential to redesign electricity rate structures to promote efficient energy use and balance grid demand, making the transition to electric technologies more economically viable (Borenstein et al., 2022).

In Seattle, the cost of electricity is comparatively low, primarily attributable to the region's extensive and cost-effective hydropower resources.[19] Consequently, in Scenario B the current unit cost advantage $D_0 = 7.94$ already exceeds the no-jump threshold $x^{**} = 6.08$, so it is economically rational to adopt immediately by investing in a fleet of EVs. Scenario C keeps the same market conditions but adds a forward-looking belief: the firm expects the EV premium to keep falling as battery costs decline and new models enter. Rather than a single dated event, this is modelled as an ongoing *stream* of cost-reducing breakthroughs: a Poisson process of exponential downward jumps in $C(t)$ at rate $\nu = 1$ per year with mean size \$3,325, an expected decline of \$3,325 per year (Table 5). The prospect of cheaper EVs tomorrow raises the value of waiting today, and, through Theorem 2.11, lifts the adoption boundary from $x^{**} = 6.08$ to $d^*(C_0) = 11.66$ (Figure 4).

Because $D_0 = 7.94$ now lies between the two boundaries, the optimal decisions diverge. In Scenario B the firm buys EVs immediately. In Scenario C it defers and buys new ICEs instead, waiting for the anticipated price declines. The deferral raises the model-implied discounted profit from \$302,573 (immediate adoption) to \$307,220. Two features of the optimal rule deserve emphasis. First, the deferral is not open-ended: the moment a cost cut is realised, the boundary on D drops by $(r/\lambda)|\Delta C|$ and adoption triggers as soon as the state crosses it: the firm waits *for* the breakthrough, not through it. Second, the precommitted benchmark understates the option value of waiting, placing its trigger at $x^{***} = 10.50$, below the closed-loop boundary, precisely because a rule that cannot react to the realised cut gains less from waiting for it.

The most compelling case is observed in Miami, where the optimal strategy hinges on the perceived security of the tax credit. Under the benchmark Scenario D, the current advantage $D_0 = 4.8$ lies below the no-jump threshold $x^{**} = 6.08$, so the firm waits, expecting the advantage to drift up before committing. Scenario E keeps the same conditions but introduces policy risk: the firm assigns a 50% chance that the \$7,500 federal credit is withdrawn within four years. This is captured as a constant hazard of an adverse policy shock, a jump that raises the adoption cost by \$7,500, arriving at rate $\nu = \frac{1}{8}$ per year, which makes the payoff process Z spectrally negative (only downward jumps). Because the risk now runs *against* waiting, it lowers the adoption boundary to $d^*(C_0) = 4.68$ (Figure 5), just below the current advantage $D_0 = 4.8$: the firm reverses its decision and adopts EVs at once, locking in the credit before it can lapse. Since the boundary sits essentially at D_0, adoption is triggered immediately, so whether the shock is treated as a one-off or as

[19] `https://www.mirskyelectric.com/electricity-cost-seattle`

a recurring hazard is immaterial for the decision. The proximity of the boundary to D_0 also shows how policy uncertainty of even moderate intensity is enough to flip the adoption decision.

Figure 6 collects the two jump scenarios in a single (C, D)-space comparison of the two solution concepts: the precommitted first-passage threshold x^{***} of Theorem 2.15 against the closed-loop boundary $d^*(C)$ of Theorem 2.11, with the constant-cost benchmark x^{**} as reference. The comparison makes the value of reacting to realised jumps visible, and shows that its magnitude depends on the direction of the anticipated shock. In Scenario C, where waiting is motivated by anticipated cost declines, the closed-loop rule demands a strictly higher hurdle ($d^*(C_0) = 11.66$ versus $x^{***} = 10.50$): a rule that can respond to a realised cost cut profits more from waiting for it, and therefore waits longer. In Scenario E the two rules nearly coincide ($d^*(C_0) = 4.68$ versus $x^{***} = 4.61$): the threat of a cost increase runs against waiting under either rule, and since the boundary already sits essentially at the current state, the ability to react adds little. Both rules agree on the direction of every response. They differ in magnitude precisely where reacting to realised jumps is valuable. The remainder of the analysis therefore reports the closed-loop values as the reference thresholds, with the precommitted rule retained as the explicit lower-bound benchmark.

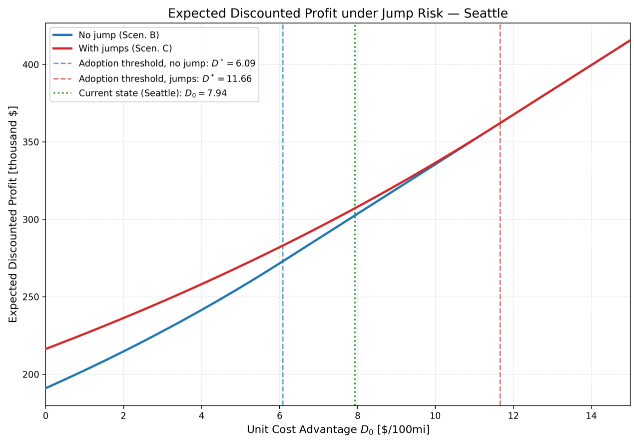

Figure 4: Expected Model-Implied Discounted Profit comparing Scenario B and Scenario C. It is shown as a function of the current unit cost advantage based on the chosen parameters outlined in Table 5 for Seattle. The solid grey line represents the value function of the baseline model without jumps (Scenario B). The solid blue line shows the value function of the closed-loop jump model of Theorem 2.11 with exponentially distributed cost reductions (Scenario C), computed via the roots β_1, β_2 of the exponent equation (2.39). Vertical lines depict the no-jump threshold x^{**} (grey dashed), the closed-loop adoption boundary $d^*(C_0)$ (blue dashed), and the current cost advantage level D_0 in Seattle (green dotted). Anticipated cost reductions raise both the value of waiting and the adoption boundary.

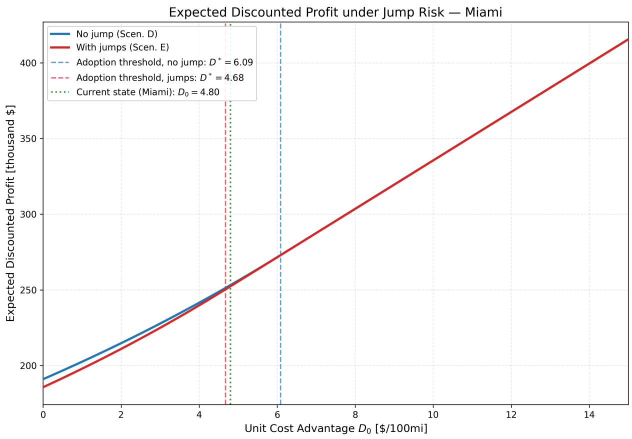

Figure 5: Expected Model-Implied Discounted Profit comparing Scenario D and Scenario E. It is shown as a function of the current unit cost advantage based on the chosen parameters outlined in Table 5 for Miami. The solid grey line represents the value function of the baseline model without jumps (Scenario D). The solid orange line shows the value function of the closed-loop jump model under the hazard of an adverse policy shock raising the adoption cost by $7,500 (Scenario E), which makes the payoff process Z jump only downward and so lowers the optimal level z^* (Proposition B.9). Vertical lines depict the no-jump threshold x^{**} (grey dashed), the closed-loop adoption boundary $d^*(C_0)$ (orange dashed), and the current cost advantage level D_0 in Miami (green dotted). The threat of subsidy withdrawal lowers both the value of waiting and the adoption boundary, which falls just below the current D_0: adoption is triggered immediately.

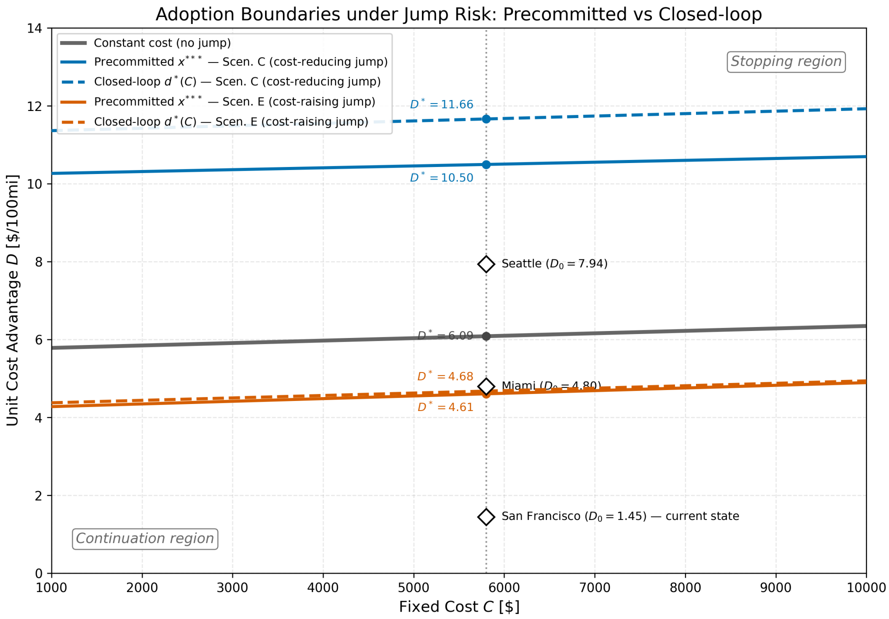

Figure 6: Adoption boundaries in (C, D) space under jump risk: precommitted first-passage thresholds x^{***} (solid) versus closed-loop boundaries $d^*(C)$ (dashed). The constant-cost benchmark x^{**} (solid grey) is shown for reference. Scenario C (cost-reducing jumps, blue) raises both boundaries above x^{**}, with the closed-loop rule demanding a higher hurdle ($d^*(C_0) = 11.66$ versus $x^{***} = 10.50$) because it can react to realised cost cuts. Scenario E (cost-increasing jumps, orange) lowers both boundaries below x^{**}, and the two rules nearly coincide ($d^*(C_0) = 4.68$ versus $x^{***} = 4.61$). Diamond markers locate the three cities at the current cost $C_0 = 5{,}800$: the current state in San Francisco ($D_0 = 1.45$) lies in the continuation region under all specifications, Miami ($D_0 = 4.80$) sits just above the Scenario E boundary, and Seattle ($D_0 = 7.94$) lies between the no-jump threshold and the Scenario C boundary.

2.4.2 Parameter Calibration and Its Impact on Decision-Making

As stated previously, although the current market conditions impact the expected time until the adoption of new technology, the optimal stopping level is determined by several factors: the forward-looking dynamics of the unit cost advantage variable process, including its drift (α) and volatility (σ), the user's annual miles driven (λ), and the fixed cost difference between the EV and the petrol car (C). A comprehensive analysis necessitates data that accurately reflects individual user characteristics, including vehicle fuel efficiency, annual mileage, local fuel and electricity prices, and the user's discount rate. Estimating the forward-looking drift and volatility of the unit cost advantage presents a significant challenge, as it requires projecting long-term market trends that are inherently uncertain and influenced by policy changes and economic fluctuations. This chapter combines historically calibrated represen-

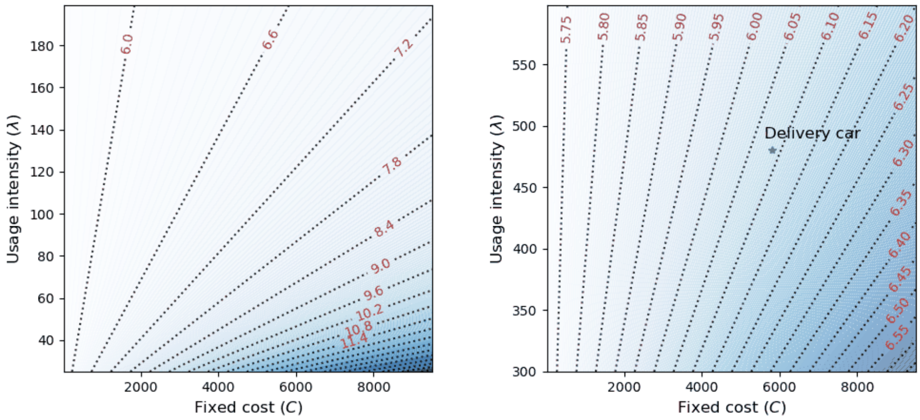

(a) (b)

Figure 7: Impact of fixed cost (C) and usage intensity (λ [x 100 miles per year]) on the Optimal Stopping Level (x^{**}) for EV Adoption. Each chart displays dotted lines representing these levels for $\alpha = 0.34$ and $\sigma = 2.42$. Two user profiles are considered with characteristic user mileage (λ). For (a), the usage intensity is characteristic of a low-usage passenger car. For (b), the usage intensity is representative of city or regional delivery cars. Note that the determined optimal stopping levels assume that there are no jumps, that is $\nu = 0$.

tative values with a range of parameter variations to account for the diversity in user profiles and geographic cost structures, thereby illustrating the model's practical implications. This subsection provides a detailed discussion of the calculation process and its impact on decision-making.

The real interest rate is set at 3% because its influence on the results is relatively minor compared to its variability. Figure 7 illustrates the impact and sensitivity of two key parameters describing individual characteristics: fixed costs and usage intensity. For an individual decision-maker with significantly lower annual mileage or a fleet of company-owned cars, the optimal stopping level exhibits high variability depending on the fixed cost. For an individual to opt for an EV in Seattle ($C_0 = 5800$, $D_0 = 7.94$), their annual mileage would need to exceed 8,000 miles. In contrast, for a high-mileage delivery company, the fixed cost has little effect on the optimal decision. Across the examined range (C between \$0 and \$10,000), the fixed cost–whether incurred directly or implied by an alternative vehicle choice–does not alter the optimal decision. Similarly, in terms of usage intensity, the optimal stopping level remains highly tolerant of variations in estimated annual mileage. Regardless of the specific combination of high mileage and fixed costs, the optimal stopping level consistently falls within the range of \$5.75 – \$6.55, ensuring stable outcomes even with uncertain input estimates.

An individual may consciously factor in externalities, such as the carbon dioxide emissions of a new car when making a decision. While monetising these effects can be challenging, a common approach is to assess their impact using a hypothetical carbon tax. To examine the effect of a hypothetical carbon tax on current costs, the variable cost advantage must be increased by the amount of potential carbon taxes. For example, in the US, the current difference in emissions between EVs and petrol cars is relatively small on average, given the current production energy mix. This translates to a 4.9 tonnes/year difference between an average EV and a petrol car, based on an average of 11,579 miles driven.[20] To provide some context, the REMIND-MAgPIE 3.0-4.4 model predicts carbon prices for 2030 to be between \$6 and \$261 per tonne of CO2 for the Current Policies and Divergent Net-Zero scenarios, respectively (in 2010 USD).[21] By approximating the carbon tax cost impact, this equates to a potential additional unit cost advantage ranging from \$0.25 to \$11 (per 100 miles) depending on the scenario. This extra consideration in our case study would only make the EV the optimal decision in that specific Divergent Net-Zero scenario, when the carbon tax is extremely high. Since carbon taxes are currently hypothetical, they are not included in the numerical example. However, it is important to consider that, depending on the energy mix and the progression of carbon taxes, they may become a significant factor in the future. Carbon taxes could potentially increase the cost advantage of EVs when policies are implemented or contribute to the expected drift of the unit cost variable process.

Parameters α (drift) and σ (volatility) were calibrated using historical data from the United States reflecting real-world conditions. The drift and volatility estimates were calculated from a synthetic history based on unit cost advantage. To calculate the unit cost advantage, the cost of one unit of fuel (gallon for gasoline and kWh for electricity) should be converted into the distance that the purchased unit allows us to travel. This is referred to as the fuel efficiency of a technology or a specific car. The fuel efficiency of cars can vary, but since in this analysis it is flat, it doesn't contribute to either drift or volatility, so we used the average fuel consumption for petrol and EVs to be 8.6 [l/100 km] and 1.9 [l/100 km], respectively, as reported by the IEA.[22] The report calculates the fuel consumption in litres of gasoline equivalents per 100 km, but using standard conversions, it can be transformed to "kWh/miles" for EVs and "gallons/miles" for petrol cars.[23] It is noteworthy that after a period of improvements in average fuel efficiency, fuel efficiency seems to have plateaued in the last five years of the report (2014–2019) and is currently stagnating. Based on this observation, it is reasonable to assume that this factor remains relatively stable.

[20] `https://afdc.energy.gov/vehicles/electric_emissions.html`

[21] NGFS Scenario Explorer `https://data.ece.iiasa.ac.at/ngfs`

[22] IEA (2021), Global Fuel Economy Initiative 2021 Data Explorer, IEA, Paris `https://www.iea.org/data-and-statistics/data-tools/global-fuel-economy-initiative-2021-data-explorer`

[23] Beyond standard unit conversions between litre to gallon and km to mile, the conversion between miles per gallon equivalent to "kWh" used is based on `https://energyeducation.ca/encyclopedia/Miles_per_gallon_gasoline_equivalent`.

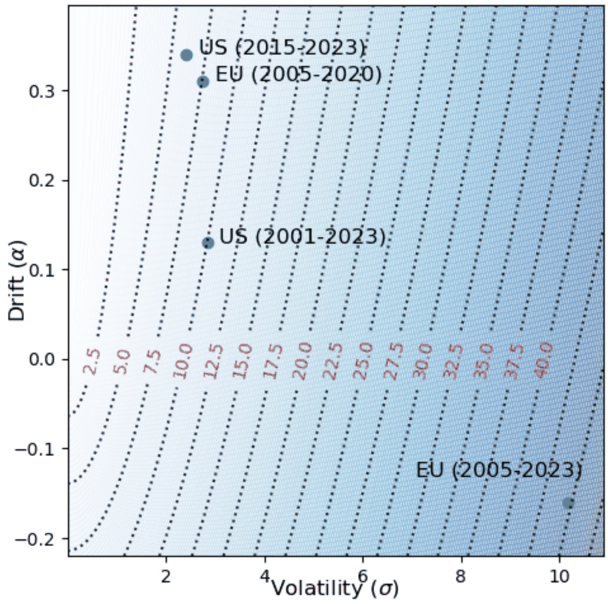

Figure 8: Impact of Drift (α) and Volatility (σ) Parameters on the Optimal Stopping Level (x^{**}) for EV Adoption. The fixed cost (C) is \$5,800, while the usage intensity (λ) is 48,000 miles annually. Note that the determined optimal stopping levels assume that there are no jumps, that is $\nu = 0$.

Once the synthetic unit cost advantage history was established, drift and volatility estimates were computed over different time periods. For our case study, the most optimistic outlook is used from the U.S. market (2015–2023). This selection is critical, as the optimal stopping level is highly sensitive to drift and volatility, underscoring the importance of expert judgment. Figure 8 illustrates how variations in drift and volatility influence the optimal stopping level, given the fixed cost and usage intensity parameters from our case study.

The dotted lines on the chart represent the optimal switching levels for different drift and volatility combinations. The four dots represent historical drift and volatility combinations calculated from the four time series presented in Figure 2. The two points for the US (2015–2023) and EU (2005–2020) seem to be close, but this is deemed a coincidence rather than proof of any long-term stable market condition. The scale of the α and σ parameters can introduce a broad range of optimal stopping levels, potentially spanning several orders of magnitude. This variability could lead to outcomes ranging from immediate adoption to never transitioning to the new technology. These findings highlight that the proposed model is highly sensitive to changes in both perceived drift and volatility, with volatility exerting a significantly greater influence on the optimal stopping level. This suggests that uncertainty in cost savings can markedly delay technology adoption decisions, emphasising the importance of reducing volatility to promote earlier adoption.

For negative drift, there might still be scenarios when it is profitable to switch to the new technology, such as when the current variable benefit is high enough or when D breaches the optimal stopping level. However, when comparing the current prices to the optimal stopping levels, all markets have a D_0 lower than the optimal level. In addition, with a negative drift, the expected time for any level higher than the current one to be breached is infinite. These are important takeaways for policymakers who can influence market stability and would like to support the green transition. If users optimise their timing for expected discounted profit and expect high volatility, they will delay their decision to switch to the new technology compared to when they anticipate stable prices. It is similar to the expected drift, although the impact is somewhat muted. Price uncertainty is detrimental in both directions. If users expect market conditions to be worse than what actually occurs, they may choose not to adopt the new technology and, therefore, miss out on potential benefits. However, if they are overly optimistic and switch to the new technology, their profits could be worse than expected, or they may even incur losses with the switch. Both options hurt the utility of people and the profitability of companies. In addition, when there is significant uncertainty in the energy markets, personal views will influence decision-making, resulting in dispersed behaviour and a potentially divergent transition, which is usually more costly than an orderly transition.

2.5 Discussion and policy implications

This section moves beyond the analytical and empirical results to critically explore their implications for EV adoption within uncertain economic environments. By examining the relationship between drift, volatility, and adoption timing, conventional assumptions regarding the effect of uncertainty on technology transition are questioned. Additionally, the findings are contextualised within the broader literature, comparing them to recent studies in technology adoption and market stability. This critical approach allows for new research directions that extend the understanding of clean technology adoption.

2.5.1 Interpreting analytical results

The outcomes of Theorems 2.6–2.7, which are related to the benchmark problem without jumps, are first interpreted in this section. Recall that the continuation and stopping regions are separated by a point

$$x^{**} = \frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\alpha}{r} + \frac{Cr}{\lambda}. \tag{2.61}$$

By investigating Equation (2.61), a few interesting findings are observed. Even though a risk-neutral decision maker was assumed, x^{**} is dependent on σ, implying that the level of volatility still influences the optimal decision-making strategy. This observation underscores the importance of incorporating stochastic factors into the unit cost advantage process (D). There is a trade-off between switching to the new technology and benefiting from potential future unit cost advantages earlier versus paying off the fixed costs at a higher present value. The net impact, as seen in Equation (2.61), is that the optimal stopping level is a positive function of σ. On the other hand, Equation (B.5) in the Appendix reveals that the expected waiting time to first hit any fixed level is not affected by the underlying volatility. However, when the optimal stopping level and the expected waiting time are combined, Equation (2.27) is the result. This equation shows that when the α value is positive and $D_0 < x^{**}$, the expected waiting time to reach the optimal level ($\mathbb{E}_{D_0}(D_{x^{**}})$) also increases as σ increases.

Note that the optimal stopping level does not depend on the current unit cost (D_0). Besides, barring degenerate cases, the expected waiting time to reach any fixed level for process D only depends on the distance to the optimal stopping level and the drift level. The limit $\lim_{\sigma\to 0} x^{**} = \frac{Cr}{\lambda}$ (for $\alpha > 0$) represents the solution to the deterministic optimisation problem, indicating that the optimal level of adoption and the expected time to reach that level ($\lim_{\sigma\to 0} \mathbb{E}_{D_0}(D_{x_{max}}) = \frac{Cr}{\lambda\alpha} - \frac{D_0}{\alpha}$, barring degenerate cases) are minimised when volatility is 0.

As an interesting result, the optimal stopping level does not depend solely on C or λ but rather on their ratio. This notion can be supported intuitively, as the benefits of increased usage must outweigh the fixed costs associated with implementing the technological change. If the user drives twice as much, they can afford to spend twice as much on the additional fixed cost of the transition, as the increased variable benefit would offset the higher cost. Another result is that the optimal stopping level is a decreasing function of α, suggesting that the expected long-term profit gained from a steeper unit cost variable process compensates for the loss of discounting effect associated with the fixed cost paid earlier. Finally, x^{**} is not a monotonic function of the interest rate r. The exact relationship between the interest rate and the time of the new technology adoption depends on the other parameters that affect the transition.

Jump risk modifies this reading in one direction: anticipated technological breakthroughs, represented by ν and $d_e < 0$, shift the adoption boundary upward from x^{**} to $d^*(C_0)$, so a state that would justify adoption in the constant-cost model no longer does, while a realised cost cut can still trigger adoption on the spot. The mechanism and its policy content are set out in Subsection 2.3.3.

2.5.2 Policy Implications and Societal Benefits

The results show that the drift and volatility expectations of the unit cost advantage play a pivotal role in determining the timing of the transition to electric vehicles (EVs). Specifically, higher drift values correlate with a greater likelihood of earlier adoption, while increased volatility tends to delay the adoption. These results align with previous research on technology adoption and optimal stopping theory, including works such as Dixit and Pindyck (1994) and Trigeorgis and Tsekrekos (2018), which suggest that uncertainties in future benefits can lead to a delay in irreversible investments. This chapter extends these conclusions by providing explicit quantifications of the stopping thresholds under stochastic energy price dynamics, contributing to real options analysis in the context of clean technology adoption.

Furthermore, the results indicate that more intensive users are more inclined to adopt new technology early due to the greater cost savings. This observation is consistent with the optimal-switching analysis of Falbo et al. (2021) and the fleet replacement programme of Kleindorfer et al. (2012), in both of which higher utilisation brings forward the switching date. Additionally, the analysis underscores the importance of government intervention to stabilise energy prices, as high volatility can deter optimal adoption decisions. These insights are especially relevant for policymakers aiming to promote clean technology through targeted incentives, as highlighted by Mercure et al. (2018) and Lam and Mercure (2021).

To support adaptation, policymakers should address high market volatility by implementing targeted energy stabilisation measures that account for costs and potential impacts. Strategic reserves, such as the U.S. Strategic Petroleum Reserve, offer a buffer against supply shocks, although they require significant investment and upkeep. Subsidies and price caps keep energy affordable for consumers and reduce production costs for providers but can distort market signals and strain budgets. Long-term supply contracts provide price stability over time but may limit responsiveness to changing market conditions. In extreme cases, market interventions such as capacity markets or temporary control over resources can be employed to maintain stability, albeit with regulatory and operational challenges. Balancing these strategies with cost-effectiveness and sustainability in mind allows policymakers to foster a resilient energy environment.

An essential focus area is the governments' Net-Zero goals and the road to get there. Transparency and smooth, controlled transition strategies matter, as anticipated or sudden disruptions in the energy market (e.g., banning fossil energy sources, without a clear alternative) can impede the voluntary adoption of new green technologies. Conversely, a gradual and predictable increase in carbon taxes can accelerate the improvement of unit cost advantages, thereby further incentivising early adoption. Moreover, the revenue generated from the tax can be allocated to fund other policy initiatives, such as market stabilisation measures.

The model emphasises that effective EV adoption policies should integrate financial incentives, price stability and public awareness. Policymakers can leverage this model to evaluate the effects of different subsidies, compare real-world adoption rates to model projections, and determine cost-efficient pathways for EV uptake. For instance, if adoption rates lag due to high upfront costs, subsidies or tax credits could offset initial expenses, encouraging early adopters. On the other hand, if actual price volatility deters adoption, policymakers could implement market-stabilisation measures such as strategic reserves or transparent trading practices. Where perceived volatility is more problematic, interventions like clear communication on price stability or price guarantees for EV electricity can boost consumer confidence. Data-driven public awareness campaigns can further inform consumers about the economic and environmental benefits of EVs, correct misconceptions, and foster a supportive adoption climate.

2.6 Chapter Summary and Extensions

This chapter examines an optimisation problem for adopting EVs in a stochastic framework, building on a body of research that underscores the urgent need for a transition from conventional fossil-fuel-based vehicles to environmentally friendly EVs. In alignment with global efforts, such as the Paris Agreement, which aims to mitigate greenhouse gas emissions, society has been increasingly inclined towards sustainable development and energy-efficient clean technology. This chapter addresses the well-motivated and timely question of when to switch to EVs, considering the stochastic nature of future price movements. The analysis addresses an endogenous, cost-based optimisation problem embedded in a stochastic framework, considering factors that drive or delay the adoption of clean technology. The analysis, grounded in empirical data, focuses on an individual or entity contemplating the replacement of gasoline vehicles with EVs, with specific attention to usage profiles, fixed and unit costs, and the stochastic nature of cost advantages associated with EV usage. Analytical solutions are derived for expected discounted profits and optimal switching regions, uncovering that the expectations of drift and volatility in unit cost advantages play a central role in determining the timing of the transition.

The results show that optimal adoption timing is primarily determined by the drift and volatility of the unit cost advantage, usage intensity, and the fixed adoption cost, as discussed in detail in Section 2.5. Under jump risk the same problem reduces exactly to one-dimensional optimal stopping of a Lévy process, which yields the optimal closed-loop boundary $d^*(C)$ in closed form (Theorem 2.11): a rule that moves with the prevailing cost, so a realised breakthrough or subsidy withdrawal can trigger adoption at the instant it occurs. High-intensity users and those anticipating rapid cost advantages are likely to adopt sooner, while high uncertainty in energy prices can significantly delay adoption decisions. This chapter extends the conclusions of Dixit and Pindyck (1994) and Trigeorgis and Tsekrekos (2018) by providing explicit quantifications of stopping thresholds under stochastic energy price dynamics, supporting the policy arguments of Mercure et al. (2018) and Lam and Mercure (2021).

To accelerate EV adoption and mitigate transportation's environmental impact, the results highlight the need for policies combining financial incentives, energy price stabilisation, infrastructure investment, and public awareness. Policymakers can use the model to assess the impact of subsidies, compare adoption rates to model predictions and find cost-effective strategies. Targeted incentives for high-mileage users, investment in charging infrastructure, and support for battery innovation can lower barriers like upfront costs and range anxiety. To manage price volatility, stabilisation measures (e.g., strategic reserves) or clear communication on pricing can reduce uncertainty for consumers. Public awareness campaigns are also essential to educate on the benefits of EVs, dispel misconceptions, and build confidence in clean technology. These policies not only support a faster transition to sustainable technology but also align with the United Nations Sustainable Development Goals, particularly SDG 7 (Affordable and Clean Energy) and SDG 13 (Climate Action).

As an empirical analysis, five distinct scenarios were examined involving three different locations, with two scenarios explicitly accounting for potential jumps in the fixed cost process. The case study shows how adoption decisions depend on location. In California, despite incentives, high electricity prices discourage immediate adoption due to insufficient annual fuel savings to offset EVs' upfront costs. In Seattle, EV adoption timing depends on anticipated price developments: immediate transition is optimal if significant price drops are unlikely, but delayed adoption is optimal if EV-ICE price parity is expected soon. Lastly, in Miami, the decision heavily depends on policy continuity. Maintaining the tax credits then in force favours delaying adoption, whereas the risk of their removal justifies immediate transition. These findings underscore the importance of localised market conditions and policy frameworks in shaping optimal technology adoption strategies.

Although the focus of the case study in this chapter is on the decision to switch to EVs, the modelling framework can be applied to the general problem of adopting a new technology that requires an initial investment but offers long-term benefits. Some examples include replacing old heating, ventilation, and air conditioning (HVAC) systems with modern alternatives like heat pumps, installing solar energy panels to supplement a building's electricity supply, or simply investing in more efficient and environmentally friendly solutions or appliances (e.g. refrigerators, dryers, and lighting). The model can also accommodate utility and personal aspects for most parameters. Because of their overall usability, the results are relevant for a broader audience of theoretical and empirical researchers. Policymakers who are studying efficient ways of incentivising individuals and companies to adopt new technologies as part of an effort to reduce carbon emissions might consider utilising these results. They should focus on implementing policies that promote market stability and growth.

Even under jump risk the adoption cost of this chapter moves only in discrete steps, so the absence of continuous cost uncertainty is the principal modelling constraint. Chapter 3 relaxes it by specifying the adoption cost $C(t)$ as an Ornstein-Uhlenbeck process correlated with $D(t)$. Unlike the jump case, that specification admits no reduction to a single state variable, and therefore requires the finite-difference methods developed there. The infinite planning horizon is likewise relaxed in Chapter 3 through a finite-horizon formulation, and Chapter 4 synthesises the single- and two-factor results through a cross-scenario comparison that quantifies the additional option value generated by cost stochasticity.

Despite the contributions of this chapter, several limitations and challenges remain for practical implementation. Key constraints arise from the modelling assumptions, particularly the assumption that the unit cost advantage ($D(t)$) follows an arithmetic Brownian motion. This approach may not fully capture real-world cost dynamics, as historical data shows that crises, demand and supply shocks, advancements in services and processes, or policy changes can lead to sudden shifts or regime switches in commodity prices. Incorporating regime switching or time-varying volatility could enhance the model's realism. The challenge is the inherent uncertainty in estimating forward-looking parameters, particularly drift (α) and volatility (σ), which are sensitive to the same unpredictable factors that might cause sudden or persistent market dislocations. Although the model's framework allows for incorporating subjective factors in the evaluation, it might not be straightforward to monetise certain aspects, like range anxiety. Further limitations relating to behavioural factors, charging infrastructure, and modal substitution are discussed in Chapter 5.

Chapter 3

EXTENDED MODELS: NUMERICAL SOLUTIONS

Chapter 2 holds the adoption cost either fixed or subject to discrete jumps, abstracting from the empirical reality that technology prices decline through production learning while fluctuating around long-run trends. Battery pack prices illustrate both forces: costs have fallen nearly 90% since 2010 through learning rates that the literature places between 6% and 27% per doubling of cumulative output (Nykvist and Nilsson, 2015; Ziegler and Trancik, 2021; BloombergNEF, 2024), while exhibiting volatility from supply-chain disruptions and input material price shocks.

This chapter extends the model to allow stochastic adoption costs. The analysis proceeds in five stages. Section 3.1 derives the governing equations for the two-factor model. Section 3.2 examines a tractable special case with deterministic cost trajectories. Section 3.3 extends the framework to finite planning horizons. Section 3.4 presents the finite-difference solution methodology, and Section 3.5 establishes numerical solution validity.

3.1 Two-Factor Model with Stochastic Cost and Benefits

3.1.1 Formulation of the Two-Factor Model

The model builds on the constant-cost framework of Chapter 2, detailed in Section 2.1. The core parameters remain unchanged: unit cost advantage D (benefit per unit of usage), usage intensity λ (annual usage level), discount rate r, benefit drift α, and benefit volatility σ. The adoption cost C, previously fixed, now follows a stochastic process governed by four additional parameters.

Mean-reversion rate of the adoption cost (κ). The parameter $\kappa > 0$ determines the speed at which deviations of the adoption cost from its long-run level dissipate. Economically, κ reflects the rate of learning or adjustment in production costs: a higher κ implies faster convergence of C_t toward its equilibrium value as experience accumulates and technology matures.

Long-run adoption cost level (θ). The parameter θ represents the asymptotic mean of the adoption cost process. It captures the structural or persistent component of cost that remains after learning-by-doing and scale economies have been exhausted, corresponding to the stable long-run equilibrium level of the technology's relative cost. In most clean-technology markets, current adoption costs remain above this expected long-run level: empirical evidence on electric vehicles indicates that further cost reductions are anticipated as battery prices continue to decline and manufacturing scales up.[1] Long-run price parity with incumbent technologies is therefore consistent with θ being below the present cost level and potentially converging toward zero. The model, however, remains flexible: a nonzero θ can represent persistent differences in total ownership cost, policy-driven subsidies, or technology-specific installation expenses.

Adoption Cost volatility (η). The parameter $\eta > 0$ measures the amplitude of short-term fluctuations in adoption cost around its mean. These fluctuations may originate from input-price volatility, policy adjustments, or temporary supply-chain disturbances that do not alter the long-run trend.

Correlation between short term fixed cost and unit cost advantage shocks (ρ). The parameter $\rho \in (-1, 1)$ represents the instantaneous correlation between the short-term stochastic shocks affecting the unit cost advantage process D_t and those influencing the fixed adoption cost C_t. It captures how contemporaneous market disturbances jointly affect the relative profitability of the clean technology and its cost of adoption.

A *negative correlation* ($\rho < 0$) indicates that shocks increasing the relative benefit of adoption tend to coincide with temporary cost reductions. This pattern may occur when a sudden rise in fossil-fuel prices increases the usage-based advantage D_t while simultaneously reducing the relative adoption cost C_t (for instance, if higher fuel prices accelerate EV demand and production scale, lowering costs). Conversely, when fossil-fuel prices fall, the reduced usage-based advantage D_t coincides with increased relative adoption costs C_t.

A *positive correlation* ($\rho > 0$) would imply that shocks improving the benefit of the clean technology also increase its relative cost. This could occur, for instance, if surging commodity prices raise fossil fuel costs (increasing D_t) while also raising

[1](Goetzel and Hasanuzzaman, 2022; International Energy Agency, 2024)

battery material costs (increasing C_t as EVs become relatively more expensive to purchase). Such co-movement characterised periods like 2021-2022 when both energy and battery material prices spiked together.

Because empirical evidence on instantaneous correlations between fuel price shocks and vehicle cost shocks is limited, the sign of ρ is guided by economic reasoning rather than direct estimation. A small or moderately negative correlation is adopted as the baseline specification, though the 2021-2022 period demonstrates that positive correlation can emerge during broad commodity market stress.

Remark 3.1 (Time-varying correlation). The model treats ρ as constant. Accommodating time-varying $\rho(t)$ would require either a regime-switching specification with Markov-modulated correlation or a stochastic correlation model in which ρ itself follows a bounded diffusion. Both extensions increase the state dimension by at least one, replacing the current two-dimensional PDE with a system of coupled PDEs (regime switching) or a three-dimensional problem (stochastic ρ). The additional computational cost is substantial, and calibrating transition rates or correlation dynamics from the limited joint data currently available (Section 4.1) is not feasible. The sensitivity analysis in Chapter 4 examines $\rho \in \{-0.3, 0, +0.3\}$ as a partial substitute, bracketing the range of plausible correlation regimes.

The decision-maker chooses between immediate adoption (paying the current cost C_t to lock in the perpetual benefit stream λD_s for $s \geq t$) and waiting (preserving flexibility to time the switch optimally as uncertainty resolves). The dynamics of the two-dimensional system are given by

$$dD_t = \alpha\, dt + \sigma\, dW_t^D, \qquad \text{(usage-based benefit, ABM)} \tag{3.1}$$

$$dC_t = \kappa(\theta - C_t)\, dt + \eta\, dW_t^C, \qquad \text{(adoption cost, OU)} \tag{3.2}$$

where the Brownian motions satisfy $d\langle W^D, W^C\rangle_t = \rho\, dt$. Setting $\eta = 0$ and $\rho = 0$ yields a deterministic cost path (covered in Section 3.2), while fixing $C_t \equiv C$ recovers the constant-cost model of Chapter 2.

The Ornstein–Uhlenbeck drift $\kappa(\theta - C_t)$ pulls the cost process toward the long-run level θ with half-life $\ln(2)/\kappa$. The explicit solution

$$C_t = \theta + (C_0 - \theta)e^{-\kappa t} + \eta \int_0^t e^{-\kappa(t-s)}\, dW_s^C$$

is a Gauss–Markov process with stationary distribution $N(\theta, \eta^2/(2\kappa))$ as $t \to \infty$ (Uhlenbeck and Ornstein, 1930). This captures mean reversion in technology prices as costs converge toward the long-run level θ through learning-by-doing and scale economies.

The closest structural antecedent to the system (3.1)–(3.2) is McDonald and Siegel (1986), in which both the value of the project and the cost of acquiring it evolve stochastically and the correlation between their innovations enters the optimal investment threshold directly. The present model retains that two-factor architecture but departs from it in the specification of each factor. The cost process is mean-reverting rather than geometric, so the strike is anchored to a structural long-run level θ determined by learning-by-doing and scale economies rather than drifting without an equilibrium; the benefit process is arithmetic rather than geometric, permitting negative realisations in periods when the incumbent technology is the cheaper one to operate.[2]

The constant-cost model of Chapter 2 admitted closed-form solutions through direct application of stochastic calculus: Laplace transforms of first-passage times and smooth-pasting conditions yielded explicit threshold formulas. A second state variable does not by itself forfeit this. The jump-cost extension of Section 2.3 also carries two state variables, yet the linear adoption payoff collapses them into the single process $Z = \frac{\lambda}{r}D + \frac{\lambda\alpha}{r^2} - C$, which inherits stationary independent increments from D and the compound Poisson cost, so the problem is solved exactly in one dimension. The two-factor system (3.1)–(3.2) precludes such treatment: with a mean-reverting cost the drift of Z becomes $\frac{\lambda\alpha}{r} - \kappa(\theta - C_t)$, which depends on the state and cannot be recovered from Z alone, so Z is no longer a process of independent increments and the reduction fails. The mean-reverting dynamics destroy the scale-invariance on which those analytical techniques depend. Solving the optimal stopping problem therefore requires recasting it as a partial differential equation via dynamic programming.[3]

The core principle is the *Bellman decomposition*: at each instant, the decision-maker compares immediate exercise against the discounted continuation value. For the two-factor system, this takes the form

$$V(c,d) = \max\Big\{\hat{g}(c,d),\ \mathbb{E}\big[e^{-r\Delta t}\,V(C_{t+\Delta t}, D_{t+\Delta t})\,\big|\,C_t = c,\ D_t = d\big]\Big\}, \qquad (3.3)$$

where the immediate-adoption payoff

$$\hat{g}(c,d) = \frac{\lambda d}{r} + \frac{\lambda\alpha}{r^2} - c \qquad (3.4)$$

is carried over unchanged from Chapter 2 (Corollary 2.5), now evaluated at the prevailing cost c rather than at a cost fixed at time zero. Taking the limit $\Delta t \to 0$ and applying Itô's formula yields the Hamilton–Jacobi–Bellman variational inequality derived in Subsection 3.1.2.

If adoption occurs at a stopping time τ, the firm pays C_τ and receives the perpetual stream of discounted benefits λD_t. The value function is

$$V(c,d) = \sup_{\tau \ge 0}\ \mathbb{E}_{c,d}\left[-e^{-r\tau}C_\tau + \int_\tau^\infty e^{-rs}\lambda D_s\,ds\right], \qquad (3.5)$$

[2] Correlated mean-reverting and driftful factors also appear in the two-factor commodity price literature, notably Schwartz and Smith (2000), though there the two components are latent and sum to a single log spot price rather than entering the payoff as separate state variables.

[3] The formal definitions of admissible stopping times, the general optimal stopping problem, and the Dynamic Programming Principle are collected in Appendix B.6.

where $(C_0, D_0) = (c, d)$. In the language of general optimal stopping theory, V is the smallest r-excessive (superharmonic) majorant of the immediate-adoption payoff $\hat{g}$, the continuous-time Snell envelope of the stopping problem (El Karoui, 1981; Peskir and Shiryaev, 2006). This characterisation is mathematically equivalent to the variational inequality derived in Subsection 3.1.2, but it makes the conceptual structure explicit: the value function is the cheapest supermartingale dominating the payoff, and stopping is optimal exactly where the two coincide.

The optimal policy is characterised by a stopping region $\mathcal{S} \subset \mathbb{R}^2$ such that $\tau^\star = \inf\{t \geq 0 : (C_t, D_t) \in \mathcal{S}\}$.[4] Note that this stopping time implies simultaneous observation of both state variables. Operationally, the decision-maker adopts when (C_t, D_t) enters the stopping region. While the precise geometry of $\mathcal{S}$ cannot be established analytically, the structure of the immediate-adoption payoff $\hat{g}(c, d) = \frac{\lambda d}{r} + \frac{\lambda \alpha}{r^2} - c$ (derived as Π_∞ in equation (2.6) of Chapter 2; the expression is unchanged by stochastic cost dynamics because C enters as the realised adoption cost at the moment of switching, not as a future trajectory), which is increasing in d and decreasing in c, suggests the boundary can be represented as $\mathcal{S} = \{(c, d) : d \geq b(c)\}$ with $b'(c) \geq 0$: higher current adoption costs require higher contemporaneous benefits to make immediate adoption optimal. This monotone-boundary conjecture extends the threshold structure from Chapter 2 to the two-dimensional setting.

3.1.2 Hamilton–Jacobi–Bellman and Variational Inequality

Taking the limit $\Delta t \to 0$ in the Bellman equation (3.3) and applying the generalised Itô formula for the joint process (C_t, D_t), the value function is characterised as the solution to a Hamilton–Jacobi–Bellman (HJB) variational inequality. The formal verification theorem guaranteeing that smooth solutions of the HJB equation coincide with the value function is stated in Appendix B.6 (Theorem B.4).

For the joint process (C_t, D_t) with Ornstein–Uhlenbeck adoption cost dynamics (3.2) and arithmetic Brownian motion unit cost advantage dynamics (3.1), the infinitesimal generator is

$$\mathcal{L}V = \kappa(\theta - c)\, V_c + \alpha\, V_d + \tfrac{1}{2}\eta^2 V_{cc} + \tfrac{1}{2}\sigma^2 V_{dd} + \rho\, \eta\sigma\, V_{cd}\,, \tag{3.6}$$

where subscripts denote partial derivatives and the model's time-homogeneity eliminates explicit time dependence. Throughout this chapter $\mathcal{L}$ denotes this two-dimensional generator of the joint process (C, D); the one-dimensional operator $\mathcal{L}_D$ of Chapter 2, which acts on D alone, is recovered by setting $\eta = \rho = 0$ and suppressing the c-derivatives.

[4]The existence of the optimal stopping region $\mathcal{S}$ is guaranteed by the theory of variational inequalities. Specifically, the value function $V(c, d)$ exists as the unique viscosity solution to the HJB equation (Oksendal, 2013). The stopping region is formally defined as the coincidence set $\mathcal{S} = \{(c, d) : V(c, d) = \hat{g}(c, d)\}$. Furthermore, due to the monotonicity of the payoff function $\hat{g}$ with respect to D, the boundary $\partial\mathcal{S}$ constitutes a free boundary curve $D^*(C)$ dividing the state space into connected continuation and stopping regions (Peskir and Shiryaev, 2006).

Following the dynamic programming principle, the value function satisfies the stationary HJB equation in the continuation region $\mathcal{C}$:

$$(\mathcal{L} - r)V(c, d) = 0, \qquad (c, d) \in \mathcal{C}\,, \tag{3.7}$$

with boundary condition in the stopping region $\mathcal{S}$:

$$V(c, d) = \hat{g}(c, d) = \frac{\lambda d}{r} + \frac{\lambda \alpha}{r^2} - c, \qquad (c, d) \in \mathcal{S}\,. \tag{3.8}$$

Note that the continuation equation is homogeneous: benefits accrue only after adoption, so no running payoff λd enters during the waiting phase; the benefit stream is captured entirely by the obstacle $\hat{g}(c, d)$. At the optimal boundary $\partial\mathcal{S}$, value matching ($V = \hat{g}$) and smooth-fit conditions ($\nabla V = \nabla \hat{g}$) are expected to hold under standard regularity assumptions (see Øksendal and Sulem, 2007, Ch. 4).

Combining (3.7) and (3.8) yields the variational inequality

$$\max\left\{ (\mathcal{L} - r)V(c, d),\;\; \hat{g}(c, d) - V(c, d) \right\} = 0, \qquad (c, d) \in \mathbb{R}^2\,. \tag{3.9}$$

Under the principle of smooth fit, the following conditions hold at the free boundary $\partial\mathcal{S}$.

Theorem 3.2 (Free-Boundary Optimality Conditions)**.** *For all $(c, d) \in \partial\mathcal{S}$:*

1. ***Value-matching:*** $V(c, d) = \hat{g}(c, d)$, *ensuring V is continuous across $\partial\mathcal{S}$.*
2. ***Smooth-pasting:*** $\nabla V(c, d) = \nabla \hat{g}(c, d)$, *ensuring the value function is tangent to the payoff manifold.*

Value matching is the continuity of V across $\partial\mathcal{S}$, which follows from $V \geq \hat{g}$ everywhere together with $V = \hat{g}$ on $\mathcal{S}$. Smooth pasting is the substantive condition, and it requires the boundary to be regular for the stopping region, which holds here because both components of (C_t, D_t) are non-degenerate diffusions. The general statement and its proof are given by Peskir and Shiryaev (2006); see also Pham (2009) for the multidimensional case.

Remark 3.3 (Geometric Intuition). Smooth-pasting implies that the marginal value of waiting vanishes at the boundary. For regular diffusions, the infinite variation of the process would allow a gain if a "kink" existed at $\partial\mathcal{S}$, contradicting optimality of τ^*.

Although the immediate-adoption payoff $\hat{g}(c, d)$ is smooth, the value function $V(c, d)$ typically attains only C^1 regularity across the free boundary $\partial\mathcal{S}$: the PDE operator governs V in $\mathcal{C}$ while the constraint $V = \hat{g}$ binds in $\mathcal{S}$, and these two conditions generically produce different second-order behaviour at the junction. Because V fails to be C^2 everywhere, it cannot satisfy the HJB equation in the classical sense. We therefore interpret (3.9) in the *viscosity sense* (Pham, 2009).

Definition 3.4 (Viscosity Solution)**.** A continuous function V is a viscosity solution to (3.9) if:

1. For any smooth test function ϕ touching V from above at (c, d): $\max\{(\mathcal{L} - r)\phi,\ \hat{g} - \phi\} \leq 0$.

2. For any smooth test function ϕ touching V from below at (c, d): $\max\{(\mathcal{L} - r)\phi,\ \hat{g} - \phi\} \geq 0$.

For the joint diffusion (3.1)–(3.2), the viscosity solution is unique and satisfies the *comparison principle*, providing a stable mathematical target for the numerical schemes developed in Section 3.4.

In the constant-cost case of Chapter 2, this variational inequality reduced to

$$\max\left\{\alpha V_d + \tfrac{1}{2}\sigma^2 V_{dd} - rV,\ \left(\frac{\lambda d}{r} + \frac{\lambda\alpha}{r^2} - C\right) - V\right\} = 0,$$

featuring a single threshold $d = b$. The full two-factor system (3.9) admits no such simple characterisation. The additional second-order terms in the generator (cost volatility $\frac{1}{2}\eta^2 V_{cc}$ and the cross-term $\rho\eta\sigma V_{cd}$) introduce convexity and correlation into the value function, transforming the optimal boundary from a simple threshold into a curved free boundary $d^*(c)$ that depends nonlinearly on the cost level.

To isolate the impact of these stochastic terms, we first examine the intermediate case where cost volatility vanishes ($\eta = 0$). Section 3.2 derives analytical solutions for this case, establishing a linear benchmark against which the curvature of the fully stochastic solution can be measured.

3.2 Deterministic Cost Trajectories

Setting $\eta = 0$ eliminates cost volatility while preserving mean reversion through the drift $\kappa(\theta - C_t)$. The resulting deterministic path $C(t) = C_0 e^{-\kappa t} + \theta(1 - e^{-\kappa t})$ captures expected cost decline toward long-run parity θ without transitory shocks. The problem reduces from a two-dimensional diffusion to a one-dimensional stopping problem with time-dependent costs, admitting semi-analytical treatment.

For a deterministic cost path $C(t)$, we define the restricted value function $V(d; x^*)$ as the expected value at $t = 0$ given current benefit $D_0 = d$ under a threshold strategy x^*:

$$V(d; x^*) = \varphi(x^*) \cdot \mathbb{E}_d[e^{-r\tau_{x^*}}] - \mathbb{E}_d[e^{-r\tau_{x^*}} C(\tau_{x^*})], \tag{3.10}$$

where $\tau_{x^*} := \inf\{t \geq 0 : D_t \geq x^*\}$ and $\varphi(x^*) = \lambda x^*/r + \lambda\alpha/r^2$. The Laplace transform $\mathbb{E}_d[e^{-q\tau_{x^*}}]$ for arithmetic Brownian motion is given in Appendix B.8.1.

Substituting the deterministic cost path into (3.10) yields the explicit value function:

$$V(d;x^*) = \left(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - \theta\right) e^{-\kappa_0(x^*-d)} - (C_0 - \theta)\, e^{-\kappa_\kappa(x^*-d)}, \tag{3.11}$$

where the discount parameters are:

$$\kappa_0 = \frac{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha}{\sigma^2}, \qquad \kappa_\kappa = \frac{\sqrt{\alpha^2 + 2(r+\kappa)\sigma^2} - \alpha}{\sigma^2}. \tag{3.12}$$

The term κ_κ discounts the cost premium $(C_0 - \theta)$ at rate $r + \kappa$, reflecting that costs decline during the waiting period.

The optimal threshold x^{**} is characterised by the first-order condition $\partial V/\partial x^*|_{x^*=x^{**}} = 0$, yielding:

$$e^{-(\kappa_\kappa-\kappa_0)(x^{**}-d)} = \frac{\frac{\lambda}{r}\left(\kappa_0 x^{**} - 1\right) + \frac{\lambda\alpha}{r^2}\kappa_0 - \kappa_0\theta}{(C_0-\theta)\kappa_\kappa}. \tag{3.13}$$

When the long-run cost target equals zero ($\theta = 0$), costs follow pure exponential decay $C(t) = C_0 e^{-\kappa t}$. Substituting $\theta = 0$ into the general value function yields:

$$V(d;x^*) = \left(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2}\right) e^{-\kappa_0(x^*-d)} - C_0 e^{-\kappa_\kappa(x^*-d)}. \tag{3.14}$$

The first-order condition for the optimal threshold simplifies to:

$$C_0\kappa_\kappa e^{-(\kappa_\kappa-\kappa_0)(x^{**}-d)} = \frac{\lambda}{r}(\kappa_0 x^{**} - 1) + \frac{\lambda\alpha}{r^2}\kappa_0. \tag{3.15}$$

Unlike the general transcendental equations typically solved numerically in the real options literature (e.g. Dixit and Pindyck, 1994), this equation can be rearranged into a canonical form admitting an explicit solution using the Lambert W function, as carried out in Appendix B.8.6:[5]

$$x^{**} = x_{\text{base}} + \frac{1}{\Delta\kappa}\mathcal{W}\left(\frac{C_0\kappa_\kappa\Delta\kappa\, r}{\lambda\kappa_0}\, e^{\Delta\kappa(d-x_{\text{base}})}\right), \tag{3.16}$$

where $x_{\text{base}} = 1/\kappa_0 - \alpha/r$ and $\Delta\kappa = \kappa_\kappa - \kappa_0$.

The second-order condition verifies that x^{**} constitutes a unique maximum. Substituting the first-order condition into $\partial^2 V/\partial(x^*)^2$ yields a strictly negative value.[6]

While the Lambert $\mathcal{W}$ solution (3.16) characterises the optimal target x^{**} from any starting state, the operationally relevant question is at what benefit level d^*

[5]The Lambert W function is the inverse of $f(w) = we^w$. As noted by Corless et al. (1996), it provides closed-form solutions for transcendental equations where the variable appears in both base and exponent. In this context, it represents the exact intersection of the exponential cost-decay path and the linear benefit-expansion boundary.

[6]This holds provided $\kappa_\kappa > \kappa_0$, guaranteed by $\kappa > 0$, which increases the effective discount rate applied to costs ($r + \kappa > r$).

adoption should occur immediately given the current cost C_t. We derive this stopping boundary by imposing the fixed-point condition $d = x^{**}$: substituting into the first-order condition eliminates the exponential term, since the expected time-value of waiting vanishes.[7] This reduction resolves an asymmetry in the model: while forecasting a future target requires path-dependent expected values, execution of the dynamic decision rule must be strictly state-dependent. Setting the expected waiting time to zero removes the non-linear structure introduced by forecasting, and the Lambert $\mathcal{W}$ expression reduces exactly to the linear operational boundary.

For the general case with $\theta \neq 0$, the boundary condition from (3.13) simplifies to:

$$(C-\theta)\kappa_\kappa = \frac{\lambda}{r}(\kappa_0 d^* - 1) + \frac{\lambda\alpha}{r^2}\kappa_0 - \kappa_0\theta. \tag{3.17}$$

Solving for d^* yields an explicit linear candidate boundary, valid for any long-run cost floor θ:

$$d^*(C) = \underbrace{\left(\frac{1}{\kappa_0} - \frac{\alpha}{r} + \frac{r\theta}{\lambda}\right)}_{\text{Intercept}} + \underbrace{\left[\frac{r\kappa_\kappa}{\lambda\kappa_0}\right]}_{\text{Slope}}(C-\theta). \tag{3.18}$$

For the special case of pure exponential decay ($\theta = 0$), this reduces to:

$$d^*(C) = x_{\text{base}} + \frac{r\kappa_\kappa}{\lambda\kappa_0}C. \tag{3.19}$$

Because $C(t)$ is deterministic, (D_t, C_t) is a Markov state, so $d^*(C)$ has the form of the optimal rule; its level comes from the constant-trigger fixed point rather than from smooth pasting on (D, C), and the numerical solutions of Section 3.4 provide the reference.

The linearity of (3.18) implies a simple decision rule: adopt when the current unit cost advantage d_t exceeds the boundary $d^*(C_t)$, whose intercept and slope are fully determined by the model primitives $(\alpha, \sigma, r, \lambda, \kappa, \theta)$ through the discount parameters κ_0 and κ_κ defined in (3.12). Table 6 summarises the boundary expressions for the deterministic cost cases alongside the constant-cost baseline from Chapter 2.

The analytical structure of the linear boundary reveals two economic insights. First, the slope of the stopping boundary is determined by the ratio κ_κ/κ_0. Since κ_κ increases with the decay rate κ, faster expected cost decline produces a steeper boundary. This can be quantified by comparing the operational boundary (3.18) to the constant-cost baseline $d^*_{\text{const}}(C) = x_{\text{base}} + \frac{r}{\lambda}C$, which corresponds to $\kappa = 0$

[7] The divergence between the static target and the dynamic boundary stems from the asymmetric effect of cost decay. Computing the ex-ante target requires evaluating the expected discount factor $\mathbb{E}[e^{-(r+\kappa)\tau_x}]$ over a random hitting time. Due to the convexity of the exponential function, Jensen's inequality embeds path-dependent volatility effects into the expectation, rendering the Lambert $\mathcal{W}$ solution conditional on the initial state (C_0, D_0).

Case	**Boundary $d^*(C)$**	**Eq.**
Constant cost (Ch. 2)	$x_{\text{base}} + \dfrac{r}{\lambda} C$	(2.19)
Exponential decay ($\theta = 0$)	$x_{\text{base}} + \dfrac{r\kappa_\kappa}{\lambda\kappa_0} C$	(3.19)
General OU mean ($\theta \neq 0$)	$x_{\text{base}} + \dfrac{r\theta}{\lambda} + \dfrac{r\kappa_\kappa}{\lambda\kappa_0}(C - \theta)$	(3.18)

Table 6: Deterministic cost boundary cases, where $x_{\text{base}} = 1/\kappa_0 - \alpha/r$. All three are linear in C. The slope increases from r/λ (constant cost) to $r\kappa_\kappa/(\lambda\kappa_0)$ (decaying cost), reflecting the waiting premium from anticipated cost reductions. Setting $\kappa = 0$ recovers the constant-cost case.

(Dixit and Pindyck, 1994). The vertical wedge between these boundaries defines the waiting premium:

$$\Delta d^*(C) = d^*_{\text{decay}}(C) - d^*_{\text{const}}(C) = \frac{r}{\lambda}\left(\frac{\kappa_\kappa}{\kappa_0} - 1\right) C. \tag{3.20}$$

For fixed current cost C_t, firms anticipating rapid cost reductions require higher current benefit d^* to justify immediate adoption, effectively "waiting for the sale" by trading delayed benefits for deeper near-term cost savings. As learning speed κ increases, the waiting premium expands, making firms in sectors with rapid technological progress significantly more selective to avoid investing just before major cost declines.

Second, the intercept and slope of (3.18) respond differently to policy interventions that target θ versus κ. The sign of $\partial d^*/\partial\theta$ and the aggregate implications for heterogeneous adopters are examined in Section 4.4 of Chapter 4, illustrated in Figure 22.

The exponential-affine family bridges the closed-form models of Chapter 2 and the stochastic two-factor system of Section 3.1. The OU mean path captures long-run cost convergence while keeping the threshold tractable analytically. These results establish benchmarks for evaluating numerical solutions: the magnitude of threshold shifts from cost trends, the interaction between learning and discounting through $r + \kappa$, and the linear boundary structure.

The strictly linear boundary $d^*(C)$ arises because the value of waiting scales linearly with cost magnitude when costs are deterministic. Moving to a genuinely stochastic cost ($\eta > 0$) introduces three structural features that, individually and jointly, preclude closed-form solutions and motivate the numerical treatment of Section 3.4.

The tractability of closed-form solutions in the real options literature rests on a narrow set of structural assumptions that constitute the boundary conditions of analytical feasibility. The canonical results of McDonald and Siegel (1986) and Dixit and Pindyck (1994) require that the state variable follow geometric Brownian motion under an infinite horizon with a time-invariant payoff. These assumptions jointly imply stationarity of the value function ($V_t = 0$), reducing the HJB variational inequality to an elliptic ordinary differential equation whose general solution has the power-function form $F(V) = AV^{\beta_1}$. The characteristic equation $\frac{1}{2}\sigma^2\beta(\beta-1) + \alpha\beta - r = 0$ then yields the investment threshold in closed form. Each departure from this baseline destroys one component of the solution structure.

Introducing a finite planning horizon $T < \infty$ reinstates the time derivative V_t, converting the elliptic problem into a parabolic one. The value function $V(t, \mathbf{x})$ becomes time-dependent, and the optimal exercise boundary $\Gamma(t)$ shifts as the horizon approaches, precluding a stationary closed-form threshold. Finite horizons are not merely a theoretical refinement: vehicle replacement cycles, organisational budget periods, and regulatory timelines impose genuinely finite decision windows on fleet operators. Collapsing T and the operational lifetime L into a single parameter, as is common in applied models, further conflates two economically distinct effects, strategic flexibility and payoff amortisation, that the two-parameter framework of Section 3.3 separates explicitly.

Replacing geometric Brownian motion with mean-reverting processes, such as the Ornstein–Uhlenbeck specification adopted for the fixed adoption cost C_t, introduces a non-constant drift term $\kappa(\theta - C_t)$ into the generator $\mathcal{L}$ of equation (3.6). This breaks the scale-invariance that makes power-function solutions exact: the resulting ODE has non-constant coefficients and does not admit elementary closed-form solutions except in a handful of special cases, such as the parabolic cylinder function solution of Falbo et al. (2021), which applies only to the perpetual single-factor problem. Mean reversion is economically necessary here because clean technology prices exhibit both short-term volatility from commodity shocks and long-run convergence toward production-cost parity.

Extending to two correlated state variables (D_t, C_t) adds a mixed second-order term $\rho\sigma\eta V_{cd}$ to the HJB equation, where ρ is the correlation between the unit cost advantage and the fixed adoption cost. The covariance matrix is constant, so a linear change of variables removes the cross-derivative itself, but the rotation mixes the arithmetic Brownian and Ornstein–Uhlenbeck drifts across both new coordinates and shears the adoption payoff away from them, leaving the problem two-dimensional. The free boundary Γ is no longer a scalar threshold x^* but a curved manifold $d = b(c)$ in the two-dimensional state space. This manifold depends on the current level of C_t relative to its long-run parity θ: when costs are far above parity, the required unit cost advantage to trigger adoption is high; when costs are near or below parity, the hurdle falls. No analytical expression for $b(c)$ is available for nonzero ρ. Economically, the sign of ρ determines whether the two sources of uncertainty amplify or offset one another in the quantity that matters, the exercise payoff. Because that payoff is a *difference*, $\frac{\lambda}{r}D_t - C_t$, its instantaneous variance is

$$\left(\tfrac{\lambda}{r}\right)^2 \sigma^2 + \eta^2 - 2\rho\,\tfrac{\lambda}{r}\,\sigma\eta, \tag{3.21}$$

which is decreasing in ρ. Positively correlated factors therefore hedge each other: a benefit shock arrives together with a cost shock in the same direction, and the difference moves less than either component. Negative correlation drives the two apart and makes the payoff more volatile. Negative correlation consequently carries the greater option value of waiting and the higher adoption hurdle, and positive correlation the lower, which is the pattern the numerical sensitivities of Figure 26 display. The direction is the one familiar from the exchange-option literature (McDonald and Siegel, 1986), and neither effect is capturable by independent single-factor models.

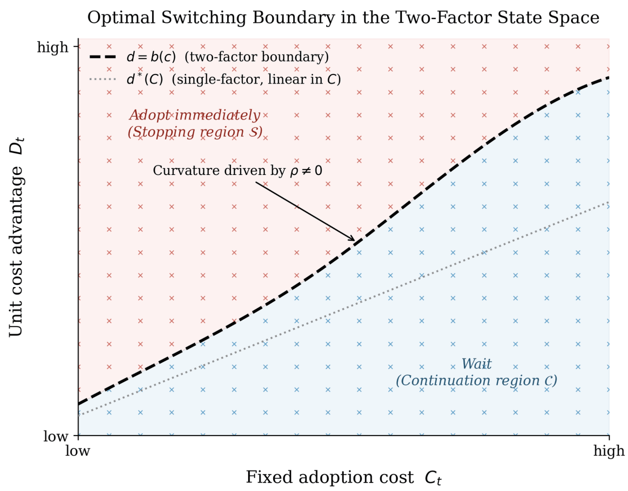

Figure 9: Optimal switching boundary $d^* = b(c)$ in the two-factor state space (D_t, C_t). The dashed curve is the two-factor boundary; the dotted line is the single-factor threshold $x^*(C)$, linear in C. Curvature is driven by nonzero correlation ρ between the unit cost advantage and the fixed adoption cost. The stopping region $\mathcal{S}$ lies above the boundary; the continuation region $\mathcal{C}$ lies below.

3.3 Finite-Horizon Extension

The infinite-horizon model assumes the adoption option never expires and generates perpetual benefits. In practice, delivery fleet operators face two binding constraints: vehicles operate for finite periods before retirement (typically 5-10 year lease terms or depreciation schedules), and strategic decisions must be made by explicit deadlines (regulatory compliance dates, subsidy programme expirations, budgeting cycles).

This section extends the two-factor model to incorporate two independent time parameters:

Decision horizon (T). The deadline by which the adoption decision must be made. This may reflect regulatory deadlines, subsidy programme expiration, or strategic planning cycles. If no adoption occurs by time T, the option expires.

Operating horizon (L). The period over which an adopted vehicle generates benefits. Once adopted at time $\tau \leq T$, the vehicle operates for L years (e.g., lease term or depreciation period), generating operating-cost savings over $[\tau, \tau + L]$.

This separation decouples the decision deadline from the asset's economic life. A vehicle adopted at $\tau < T$ continues operating beyond the decision horizon, contributing residual value over $[T, \tau + L]$. Unlike the infinite-horizon case where adoption value is a perpetuity, the finite-horizon payoff is the present value of a finite annuity with stochastic drift.

The finite-horizon problem seeks the optimal stopping time $\tau \in [0, T]$ maximising the expected discounted payoff. The value function $V(t, c, d)$ represents the value of the adoption option at time t, given state $(C_t, D_t) = (c, d)$:

$$V(t, c, d) = \sup_{\tau \in \mathcal{T}_{t,T} \cup \{\infty\}} \mathbb{E}_{t,c,d}\left[e^{-r(\tau - t)} \hat{g}^L(C_\tau, D_\tau) \mathbf{1}_{\{\tau \leq T\}}\right], \tag{3.22}$$

where $\mathcal{T}_{t,T}$ is the set of stopping times taking values in $[t, T]$ and $\hat{g}^L(c, d)$, given in (3.24) below, is the net present value of adoption at the moment of switching. Admitting $\tau = \infty$ represents the decision never to adopt, which pays nothing; without it the supremum would force adoption by T even at states where $\hat{g}^L$ is negative, and would be inconsistent with the terminal condition (3.29).

The payoff function accounts for the drift in unit cost advantage over the finite operating horizon L. Given that D_t follows arithmetic Brownian motion with drift α, the expected benefit stream over $[\tau, \tau + L]$ is:

$$\hat{g}^L(c, d) = \mathbb{E}\left[\int_0^L e^{-rs} \lambda D_{\tau+s}\, ds \,\middle|\, D_\tau = d\right] - c. \tag{3.23}$$

Evaluating the integral, which is carried out in Appendix B.8.3, yields:

$$\hat{g}^L(c, d) = \lambda d \cdot A_L + \lambda \alpha \cdot B_L - c, \tag{3.24}$$

where the annuity factors are:

$$A_L = \frac{1 - e^{-rL}}{r}, \tag{3.25}$$

$$B_L = \frac{1 - e^{-rL}(1 + rL)}{r^2}. \tag{3.26}$$

The term A_L captures the present value of the constant component d, while B_L accounts for the expected growth α over the operating period. As $L \to \infty$, these expressions recover the perpetuity formulas: $A_L \to 1/r$ and $B_L \to 1/r^2$.

The finite-horizon payoff (3.24) differs from the infinite-horizon benchmark. In the infinite-horizon model developed in Subsection 3.1.2, adoption yields a perpetual benefit stream:

$$\hat{g}^\infty(c, d) = \frac{\lambda d}{r} + \frac{\lambda \alpha}{r^2} - c. \tag{3.27}$$

For typical parameter values ($r \approx 0.03$, $L = 10$ years), the annuity factors $A_{10} \approx 8.64$ and $B_{10} \approx 41.0$ represent approximately 26% and 3.7% of their perpetuity counterparts $1/r \approx 33.3$ and $1/r^2 \approx 1{,}111$. At the baseline state (C_0, D_0) = (\$5,800, 1.45 \$/100mi) the finite-horizon payoff is about 6% of the perpetual payoff gross of the adoption cost, and about 3.5% of it net of C_0, the difference between the two reflecting that the same fixed cost is subtracted from a much smaller benefit stream. This lower intrinsic value implies that the optimal switching threshold must be higher in the finite-horizon case to justify the same upfront cost.

During the waiting phase, the firm receives no flow benefits; all value derives from the option to adopt at a favourable future state. The value function therefore satisfies the homogeneous[8] Hamilton–Jacobi–Bellman (HJB) partial differential equation (PDE) subject to the exercise constraint:

$$\min\left\{ rV - \frac{\partial V}{\partial t} - \mathcal{L}V,\;\; V(t,c,d) - \hat{g}^L(c,d) \right\} = 0, \tag{3.28}$$

with terminal condition:

$$V(T,c,d) = \max\{\hat{g}^L(c,d), 0\}. \tag{3.29}$$

Here $\mathcal{L}$ is the infinitesimal generator defined in equation (3.6).[9] The absence of a flow term λd in (3.28) reflects that benefits accrue only post-adoption, as captured by the obstacle $\hat{g}^L(c,d)$.

3.4 Numerical Solution Methodology

The analytical methods of Chapter 2 provide exact solutions for the constant-cost case, and the deterministic cost trajectories of Section 3.2 admit semi-analytical treatment. The two-factor stochastic model and its finite-horizon extension require numerical solution. This section presents the computational framework: general principles of finite-difference discretisation, variational inequality formulation, and iterative solution, followed by specific implementation details for the infinite-horizon case and extensions required for finite horizons. Figure 10 provides an overview of both solution pipelines.

3.4.1 General Numerical Framework

The numerical approach consists of three main components. First, the continuous state space is discretised using finite difference approximations, replacing derivatives

[8] In the context of PDEs, “homogeneous” refers to the absence of a source term (reflecting zero cash flow while waiting). This should not be confused with time-homogeneity; the finite-horizon problem is inherently time-dependent due to the fixed deadline T.

[9] The infinitesimal generator $\mathcal{L}$ encodes the local stochastic dynamics of the state variables C and D (drift, volatility, and mean reversion). These dynamics describe the evolution of the underlying market variables and are independent of the decision-maker’s planning horizon T. Consequently, the spatial operator $\mathcal{L}$ is identical in both the finite- and infinite-horizon formulations; the mathematical difference lies solely in the presence of the time-derivative term $\partial V/\partial t$ in the finite-horizon HJB equation.

in the infinitesimal generator with algebraic expressions on a spatial grid. Second, the temporal structure is handled differently depending on whether the problem is infinite-horizon (stationary) or finite-horizon (time-dependent), leading to either an elliptic or parabolic variational inequality. Third, the resulting discrete system is formulated as a linear complementarity problem and solved iteratively using projected successive over-relaxation. While these general principles apply universally, specific implementation details: grid parameters, convergence tolerances, and boundary conditions, are tailored to each model variant and presented in the model-specific sections that follow.

The structure of this section follows the logical sequence of the numerical solution process: we begin with spatial discretisation and domain truncation, then distinguish between stationary and time-marching formulations (temporal formulation), present the linear complementarity problem and its iterative solution (the linear complementarity formulation), and conclude with validation and convergence assessment.

Finite Difference Discretisation

The finite difference method approximates the continuous differential operator $\mathcal{L}$ by replacing derivatives with discrete algebraic expressions evaluated on a spatial grid (Duffy, 2013). For a state vector $\mathbf{X} \in \mathbb{R}^n$, we construct a computational mesh ω_h with grid spacings Δx_j in each dimension $j = 1, \ldots, n$. The value function $V(\mathbf{X}, t)$ is then approximated at discrete grid points, with V_i^k denoting the approximation at spatial location i and time step k.

Since the state processes in this thesis have unbounded domains, such as Brownian motions or mean-reverting diffusions, the computational domain must be truncated to a finite region $[\mathbf{X}_{\min}, \mathbf{X}_{\max}]$. Following established practice (Tavella and Randall, 2000; Kangro and Nicolaides, 2000), these bounds are chosen such that boundary effects do not materially affect the solution in the region of economic interest. For diffusion processes, typical truncation extends to several multiples of the standard deviation around the current value or long-run mean. For mean-reverting processes, the domain extends several standard deviations around the mean-reversion level. These heuristics ensure the free boundary lies well within the computational domain while keeping the problem size manageable. The adequacy of domain truncation is verified through grid refinement studies, as discussed in the convergence assessment below. Specific truncation bounds for each model are reported in the respective chapters.

The discretisation of the infinitesimal generator $\mathcal{L}$ requires careful treatment to ensure stability and convergence to the correct viscosity solution. Second-order diffusion terms are discretised using standard central differences:

$$\frac{1}{2}\sigma^2 \frac{\partial^2 V}{\partial x^2} \approx \frac{1}{2}\sigma^2 \frac{V_{i+1} - 2V_i + V_{i-1}}{(\Delta x)^2}. \tag{3.30}$$

However, first-order drift terms pose a more delicate issue. Standard central differencing, while second-order accurate, can produce spurious oscillations when convection dominates diffusion (Hundsdorfer and Verwer, 2013). To guarantee monotonicity, a necessary condition for convergence to the unique viscosity solution (Barles and Souganidis, 1991), we employ upwind differencing for all drift derivatives:

$$\mu \frac{\partial V}{\partial x} \approx \begin{cases} \mu \dfrac{V_{i+1} - V_i}{\Delta x} & \text{if } \mu > 0, \\ \mu \dfrac{V_i - V_{i-1}}{\Delta x} & \text{if } \mu < 0. \end{cases} \tag{3.31}$$

This choice biases the approximation in the direction of information flow, ensuring that the discretised system satisfies the positive coefficient property required for the maximum principle.

For two-dimensional problems with correlated state variables, the discretisation includes central differences for the pure second derivatives in each dimension (as in equation 3.30), while mixed derivative terms arising from correlation require careful discretisation. While the natural 9-point stencil provides a centered approximation to the cross-derivative $\frac{\partial^2 V}{\partial x_1 \partial x_2}$, it can lead to instability when the correlation is strong. We therefore employ a 7-point stencil that retains only the diagonal pair compatible with the sign of the correlation and drops the other, which has been shown to preserve positivity of coefficients and ensure monotonicity even for high correlation (Pooley et al., 2003), subject to a condition on the grid aspect ratio derived in Appendix C.2. The complete discretisation yields a sparse linear system whose structure depends on the dimensionality and correlation of the underlying state processes.

Temporal Treatment: Stationary and Time-Dependent Problems

The numerical framework must accommodate both infinite-horizon (time-stationary) and finite-horizon (time-dependent) formulations, as these lead to different computational structures and reflect distinct economic scenarios. In the infinite-horizon setting, the optimal stopping rule is time-stationary: the decision to adopt depends only on the current state $\mathbf{X}$, not on calendar time. The HJB variational inequality reduces to an elliptic problem:

$$\max\{(\mathcal{L} - r)V(\mathbf{X}) + f(\mathbf{X}),\ \hat{g}(\mathbf{X}) - V(\mathbf{X})\} = 0, \tag{3.32}$$

where $f(\mathbf{X})$ represents any running payoff and $\hat{g}(\mathbf{X})$ is the terminal payoff upon exercise. This is solved as a global system over the entire spatial grid simultaneously. The absence of a time-stepping term ($\partial V/\partial t$) means there is no natural damping in the discrete system, which can lead to slow convergence or ill-conditioning. This motivates the use of successive over-relaxation and continuation methods discussed in Subsection 3.4.1.

When a finite maturity T is imposed, the value function becomes explicitly time-dependent: $V(\mathbf{X}, t)$. The problem becomes parabolic:

$$\max\left\{\frac{\partial V}{\partial t} + (\mathcal{L} - r)V + f,\ \hat{g} - V\right\} = 0, \quad t < T, \tag{3.33}$$

with terminal condition $V(\mathbf{X}, T) = \max\{\hat{g}(\mathbf{X}), 0\}$. Following Brennan and Schwartz (1978), we discretise the time derivative using a θ-scheme:

$$\frac{V^k - V^{k+1}}{\Delta t} = \theta(\mathcal{L}_h - r)V^k + (1-\theta)(\mathcal{L}_h - r)V^{k+1} + f^k, \tag{3.34}$$

where $\mathcal{L}_h$ is the discrete spatial operator of Subsection 3.4.2 and, under backward induction, V^{k+1} is the already computed later time level while V^k is the unknown. The parameter $\theta \in [0, 1]$ therefore controls the implicitness: $\theta = 0$ places the operator entirely on the known level and yields explicit Euler (conditionally stable), $\theta = 1$ places it entirely on the unknown level and yields fully implicit Euler (unconditionally stable but first-order accurate), and $\theta = 1/2$ yields the Crank–Nicolson scheme (unconditionally stable and second-order accurate). The solution proceeds via backward induction from the known terminal condition, solving a sequence of spatial problems at each discrete time step Δt. We employ fully implicit discretisation ($\theta = 1$) to ensure unconditional stability, allowing time steps determined by accuracy rather than stability constraints. While explicit schemes ($\theta = 0$) would impose severe stability restrictions $\Delta t \lesssim \Delta C^2/\eta^2$, the implicit scheme removes this limitation at the cost of solving a linear system at each time step.

The fully implicit scheme produces, at each backward time step, a linear system of the form $(\mathbf{I} - \Delta t(\mathcal{L}_h - r))\, V^k = V^{k+1} + \Delta t\, f^k$, where $k+1$ denotes the known future time level and k the unknown current level. The optimal stopping constraint $V^k \geq \hat{g}^L$, however, prevents direct linear inversion: grid points where immediate adoption is optimal must satisfy $V^k_{i,j} = \hat{g}^L_{i,j}$ rather than the interior PDE residual. This constraint transforms the linear system into a linear complementarity problem (LCP), which PSOR solves at each time step by projecting iterates onto the admissible set. In the stationary (infinite-horizon) formulation, no time-stepping is required; PSOR solves the resulting elliptic LCP directly on the full spatial grid.

Linear Complementarity Formulation and Iterative Solution

The variational inequality characterising the optimal stopping problem can be equivalently expressed as a linear complementarity problem (LCP). This reformulation is particularly natural for American-style exercise problems because it explicitly enforces the constraint that the value function must remain above the exercise payoff (Wilmott et al., 1995).

Definition 3.5 (Linear Complementarity Problem)**.** Given a matrix $\mathbf{M} \in \mathbb{R}^{N\times N}$ and a vector $\mathbf{q} \in \mathbb{R}^N$, the LCP seeks a vector $\mathbf{w} \in \mathbb{R}^N$ satisfying:

$$\mathbf{w} \geq \mathbf{0}, \quad \mathbf{M}\mathbf{w} + \mathbf{q} \geq \mathbf{0}, \quad \mathbf{w}^T(\mathbf{M}\mathbf{w} + \mathbf{q}) = 0. \tag{3.35}$$

In the context of optimal stopping, we define the slack variable $\mathbf{w} = \mathbf{V} - \hat{\mathbf{g}}$, representing the continuation premium. The complementarity condition ensures that at each grid point either the firm waits ($w_i > 0$ and the PDE holds with equality) or it exercises immediately ($w_i = 0$ and $V_i = \hat{g}_i$). This formulation is mathematically equivalent to the obstacle problem in variational calculus, where the solution must lie above the obstacle surface defined by $\hat{g}(\mathbf{X})$ (Wilmott et al., 1995).

We solve the LCP using the projected successive over-relaxation (PSOR) method (Cryer, 1971), an iterative algorithm that alternates between a Gauss–Seidel update and a projection onto the feasible set. At each grid point i and iteration m, the algorithm first computes the unconstrained iterate using the most recently available values:

$$\tilde{V}_i^{(m)} = \frac{1}{M_{ii}} \left(b_i - \sum_{j<i} M_{ij} V_j^{(m)} - \sum_{j>i} M_{ij} V_j^{(m-1)} \right), \qquad \mathbf{b} := \mathbf{M}\hat{\mathbf{g}} - \mathbf{q}, \tag{3.36}$$

where $\mathbf{b}$ is the right-hand side of the equality $\mathbf{M}(\mathbf{V} - \hat{\mathbf{g}}) + \mathbf{q} = \mathbf{0}$ that holds in the continuation region. This intermediate value is then relaxed by a factor $\omega \in (1, 2)$ and projected onto the obstacle:

$$V_i^{(m)} = \max \left\{ \hat{g}_i,\ V_i^{(m-1)} + \omega \left(\tilde{V}_i^{(m)} - V_i^{(m-1)} \right) \right\}. \tag{3.37}$$

The relaxation parameter ω accelerates convergence by extrapolating in the direction of the Gauss–Seidel update. The projection step ensures that the complementarity constraint is satisfied at every iteration and implicitly enforces the smooth-pasting condition at the free boundary, as the value function meets the exercise payoff tangentially (Wilmott, 2006). Iterations continue until the maximum pointwise change $\max_i |V_i^{(m)} - V_i^{(m-1)}|$ falls below a prescribed tolerance, 10^{-3} for infinite-horizon problems, which at $V \sim 10^5$ is a relative change of about 5×10^{-9}, and 10^{-9} per time slice for time-marching schemes, where the $1/\Delta t$ term added to the diagonal damps the iteration and a tighter tolerance is reached within a few hundred sweeps. The values used are collected in Appendix C.4.

For models with correlation between state variables or complex drift structures, the matrix $\mathbf{M}$ can become ill-conditioned, leading to slow or failed convergence. To address this, we employ a homotopy (continuation) method (Allgower and Georg, 2012). The solver begins with a simplified problem in which the cross-derivative term is omitted, and restores the full operator once the iterate has settled into a stable convergence basin, so that the simplified solution serves as the initial guess for the complete one. The schedule used for the infinite-horizon solves is given in Appendix C.4. This technique proves particularly useful for two-dimensional problems with correlated state variables, where mixed derivatives can destabilise direct iterative methods.

The specific implementation of this framework for the two-factor model is detailed in Subsection 3.4.2.

Convergence, Validation, and Boundary Extraction

The numerical solution is validated against grid refinement, Richardson extrapolation, and the analytical limits of Section 3.2.

To assess discretisation error, we solve each problem on a sequence of progressively refined grids, typically halving the step size Δx in each dimension. Following the framework of Barles and Souganidis (1991), monotone finite difference schemes for viscosity solutions exhibit first-order convergence in L^∞ norm:

$$\|V_h - V\|_\infty \leq Ch^\alpha, \tag{3.38}$$

where h is the grid spacing, $\alpha \approx 1$ for upwind schemes, and C is a problem-dependent constant. We verify this convergence by computing the value at a reference point, typically the current state, across multiple grid resolutions, and confirming that the error decreases monotonically.

To accelerate convergence and improve accuracy, we employ two enhancements. First, a multigrid initialisation strategy reduces iterations required for the Projected Successive Over-Relaxation (PSOR) algorithm. The problem is solved on a coarse mesh ω_{2h}, and the resulting value function is bilinearly interpolated onto the fine mesh ω_h to serve as an initial guess V_0, reducing high-frequency errors early in the iteration process. Second, we apply Richardson extrapolation (Richardson and Gaunt, 1927) to the sequence of solutions to obtain higher-order estimates. Combining results from coarse and fine grids, the extrapolated value function is:

$$V_{\text{extrap}} = \frac{4V_{\text{fine}} - V_{\text{coarse}}}{3} \tag{3.39}$$

This cancels the leading-order error term, yielding a smoother gradient $\partial V/\partial d$ that is critical for accurate identification of the optimal switching boundary. We report results from grids where further refinement yields only marginal improvements in accuracy at disproportionate computational cost, representing a practical balance between precision and efficiency.

The optimal stopping boundary is identified as the locus of points where the continuation value equals the exercise payoff: $\{\mathbf{X} : V(\mathbf{X}) = \hat{g}(\mathbf{X})\}$. On a discrete grid, this condition is satisfied only approximately, yielding an uncertainty corridor of width proportional to the grid spacing. To extract a continuous economic threshold suitable for interpretation and policy analysis, the quadratic decay of the continuation premium implied by smooth pasting is extrapolated to its root, and the resulting boundary points are smoothed with a monotonic spline (de Boor, 1978). The procedure and its justification are given in Subsection 3.5.2.

Whenever possible, numerical results are cross-validated against analytical benchmarks. For models without closed-form solutions, we validate by examining limiting cases. For instance, in the two-factor model with stochastic costs, setting cost volatility to zero ($\eta \to 0$) should recover the deterministic cost thresholds from Section 3.2, while additionally fixing mean reversion ($\kappa \to 0$) should recover the constant-cost threshold from Chapter 2. Confirming such limiting behaviour provides strong evidence that the numerical implementation correctly captures the underlying economic structure across different model specifications. Additionally, the adequacy of domain truncation is verified by solving problems with different window sizes and confirming that results in the region of interest remain unchanged, ensuring boundary conditions do not introduce artefacts. Monte Carlo simulation is used as a cross-check on selected results: a least-squares Monte Carlo solve of the finite-horizon problem, reported in Chapter 4 and parameterised in Appendix C.5, which is independent of the finite-difference scheme. The specific convergence tolerances, grid dimensions, and computational costs for each model are reported in the respective chapters, along with convergence analysis demonstrating that solutions achieve the expected accuracy.

3.4.2 Two-Factor Finite-Difference Solution

The variational inequality (3.9) admits no analytical solution for nonzero correlation $\rho \neq 0$, in contrast to the constant-cost model of Chapter 2. Building on the general numerical framework established in Subsection 3.4.1, we discretise the problem on a rectangular grid and formulate the obstacle constraint $V \geq \hat{g}$ as a linear complementarity problem (LCP). The main algorithmic steps are:

1. truncate the unbounded domain and impose boundary conditions consistent with asymptotic behaviour;
2. discretise the PDE on a rectangular grid using finite-difference stencils;
3. apply the discrete stationary operator $(\mathcal{L}_h - r)$ node-by-node (matrix-free);
4. impose the obstacle constraint $V \geq \hat{g}$ via an LCP formulation;
5. solve the resulting LCP using Projected Successive Over-Relaxation (PSOR) with adaptive relaxation parameter ω;
6. extract the optimal adoption boundary $d = b(c)$ and refine estimates using Richardson extrapolation.

The following subsections detail each stage, emphasising the specific numerical challenges and design choices required for the two-dimensional correlated diffusion setting.

Domain Truncation and Grid Discretisation

The continuous state space is truncated to a finite rectangular computational domain $\Omega = [C_{\min}, C_{\max}] \times [D_{\min}, D_{\max}]$. The boundaries are chosen heuristically to ensure

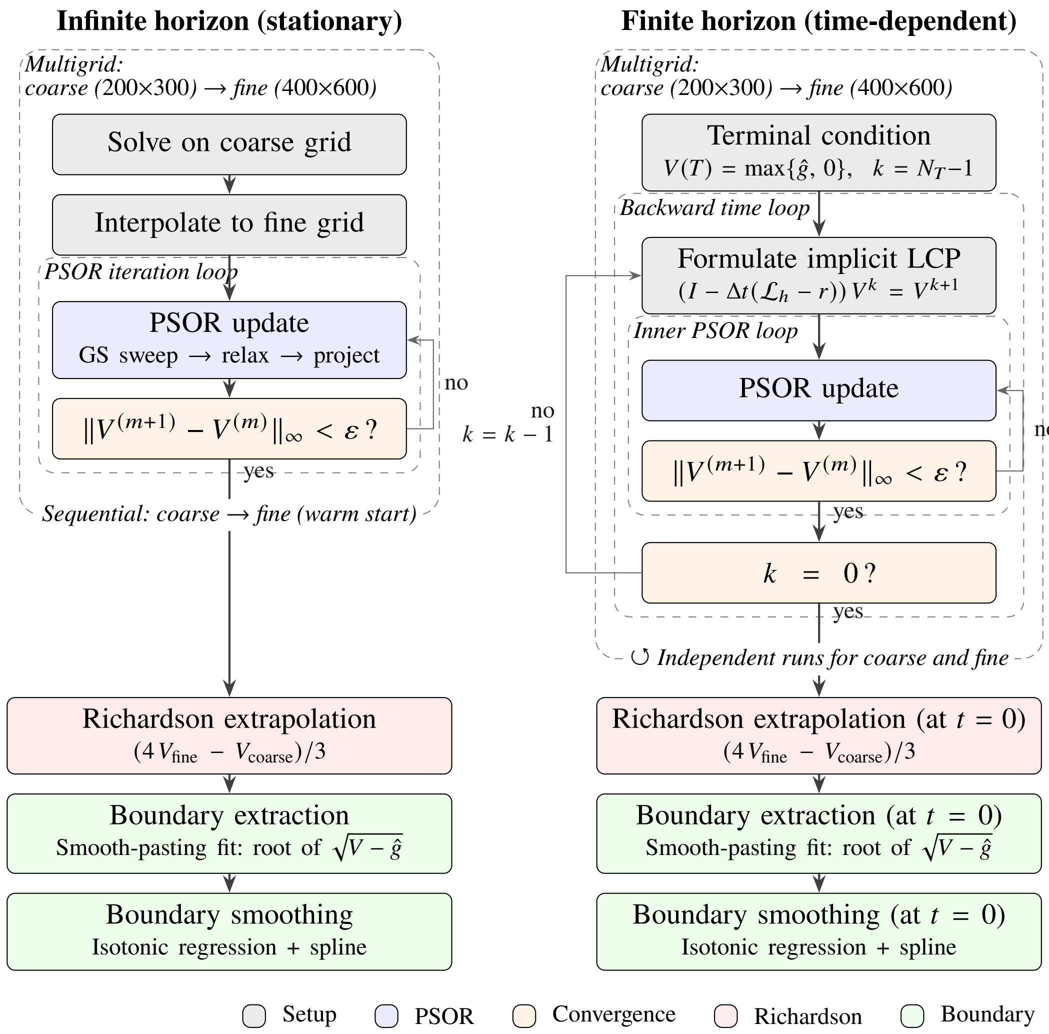

Figure 10: Computational pipeline for the infinite-horizon (left) and finite-horizon (right) numerical solutions. The stationary problem requires a single PSOR iteration loop with multigrid warm-starting (Subsection 3.4.2). The time-dependent problem nests the same inner PSOR solver within a backward time march from $t = T$ to $t = 0$ (Subsection 3.4.3). Both pipelines solve on two grid resolutions and combine the results via Richardson extrapolation before extracting the optimal boundary. Full implementation details are provided in Appendix C.

they lie sufficiently far from the region of interest, centered around the initial cost $C_0 \approx 5800$, to minimise boundary artefacts. For the baseline calibration, we set

$$C_{\min} = -8000, \quad C_{\max} = 12000, \quad D_{\min} = -30, \quad D_{\max} = 30. \tag{3.40}$$

These bounds extend approximately two to five standard deviations from the expected range of the state variables, ensuring that the expected switching boundary lies well inside the computational domain.

A uniform grid is employed for both state variables:

$$C_i = C_{\min} + i\,\Delta C, \quad i = 0, \dots, N_C, \qquad D_j = D_{\min} + j\,\Delta D, \quad j = 0, \dots, N_D, \tag{3.41}$$

with discretisation steps

$$\Delta C = \frac{C_{\max} - C_{\min}}{N_C}, \qquad \Delta D = \frac{D_{\max} - D_{\min}}{N_D}. \tag{3.42}$$

The fine grid uses $N_C \times N_D = 400 \times 600$ points, yielding $\Delta C = 50$ and $\Delta D = 0.1$. The benefit resolution is set by the positive-coefficient condition of Appendix C.2 rather than by accuracy alone: an aspect ratio $\Delta C/\Delta D = 500$ keeps every off-diagonal weight non-negative across the whole sensitivity range, including the largest cost volatility examined, $\eta = \$3{,}000/\sqrt{\text{year}}$. A coarser 200×300 grid is employed for initialisation via multigrid warm-starting (see Subsection 3.4.2). While non-uniform meshes could increase resolution near the switching boundary, grid-refinement tests confirm that uniform grids achieve adequate accuracy while simplifying matrix assembly and preserving sparsity, which is critical for the efficiency of iterative solvers.

Finite Difference Stencils

For any twice continuously differentiable test function v, the infinitesimal generator of the joint Ornstein–Uhlenbeck–Brownian diffusion is

$$\mathcal{L}v = \kappa(\theta - c)\, v_c + \alpha\, v_d + \tfrac{1}{2}\eta^2 v_{cc} + \tfrac{1}{2}\sigma^2 v_{dd} + \rho\,\eta\sigma\, v_{cd}\,, \tag{3.43}$$

where subscripts denote partial derivatives. In the continuation region, the stationary PDE for V is therefore

$$(\mathcal{L} - r)\, V(c, d) = 0\,. \tag{3.44}$$

Because all coefficients are constant or at most linear in the state variables, a centered finite-difference discretisation provides second-order local accuracy. Special care, however, is required for the cross-derivative term to maintain the monotonicity of the discrete operator.

At interior nodes (i, j) with $1 \le i \le N_C - 1$ and $1 \le j \le N_D - 1$, second derivatives are computed by standard central differences, while the first derivatives use directional (upwind) splitting of the drift to preserve monotonicity of the discrete operator:

$$V_{cc} \approx \frac{V_{i+1,j} - 2\,V_{i,j} + V_{i-1,j}}{\Delta C^2}, \qquad V_{dd} \approx \frac{V_{i,j+1} - 2\,V_{i,j} + V_{i,j-1}}{\Delta D^2},$$

$$\kappa(\theta - c)\,V_c \;\rightsquigarrow\; \frac{\max\{\mu_C, 0\}}{\Delta C}\big(V_{i+1,j} - V_{i,j}\big) + \frac{\max\{-\mu_C, 0\}}{\Delta C}\big(V_{i-1,j} - V_{i,j}\big),$$

$$\alpha\,V_d \;\rightsquigarrow\; \frac{\max\{\mu_D, 0\}}{\Delta D}\big(V_{i,j+1} - V_{i,j}\big) + \frac{\max\{-\mu_D, 0\}}{\Delta D}\big(V_{i,j-1} - V_{i,j}\big),$$

with $\mu_C = \kappa(\theta - C_i)$ and $\mu_D = \alpha$. In each expression the first term is active when the drift is positive and the second when it is negative, so that the difference quotient is taken in the direction of information flow in both branches and the product reproduces $\mu\,V_x$ with the correct sign. Written this way the upwind terms contribute only non-negative off-diagonal weights, preserving the M-matrix property on which PSOR convergence rests.

To promote the M-matrix property of the discrete operator, the stencil for the mixed derivative $\rho\,\eta\sigma\,V_{cd}$ depends on the sign of the correlation coefficient. For negative correlation $\rho < 0$ (the baseline case $\rho = -0.3$), a 7-point upwind stencil utilising the anti-diagonal neighbors is employed:

$$V_{cd}\big|_{i,j} \approx \frac{V_{i+1,j} + V_{i-1,j} + V_{i,j+1} + V_{i,j-1} - V_{i-1,j+1} - V_{i+1,j-1} - 2V_{i,j}}{2\,\Delta C\,\Delta D}. \tag{3.45}$$

Multiplied by the mixed coefficient $\rho\,\eta\sigma$, which is negative in this branch, the two anti-diagonal terms acquire positive weights $|\rho|\eta\sigma/(2\,\Delta C\,\Delta D)$ while the four cardinal neighbours acquire negative weights of the same magnitude, so the cross term draws down the positive off-diagonal contributions supplied by the unmixed diffusion terms. Keeping every net off-diagonal weight non-negative therefore constrains the grid aspect ratio $\Delta C/\Delta D$ relative to the correlation magnitude $|\rho|$; the condition is derived and evaluated at the baseline calibration in Appendix C.2. Boundary stabilisation treatments at the cost edges are detailed in Subsection 3.4.2.

The discrete operator and the projected iteration

Rather than assembling a global sparse matrix, the solver applies the discrete stationary operator $(\mathcal{L}_h - r)$ node-by-node and advances the solution with a *projected successive over-relaxation* (PSOR) sweep. This matrix-free choice is dictated by the nature of the problem: the value function does not satisfy a linear equation everywhere, but a linear complementarity problem (LCP) arising from the optimal-stopping (free-boundary) structure.

For an interior grid node (i, j), the discrete continuation operator acting on $V_{i,j}$ couples the node to its four axial neighbours and to two diagonal neighbours introduced by the mixed second derivative,

$$(\mathcal{L}_h - r)\, V_{i,j} = a_P\, V_{i,j} + a_W\, V_{i-1,j} + a_E\, V_{i+1,j} \\ + a_S\, V_{i,j-1} + a_N\, V_{i,j+1} + a_{D_1}\, V_{i+s,\,j+1} + a_{D_2}\, V_{i-s,\,j-1}, \tag{3.46}$$

where the coefficients $a_\cdot$ depend on the model parameters α, κ, θ, η, σ, ρ, r and on the grid spacings ΔC, ΔD. The orientation $s \in \{+1, -1\}$ of the diagonal stencil is selected from the sign of the correlation: for $\rho \geq 0$ the cross term uses the NE/SW neighbours ($s = +1$, i.e. $(i + 1, j + 1)$ and $(i - 1, j - 1)$), while for $\rho < 0$ it uses the NW/SE neighbours ($s = -1$). This sign-dependent orientation preserves the monotonicity (M-matrix) structure of the diagonal contribution.

Because the running payoff enters through the early-adoption obstacle rather than as a forcing term, the continuation equation is homogeneous, $(\mathcal{L}_h - r)V_{i,j} = 0$, and the value function solves the LCP

$$\min\Big\{ (r - \mathcal{L}_h)\, V_{i,j},\ \ V_{i,j} - \hat{g}_{i,j} \Big\} = 0, \qquad \text{for all interior } (i, j), \tag{3.47}$$

where $\hat{g}_{i,j}$ is the adoption payoff evaluated at (C_i, D_j). In the infinite-horizon case $\hat{g}_{i,j} = \lambda D_j/r + \lambda\alpha/r^2 - C_i$.

The LCP (3.47) is solved by iterating a Gauss–Seidel relaxation with over-relaxation parameter ω, followed at each node by a projection onto the obstacle. Writing $\tilde{a}_P := -a_P > 0$ for the diagonal coefficient of $(r - \mathcal{L}_h)$, the sign-reversed counterpart of the diagonal entry a_P in (3.46), and $R_{i,j}$ for the sum of all off-diagonal neighbour contributions in (3.46), one PSOR update reads

$$V^{\mathrm{GS}}_{i,j} = \frac{R_{i,j}}{\tilde{a}_P}, \qquad V^{\mathrm{new}}_{i,j} = \max\Big\{ (1 - \omega)\, V^{\mathrm{old}}_{i,j} + \omega\, V^{\mathrm{GS}}_{i,j},\ \ \hat{g}_{i,j},\ 0 \Big\}. \tag{3.48}$$

Neighbour values are read in place, so updates propagate within a single sweep. The sweeps are repeated until the maximum nodal change $\max_{i,j}|V^{\mathrm{new}}_{i,j} - V^{\mathrm{old}}_{i,j}|$ falls below a prescribed tolerance, at which point the stationary solution of the free-boundary problem has been reached. The projection $\max\{\cdot, \hat{g}_{i,j}, 0\}$ enforces the early-adoption constraint and, implicitly, locates the free boundary as the set where the two arguments of (3.47) switch.

Boundary conditions

Boundary conditions are enforced to ensure well-posedness of the discrete system while reflecting the asymptotic economic behaviour of the value function.

Unit Cost Advantage boundaries ($D_{\min}, D_{\max}$). At the lower and upper boundaries in the benefit direction, Dirichlet conditions are imposed after projection onto the feasible set $V \geq \max(\hat{g}, 0)$. The implemented conditions are

$$V(C, D_{\min}) = \max\bigl(\hat{g}(C, D_{\min}),\ 0\bigr) = 0, \tag{3.49}$$

$$V(C, D_{\max}) = \max\bigl(\hat{g}(C, D_{\max}),\ 0\bigr) = \hat{g}(C, D_{\max}), \tag{3.50}$$

where $\hat{g}(c, d) = \frac{\lambda d}{r} + \frac{\lambda \alpha}{r^2} - c$ is the immediate-adoption payoff. At $D_{\min}$ the payoff is negative throughout the computational cost range, so the projected value is identically zero; at $D_{\max}$ the payoff is positive, so the boundary coincides with $\hat{g}$. These boundary nodes (corresponding to $j = 0$ and $j = N_D$) are excluded from the PSOR update sweeps, with their values held fixed throughout the iteration.

Fixed Cost boundaries ($C_{\min}, C_{\max}$). To ensure the numerical domain remains self-contained without fictitious nodes, structural boundary conditions are imposed at the extreme cost edges $i = 0$ and $i = N_C$. The second-order cost-diffusion and cross-diffusion terms are set to vanish,

$$V_{CC}(C_{\min}, D) = V_{CD}(C_{\min}, D) = 0, \tag{3.51}$$

$$V_{CC}(C_{\max}, D) = V_{CD}(C_{\max}, D) = 0, \tag{3.52}$$

which imposes linearity of the value function in the cost direction at the edges, the standard far-field condition once the domain has been placed well outside the switching region. In the discrete operator this sets $a_{cc} = a_{cd} = 0$ at $i = 0$ and $i = N_C$, while the cost drift, the benefit-direction diffusion and the discount term are retained; the outward neighbour, which would lie outside the grid, is replaced by the edge node itself. Unlike the benefit-direction edges, these nodes are not held at fixed Dirichlet values but remain in the PSOR sweep subject to the obstacle $V \geq \max(\hat{g}, 0)$, so the adoption region is free to extend to the cost boundaries. To mitigate numerical shocks originating from the non-smooth payoff function at these boundaries, the finite-horizon solver applies a homotopy-based continuation during the initial time steps. Boundary node updates are constructed as a convex combination of the frozen previous-step value and the fully implicit spatial PDE update:

$$V_{\text{new}} = (1 - \lambda_h)\, V_{\text{prev}} + \lambda_h\, V_{\text{PDE}}\,, \tag{3.53}$$

where the homotopy parameter λ_h ramps linearly from near zero to 1.0 over the first $N_{\text{ramp}} = 10$ initial time steps, smoothly deforming the stable boundary state into the full differential equation and suppressing high-frequency spurious oscillations that can otherwise propagate from singular boundaries (Giles and Carter, 2006).

Discrete linear complementarity formulation and PSOR solution

Discretising the variational inequality (3.9) produces a componentwise linear complementarity problem (LCP) in terms of the discrete values of V:

$$\mathbf{A}\,v - f \;\geq\; 0, \qquad v \;\geq\; g, \qquad (v-g)^{\top}(\mathbf{A}\,v - f) \;=\; 0\,, \tag{3.54}$$

where v is the vector of unknown values $V_{i,j}$ on the interior grid, g is the vector of corresponding payoff values $\hat{g}(C_i, D_j)$, the matrix $\mathbf{A}$ collects the stencil coefficients of the discrete operator $(r - \mathcal{L}_h)$ in (3.46), and f collects the boundary contributions; (3.54) is the matrix statement of the componentwise LCP (3.47). Because $\mathbf{A}$ discretises $(r - \mathcal{L}_h)$, the sign convention here is the same as in the componentwise form (3.47): the inequalities encode the free boundary through $v > g$ and $\mathbf{A}v - f = 0$ in the continuation region, where the PDE holds with equality, and $v = g$ with the strict inequality $\mathbf{A}v - f > 0$ active in the stopping region, where continuing would destroy value.

The stationary system (3.54) is solved using a *Projected Successive Over-Relaxation* (PSOR) scheme, which extends the classical Gauss–Seidel method by combining over-relaxation and projection onto the feasible set $v \geq g$. At each grid node k, the discrete equation for that node is solved locally to obtain a provisional Gauss–Seidel value v_k^{GS}, after which the relaxed and projected update is applied:

$$v_k^{\mathrm{new}} = \max\Big\{\,(1-\omega)\,v_k^{\mathrm{old}} + \omega\,v_k^{\mathrm{GS}},\;\; g_k\Big\},$$

where $\omega \in (1,2)$ is the relaxation parameter and g_k is the discrete payoff at node k. When the obstacle is inactive (continuation region), the projection has no effect and the method reduces to the standard SOR iteration; when the obstacle is active (stopping region), the projection enforces $v_k = g_k$. Iterations proceed in lexicographic order consistent with the matrix assembly, and convergence is declared once $\|v^{(n+1)} - v^{(n)}\|_\infty \leq \varepsilon$ for a prescribed tolerance ε. Under the M-matrix property of A, convergence to the unique discrete solution of (3.54) is guaranteed (Cryer, 1971; Ikonen and Toivanen, 2009).

The convergence rate is sensitive to the choice of relaxation factor. The implementation provides an adaptive schedule, in which ω starts at $\omega_0 = 1.45$ and is decreased by $\delta\omega = 0.005$, subject to $\omega \geq 1$, whenever the residual reduction stalls.[10]

To accelerate convergence, we employ a multigrid warm-start strategy: the problem is first solved on a coarse 200×300 grid, and the result is interpolated (using bilinear interpolation) to the fine 400×600 grid to serve as the initial guess $V^{(0)}$. The termination criterion is the one stated in Subsection 3.5.1; the warm start reduces the number of sweeps needed before the reported outputs stabilise.

[10] The adaptive rule is active only on the warm-started fine solve. The sensitivity solves and the coarse leg of the extrapolation pair run at a fixed $\omega = 1.45$, and the finite-horizon time march at a fixed $\omega = 1.2$; on the grids used here the schedule changes no reported quantity.

Although a pseudo-time-marching approach could in principle be used to reach the stationary solution, preliminary tests showed no improvement in wall-clock performance. The stationary LCP formulation for V solved directly by PSOR was therefore adopted throughout.

Boundary Extraction and Richardson Extrapolation

The discrete solution provides the value function $V_{i,j}$ on the computational grid. Two additional steps refine the solution and extract the optimal adoption boundary.

To improve accuracy from $O(h^2)$ to higher order, we combine the fine grid solution (V_{fine}) and the coarse grid solution (V_{coarse}) using Richardson extrapolation:

$$V_{\text{final}} = \frac{4V_{\text{fine}} - V_{\text{coarse}}}{3}. \tag{3.55}$$

This formula assumes second-order local truncation error in the interior and yields a more accurate approximation of the continuous value function.[11]

The optimal adoption threshold $d = b(c)$ is identified as the contour where $V(c, d) = \hat{g}(c, d)$. Numerically, this is achieved by:

1. Forming the continuation premium $S(c, d) = V(c, d) - \hat{g}(c, d)$, which is non-negative and vanishes identically in the stopping region.
2. For each cost level C_i, regressing $\sqrt{S(C_i, \cdot)}$ on d over the band in which the quadratic smooth-pasting regime holds, and taking the root of the fitted line as d_i^* (Subsection 3.5.2).
3. Smoothing the resulting boundary points (C_i, d_i^*) using isotonic regression, which enforces monotonicity $b'(c) \geq 0$ consistent with the theoretical conjecture from Subsection 3.1.1, followed by a smoothing spline, with monotonicity re-imposed on the fitted curve.

The resulting boundary curve is used for all numerical analysis and visualisation in Chapter 4.

3.4.3 Finite Horizon Model via Backward Time-Marching

The finite-horizon problem differs structurally from the infinite-horizon benchmark solved in Subsection 3.4.2. The infinite-horizon case yields a time-invariant elliptic PDE with a single stationary value surface $V(c, d)$ and time-independent boundary $d = b(c)$, solved by projected successive over-relaxation (PSOR) with no temporal interpretation. The finite-horizon problem introduces explicit time dependence

[11] The assumption is not automatic for an upwind scheme, whose drift terms are formally first order. Subsection 3.5.3 quantifies the first-order contribution, artificial diffusion of 1.4% and 0.6% of the physical diffusion coefficients on the production grid, and verifies that the observed error behaves as $O(h^2)$ over the resolutions tested, which is what the weights in (3.55) require.

through the parabolic evolution equation (3.28), yielding a family of value surfaces $V(t, c, d)$ and time-dependent boundaries $d = b(t, c)$ that evolve as the deadline approaches. This requires backward induction from the terminal condition (3.29), strictly respecting the direction of time (Detemple, 2006).

As $t \to T$, the option value of waiting vanishes: at the deadline, the boundary collapses to the zero-NPV contour $\hat{g}^L(c, d) = 0$, where the firm adopts immediately if and only if the project has positive intrinsic value. Far from maturity ($t \ll T$), the boundary resembles the infinite-horizon benchmark, shifted upward due to the reduced asset value discussed above.

The finite-horizon solution involves a nested computational structure: an outer loop marching backward in time, with an inner loop at each time step solving the constrained optimisation problem to convergence (Wilmott, 2006; Duffy, 2013). We discretise the time horizon into N_T intervals $\Delta t = T/N_T$. At each time step $t_n = n\Delta t$, we solve the fully implicit system:

$$\frac{V_{i,j}^{n+1} - V_{i,j}^{n}}{\Delta t} + (\mathcal{L}_h - r)V_{i,j}^{n} = 0, \tag{3.56}$$

where $V_{i,j}^n$ approximates $V(t_n, c_i, d_j)$ and $\mathcal{L}_h$ is the spatial discretisation operator from Subsection 3.4.2. The nested computational structure is illustrated in Figure 10 (right panel).

We cannot simply invert a matrix to obtain V^n from V^{n+1}, as one would in a standard parabolic PDE. The American-style constraint $V^n \geq \hat{g}^L$ introduces a complementarity condition that makes the problem nonlinear: at each grid point, the solution must satisfy either the PDE (continuation region) or the obstacle constraint (stopping region), but we do not know in advance which condition binds where (Detemple, 2006). This inequality transforms the linear system into a linear complementarity problem (LCP), which requires iterative solution (Cottle et al., 2009).

At each time step t_n, we therefore employ the Projected Successive Over-Relaxation (PSOR) method described in Subsection 3.4.2 to solve the LCP subject to the obstacle constraint $V^n \geq \hat{g}^L$ (Cryer, 1971). PSOR iterates until convergence (typically $\|V^{(k+1)} - V^{(k)}\|_\infty < \varepsilon$), effectively solving a nonlinear fixed-point problem at every time slice. Only after PSOR converges at time t_n do we proceed to the previous time step t_{n-1}, carrying the converged solution V^n as the "future value" input to the next implicit system.

The nested structure implies that the total computational cost is the product of the number of time steps and the number of PSOR iterations per step. The number of PSOR sweeps required per time slice is close to constant along the march, because the $1/\Delta t$ added to the diagonal of (3.34) is the same at every slice. Measured over the finite-horizon march, the inner iteration count runs from 119 at the first slices to 115 at $t = 0$ on the coarse 200×300 grid, and from roughly 420 to 400 on the fine grid: it falls slightly going backward rather than rising, and the cap of 5,000 inner sweeps is never approached. The cost of the finite-horizon solve is therefore governed by the number of time slices rather than by the growth of the continuation region.

We employ a fully implicit discretisation in (3.56) to ensure unconditional stability (Strikwerda, 2004). The spatial operator $\mathcal{L}_h$ contains second-order diffusion terms proportional to η^2 and σ^2, as well as a reaction term $-rV$ that can be stiff when the discount rate is moderate to high relative to the time step. An explicit scheme (evaluating $\mathcal{L}_h$ at time t_{n+1} rather than t_n) would impose a severe stability constraint $\Delta t \lesssim \Delta C^2/\eta^2$, requiring prohibitively small time steps. The fully implicit scheme removes this restriction, allowing time steps determined by accuracy rather than stability considerations. Typical choices are $N_T = 250$ steps for the calibrated planning horizon $T = 5$ years (50 steps per year), yielding $\Delta t = 0.02$ year (approximately weekly resolution).[12]

The complete backward-marching algorithm proceeds as follows:

1. Compute the payoff grid $\hat{g}^L_{i,j}$ using (3.24).
2. Initialise the terminal condition: $V^{N_T}_{i,j} = \max(\hat{g}^L_{i,j}, 0)$.
3. For each time step $n = N_T - 1$ down to 0:
    a) Formulate the implicit LCP: $(I - \Delta t(\mathcal{L}_h - r))V^n = V^{n+1}$ subject to $V^n \geq \hat{g}^L$.
    b) Solve via PSOR until $\|V^{(k+1)} - V^{(k)}\|_\infty < \varepsilon$ ($\varepsilon = 10^{-9}$; see Appendix C.4).
    c) Extract the free boundary $\{(c,d) : V^n(c,d) = \hat{g}^L(c,d)\}$ using the methods from Subsection 3.4.2.

The spatial stencils and boundary conditions remain as specified in Subsection 3.4.2; the PSOR settings of the time march ($\omega = 1.2$, $\varepsilon = 10^{-9}$) and the extended benefit range $D_{\max} = 60$ used for Scenario H are recorded in Appendix C.4. Section 4.6 presents numerical results for the delivery fleet case study with specific choices of T and L.

3.5 Numerical Validation

The two-factor model requires numerical solution via finite difference methods. This section validates the numerical implementation through grid convergence analysis, iteration convergence diagnostics, boundary extraction accuracy assessment, and Richardson extrapolation. These tests establish that discretisation errors remain

[12] The first two time steps use half-steps $\Delta t/2$ (Rannacher-style smoothing) to damp high-frequency start-up oscillations near the terminal condition before reverting to the full step size.

relatively small, confirming numerical reliability for the comparative statics and empirical calibration that follow.

Numerical solutions to free boundary problems introduce three distinct sources of error: spatial discretisation error from approximating continuous operators on a finite grid, iterative error from terminating the PSOR solver after finitely many sweeps, and boundary extraction error from locating the free boundary on a discrete mesh. The validation strategy addresses each source independently. Grid convergence analysis (Subsection 3.5.1) verifies that spatial discretisation error decreases as the mesh is refined. Iteration convergence diagnostics confirm that the PSOR residual falls below the prescribed tolerance, so that iterative error does not contaminate the spatial solution. Boundary extraction and uncertainty quantification (Subsection 3.5.2) assess the accuracy of the procedure that maps the discrete value function to a continuous boundary curve. Finally, Richardson extrapolation provides an independent, higher-order accuracy estimate by exploiting the known convergence rate of the finite difference scheme.

3.5.1 Convergence Analysis

Grid Convergence

A finite difference solution is spatially convergent if the computed quantities approach definite limits as the grid spacing $\Delta C, \Delta D \to 0$. For a second-order scheme, the discretisation error is $O(h^2)$ where h denotes the characteristic mesh width, so halving the grid spacing should reduce the error by approximately a factor of four (LeVeque, 2007; Pooley et al., 2003). We verify this property by solving the baseline infinite-horizon problem on a sequence of uniformly refined grids from 50×50 to 200×200, monitoring both a pointwise diagnostic (the value function at the current state, $V(C_0, D_0)$) and a global diagnostic (the optimal boundary threshold at C_0).

Figure 11 examines how the computed value function and optimal boundary converge as the computational grid is refined. Panel (a) displays the option value at the current state $V(C_0, D_0)$ for grid resolutions ranging from 50×50 to 200×200. The value exhibits monotonic convergence, with successive refinements producing smaller changes. The coarsest grid (50×50) yields V = \$212,605, while the finest grid (200×200) yields V = \$211,839, a difference of \$766 or 0.36%. The convergence pattern suggests that discretisation error decays systematically with grid refinement.

Panel (b) examines boundary convergence at the current cost level C_0 = \$5,800. The threshold decreases monotonically with refinement, from d^* = 8.24 \$/100mi on the 50×50 grid through 8.17, 8.13 and 8.00 to 7.96 \$/100mi on the 200×200 grid, with successive differences of 0.07, 0.04, 0.13 and 0.04 \$/100mi. Monotonicity here is a property of the extraction method: because the estimate is obtained by

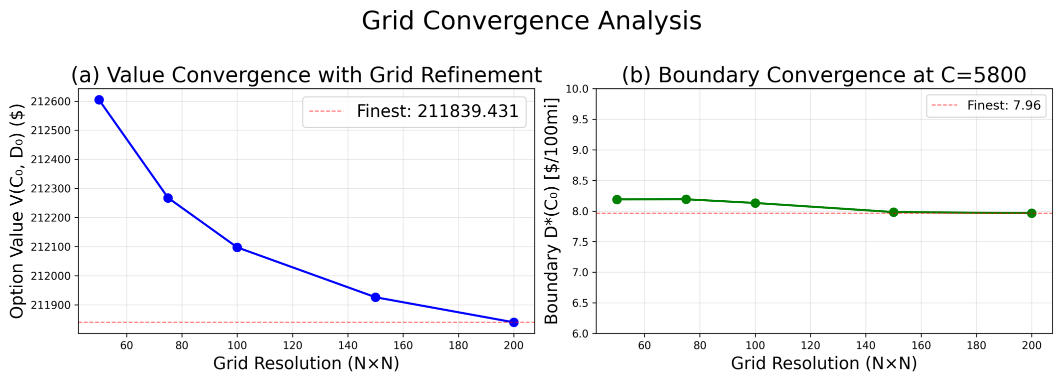

Figure 11: Grid convergence analysis for the infinite-horizon model. Panel (a) shows the option value at the current state converging as grid resolution increases. Panel (b) shows the optimal boundary threshold at C_0 = \$5,800 converging to $d^* \approx 7.96$ \$/100mi. The finest grid (200×200) serves as the reference solution.

extrapolating the quadratic decay of the continuation premium rather than by locating a level crossing, it is not quantised to the spacing of its own mesh, and successive refinements are directly comparable. A level-crossing estimate on the same solutions runs 8.32, 8.47, 7.99, 8.07, 8.11, with each entry pinned to a different grid. The convergence study uses square grids with bilinear interpolation for value extraction but without multigrid warm-start or Richardson extrapolation, isolating the effect of uniform refinement on spatial discretisation error. The production results reported in Chapter 4 employ a 400×600 grid (ΔC = \$50, ΔD = 0.10 \$/100mi), doubling the cost-dimension resolution and tripling the benefit-dimension resolution relative to the 200×200 reference.[13] Richardson extrapolation on the production grid pair (200×300 coarse, 400×600 fine) provides a direct accuracy assessment for the baseline thresholds and option values. Sensitivity analyses use a single-grid solve at 400×600 with no multigrid warm-start and no Richardson extrapolation, as the relevant quantity is the threshold difference across parameter values rather than the absolute level.

Iteration Convergence

At each grid resolution, the linear complementarity problem (3.54) is solved iteratively by Projected Successive Over-Relaxation (PSOR). The iteration is terminated when the maximum absolute change between successive iterates falls below a prescribed tolerance, $\|V^{(k+1)} - V^{(k)}\|_\infty < \varepsilon$, or when a cap of 150,000 sweeps is reached, whichever comes first. Because the value function is of order 10^5 dollars,

[13] The square-grid convergence family isolates spatial refinement and yields $d^* \approx 7.96$ \$/100mi at C_0 on its finest member. The production threshold $d^*(C_0) = 7.93$ \$/100mi reported in Chapter 4 comes from the finer and more anisotropic 400×600 pair after Richardson extrapolation, and the remaining 0.03 \$/100mi between the two is consistent with the refinement sequence above continuing toward its limit.

the criterion that matters is the relative change: the tolerance $\varepsilon = 10^{-3}$ is a relative change of about 5×10^{-9}, and every run reported here terminates on the tolerance rather than on the cap (Appendix C.4). On the 200×200 convergence member, for instance, the iteration stops after 24,118 sweeps at a maximum nodal change of 9.996×10^{-4}. Two convergence diagnostics are monitored: the residual norm (measuring how well the current iterate satisfies the complementarity system) and the evolution of economically meaningful outputs ($V(C_0, D_0)$ and $d^*(C_0)$) across iterations. The latter ensures that the quantities reported in the results have stabilised, not merely that the algebraic residual is small. Figure 12 documents the convergence behaviour of this iterative solver. Panel (a) displays the maximum nodal change across all grid points as a function of PSOR iteration count, for the 200×200 grid of the convergence family. It decays from an initial value near 10^3 to below 10^{-3} over roughly 24,000 iterations, a relative change of order 10^{-8}. The smooth decay indicates that the relaxation parameter is well-tuned and that the complementarity conditions are being satisfied consistently across the computational domain.

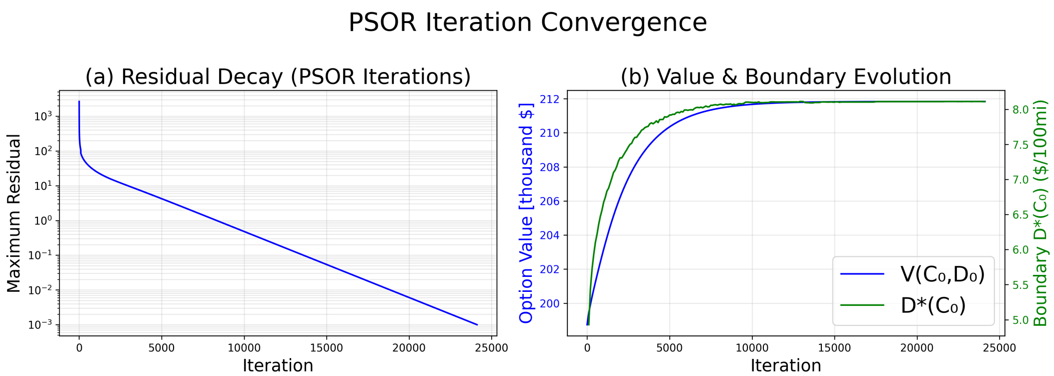

Figure 12: PSOR iteration convergence for the infinite-horizon model. Panel (a) shows exponential decay of the maximum residual to below 10^{-3} after approximately 24,000 iterations. Panel (b) tracks the evolution of the option value (blue) and boundary threshold (green) during iteration, demonstrating rapid initial adjustment followed by asymptotic refinement.

Panel (b) tracks the evolution of the option value $V(C_0, D_0)$ and boundary threshold $d^*(C_0)$ during the iterative process. Both quantities exhibit rapid initial adjustment: the option value rises from approximately \$199,000 to \$210,000 in the first 5,000 iterations, while the boundary rises from 4.93 to 7.92 \$/100mi over the same span. After this initial transient, both quantities approach their final values asymptotically, with negligible changes after iteration 15,000.[14] The concurrent con-

[14]The boundary trace in this diagnostic is produced by the level-crossing extractor rather than by the smooth-pasting fit of Subsection 3.5.2, which is why its terminal value, 8.11 \$/100mi on this grid, sits above the 7.96 \$/100mi reported for the same solve in Subsection 3.5.1. The panel is a convergence diagnostic and the offset is the quantisation discussed there; no reported threshold is taken from it.

vergence of value and boundary confirms that the iterative scheme is simultaneously satisfying the value function equation and the free boundary conditions.

3.5.2 Boundary Extraction and Uncertainty Quantification

The optimal boundary $d^*(c)$ is defined as the locus of states where the decision-maker is indifferent between immediate adoption and continued waiting: $V(c, d^*(c)) = \hat{g}(c, d^*(c))$. On the discrete grid this contour is not resolved exactly, and how it is extracted matters more than the grid resolution does. The projection step enforces $V = \max\{\hat{g}, 0\}$ exactly in the stopping region, so the continuation premium $V - \hat{g}$ is identically zero above the boundary rather than merely small. Locating the boundary as the last node whose premium exceeds a threshold therefore returns a grid node exactly, whatever the threshold: the estimate is quantised to ΔD, and refining the mesh moves it one node at a time instead of converging.[15]

Smooth pasting supplies the remedy. Since V meets $\hat{g}$ tangentially, the premium vanishes quadratically as the boundary is approached, $V - \hat{g} \sim k\,(d^*(c) - d)^2$, so $\sqrt{V - \hat{g}}$ is linear in d over a neighbourhood below the boundary and its root locates $d^*(c)$ without reference to any threshold. For each cost slice C_i the premium is windowed to the band where the quadratic regime holds, $\sqrt{V - \hat{g}}$ is regressed on d over that band, and the root of the fitted line is taken as the raw boundary $d_{\text{raw}}(c)$; the relative residual of the fit, around one per cent at the production resolution, doubles as a diagnostic that the quadratic regime was reached. Isotonic regression then enforces monotonicity and a smoothing spline with smoothing parameter s equal to the number of valid boundary points (or 1.5 times that number in the finite-horizon extractor) produces the final reported boundary $d^*(c)$.

Halving ΔD from 0.30 to 0.15 moves the smooth-pasting estimate at C_0 by 0.026 \$/100mi, against 0.107 for a level-crossing estimate on the same solutions, and across the whole cost range the mean discrepancy between the two resolutions is 0.061 \$/100mi, against 0.109. Every threshold value quoted in this dissertation comes from the smooth-pasting extractor.[16]

Figure 13 examines the accuracy of this extraction procedure on the production 400×600 grid. Panel (a) illustrates the computational grid, the raw extracted boundary (blue dashed), and the smoothed boundary (solid red).

Panel (b) quantifies extraction uncertainty through a corridor combining two components in quadrature: grid quantisation uncertainty ($w_{\text{grid}} = \Delta D/2$) and smoothing model uncertainty ($w_{\text{model}} = |d^* - d_{\text{raw}}|$). A gradient-based term

[15] Interpolating between that node and the one above it does not help, because the premium there is exactly zero, so the interpolant collapses onto the node and a root-finder has no sign change to bracket.

[16] A residual column-to-column ripple survives the fit, because the $O(h^2)$ truncation error of the PSOR solution has a grid-frequency component that perturbs the fitted premium slightly differently in adjacent cost columns. It is an order of magnitude smaller than the grid spacing ΔD, and the isotonic regression and smoothing spline remove what remains.

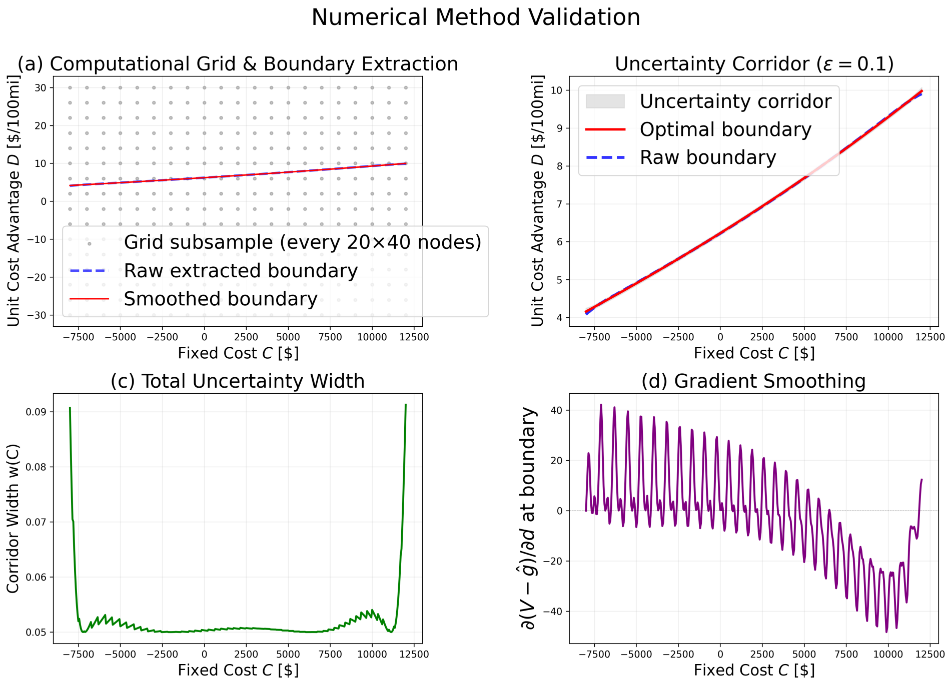

Figure 13: Numerical method validation for boundary extraction on the production 400×600 grid. Panel (a) shows the computational grid, the raw smooth-pasting fit, and the smoothed boundary. Panel (b) displays the uncertainty corridor (grey shading) around the optimal boundary, with the two uncertainty components combined in quadrature. Panel (c) shows the corridor half-width $w(C)$ (the corridor is $d^* \pm w$), essentially constant at 0.05 \$/100mi across the central cost range and at most 0.09 \$/100mi at the domain endpoints. Panel (d) displays the gradient difference $\partial(V - \hat{g})/\partial d$ along the boundary.

$w_{\text{grad}} = \epsilon/|\partial(V - \hat{g})/\partial d|$ is not included in the corridor: it is evaluated at the boundary, where smooth pasting makes the derivative vanish, so it diverges precisely when the boundary has been located accurately. It is reported as a diagnostic instead. Panel (c) displays the total uncertainty width across the cost range. In the central region ($-2{,}500 \lesssim C \lesssim 7{,}500$), the width is essentially constant at 0.05 \$/100mi, rising to at most 0.09 \$/100mi at the domain endpoints. It is dominated by the grid term $\Delta D/2$, with w_{model} contributing a median of 0.008 \$/100mi, which reflects a raw fitted boundary that varies smoothly between adjacent cost columns rather than jumping by a full grid cell.

Panel (d) examines the smooth-pasting condition by plotting $\partial(V - \hat{g})/\partial d$ along the boundary. Optimality requires this gradient difference to vanish at $d^*(c)$. The computed gradients exhibit oscillations with amplitude approximately -50 to $+100$, reflecting the same grid-frequency discretisation error. Smooth pasting makes V continuously differentiable across the boundary, so the discontinuity there is one of curvature rather than of gradient, and the residual plotted in Panel (d) measures how far the extracted boundary sits from the true indifference contour rather than any failure of the value function itself. It is therefore a direct read-out of boundary-extraction error, and it is the component the uncertainty corridor of Subsection 3.5.2 is constructed to quantify.

3.5.3 Richardson Extrapolation

Richardson extrapolation (3.55) cancels the leading $O(h^2)$ term of the discretisation error by combining coarse- and fine-grid solutions. The scheme is not uniformly second order: the upwind treatment of the drift contributes an $O(h)$ term that the weights $(4, -1)/3$ do not cancel. At the baseline calibration that term is small, because the artificial diffusion it introduces, $|\mu_C|\Delta C/2$ in the cost direction and $\alpha\Delta D/2$ in the benefit direction, amounts to 1.4% and 0.6% of the corresponding physical diffusion coefficients on the production grid $\eta^2/2$ and $\sigma^2/2$. Over the range of resolutions tested the observed error therefore behaves as $O(h^2)$ and the extrapolation is informative, but the residual is estimated empirically from the grid sequence rather than claimed to be of any particular higher order. The extrapolation is applied to the raw value function arrays; the optimal boundary is then extracted from the extrapolated surface using the procedure described in Subsection 3.5.2. Figure 14 applies this technique to the production grid pair: coarse (200×300) and fine (400×600).[17]

Panel (a) compares the boundaries extracted from all three solutions. Both the coarse and fine boundaries exhibit the grid-frequency oscillations described in Subsection 3.5.2, while the Richardson-extrapolated boundary is smooth. The cancellation of grid-scale noise is a direct consequence of the extrapolation: the $O(h^2)$ truncation error that drives the oscillations is eliminated by the linear combination, leaving a boundary dominated by the higher-order remainder. Panel (b) shows that the fine-grid boundary sits approximately 0.008 \$/100mi above the Richardson estimate at C_0. Panel (c) displays value function slices at $C = \$5{,}800$ across the benefit dimension; the three solutions are visually indistinguishable. Panel (d) quantifies value errors: the fine grid deviates by at most \$120 from the Richardson estimate, while the coarse grid peaks at approximately \$500 near $d \approx -10$ \$/100mi, with a sawtooth pattern reflecting interpolation of the coarse solution onto the fine grid. At this grid resolution, the Richardson correction shifts the baseline threshold by 0.008 \$/100mi and the value function by \$25 (0.012%), the figures recorded in Table C.5. Both are an order of magnitude below the width of the boundary uncertainty corridor of Subsection 3.5.2, which is the appropriate error bar on d^* and which the following subsection quotes.

[17] The cancellation of the $O(h^2)$ term is exact at grid points shared by both meshes. At intermediate fine-grid nodes, the coarse-grid solution is obtained by linear interpolation, which introduces an $O(h^2)$ interpolation error that does not share the structure of the PDE truncation error. The extrapolated solution at these points therefore retains a small $O(h^2)$ component. In practice, the interpolation error is comparable in magnitude to the discretisation error it replaces, so the net improvement remains meaningful.

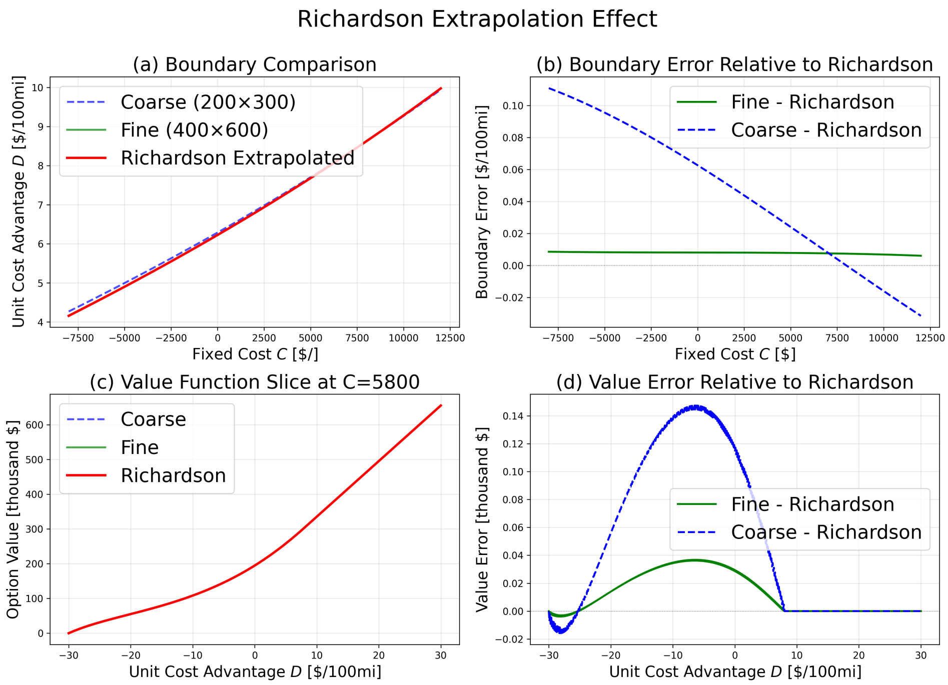

Figure 14: Richardson extrapolation analysis on the production grid pair (200×300 coarse, 400×600 fine). Panel (a) compares boundary estimates from coarse, fine, and Richardson-extrapolated grids. Panel (b) quantifies boundary errors relative to the Richardson estimate. Panel (c) displays value function slices at $C = \$5{,}800$ for the three resolutions. Panel (d) quantifies value function errors relative to the Richardson estimate.

3.5.4 Summary of Numerical Reliability

The validation exercises establish that the reported numerical results achieve acceptable accuracy for economic analysis. On the production grid (400×600) the Richardson correction to the baseline threshold is 0.008 \$/100mi and to the value function \$25, and the boundary uncertainty corridor, which additionally absorbs the grid quantisation and the residual smoothing spread, has a half-width of 0.05 \$/100mi across the central cost range. The corridor is the appropriate error bar on a reported threshold, and it is the figure quoted here. These error magnitudes remain an order of magnitude smaller than the economic effects examined in the sensitivity analysis, where parameter variation produces threshold changes of 0.4 to 19 \$/100mi. The numerical uncertainty is also small relative to the gap between the current state and the boundary (6.5 \$/100mi) for the baseline calibration with $C_0 = \$5{,}800$ and $D_0 = 1.45$ \$/100mi, confirming that the qualitative adoption decision is not sensitive to discretisation choices for the parameter ranges examined. Richardson extrapolation confirms that the production grid is well into the asymptotic convergence regime, with no evidence of numerical instability or spurious boundary behaviour.

Chapter 4

NUMERICAL RESULTS AND COMPARISON

This chapter presents numerical analysis of the two-factor optimal stopping model applied to the delivery fleet replacement problem by extending the calibration from Chapter 2 to include fixed cost dynamics. Section 4.1 estimates the Ornstein–Uhlenbeck parameters from UK vehicle price data to demonstrate the estimation methodology. Sections 4.2–4.6 present the optimal switching boundaries under deterministic cost decay, stochastic mean-reverting costs, and finite planning horizons, with sensitivity analysis across key parameters.

4.1 Calibrating the Stochastic Cost Process

The two-factor optimal stopping framework developed in Section 3.1 of Chapter 3 requires concrete values for the fixed cost ($C(t)$) dynamics parameters (κ, θ, η). As reviewed in Subsection 1.2.3, EV–ICE price convergence operates on timescales of roughly 2.5 to 6 years with segment-specific learning rates. This section translates those cross-sectional findings into time-series parameter estimates suitable for the Ornstein–Uhlenbeck specification. We estimate (κ, θ, η) from UK vehicle price data for 2018–2024 to demonstrate the methodology; practical applications require decision-specific calibration reflecting the relevant market, vehicle segment, and time period. We validate the implied half-life consistency against the learning-curve benchmarks reviewed in Chapter 1 and demonstrate that the OU specification successfully captures non-monotonic price dynamics observed during this period.

The calibration uses UK market data for a representative C-segment vehicle pair: the Nissan Leaf (battery electric) and Ford Focus (petrol). The C-segment represents 22% of EU new car sales (ACEA, 2025) and permits like-for-like comparison between vehicles with comparable size, features, and target demographics.[1]

[1] The VW e-Golf versus Golf comparison offers the cleanest platform match but ends in 2020 when VW discontinued the e-Golf in favour of the ID.3. The Tesla Model 3 versus BMW 3-Series

The pairing reflects practical data availability; while current prices are publicly observable, historical price series for specific models are rarely available. Comcar [2], a UK company car tax calculator, maintains historical specifications and list prices for Benefit-in-Kind calculations accessible through its web interface. This single-pair analysis is illustrative rather than definitive. The methodology applies equally to other vehicle comparisons or segment-level indices, and the specific example demonstrates feasibility and delivers parameters directly usable in the two-factor numerical analysis.

The theoretical object of interest is $C(t) = P_{\mathrm{EV}}(t) - P_{\mathrm{ICE}}(t)$: the additional upfront cost a household pays to choose the electric option. We use the Nissan Leaf as the EV representative and the Ford Focus as the ICE comparator. Both occupy the C-segment (compact family cars), both have been continuously available in the UK market throughout the period covered by the price histories (2011–2024), and both represent mainstream rather than premium positioning within their respective powertrains.

Price data are derived from On-The-Road (OTR) list prices collected from Comcar, converted to constant 2022 USD using FRED exchange rates and CPI adjustments. Minimum list prices are used for each model to maintain comparability across years; where multiple concurrent listings share a revision date, their minimum prices are averaged. The Ford Focus transitioned to mild-hybrid variants in 2022, which are treated as appropriate ICE comparators since they run exclusively on petrol and represent evolutionary refinement rather than fundamental powertrain substitution.

Figure 15 displays the price histories. The Leaf exhibits learning-curve-driven decline from \$60,000 (2011) to \$33,000 (2022),[3] while the Focus shows modest price increases reflecting emissions compliance costs. Figure 16 plots the EV–ICE price differential $C(t)$, which falls from \$17,000 (2018) to \$640 (late 2022), then widens to \$4,300 by spring 2024 as the repricing associated with the Leaf's 2025 relaunch takes effect.[4] This non-monotonic pattern (rapid convergence followed

comparison captures the premium end of the segment but introduces brand-positioning complications. The Leaf–Focus pairing balances data availability, segment consistency, and temporal coverage.

[2] `https://comcar.co.uk`

[3] Cumulative global Leaf sales grew from approximately 22,000 units in 2011 to approximately 577,000 units by early 2022, reaching 650,000 by mid-2023. Sources: Nissan Motor Corporation, "Renault-Nissan Alliance Posts Record Sales in 2011 for Third Consecutive Year," press release, February 1, 2012, `https://usa.nissannews.com/en-US/releases/renault-nissan-alliance-posts-record-sales-in-2011-for-third-consecutive-year`; Nissan Europe, "Nissan LEAF Gets a New Glow for 2022," press release, February 23, 2022; Nissan Motor Corporation, "Nissan Global EV Sales Surpass 1-Million-Unit Milestone," press release, July 25, 2023, `https://global.nissannews.com/en/releases/nissan-global-ev-sales-surpass-1-million-unit-milestone`. This production growth is essential for calculating learning rates, which quantify cost reductions as a function of cumulative output.

[4] The final two evaluation dates interpolate the Leaf price between its November 2022 and November 2025 revisions. Reconstructing the latent path at interior dates from the full sample is appropriate for ex-post calibration; no real-time forecasting claim is made. The November 2025 posted prices imply a differential above \$9,000, confirming that the reversal is a feature of the data rather than an artefact of the interpolation.

by partial reversal) motivates the stochastic OU specification over deterministic learning curves. Since the two models do not update prices on the same dates, the price differential $C(t_i)$ is evaluated at each date on which either model records a price revision. Posted OTR prices are staircase functions of time: administrative snapshots of an underlying cost differential that evolves continuously with battery costs, exchange rates, and competitive conditions. The calibration targets this latent differential rather than the posted staircase. At each evaluation date the revising model contributes its freshly posted price, while the other model's price is estimated by linear interpolation between its two adjacent revisions. For driftless Brownian motion the linear interpolant is exactly the conditional expectation of the path given the bracketing observations. For the Ornstein–Uhlenbeck process estimated here it is an approximation to it: the OU bridge weights the endpoints by sinh terms rather than linearly, though at the estimated κ and the observed spacings ($\kappa\Delta_i \leq 0.25$) the two differ by about one per cent. The linear reconstruction is therefore used as a minimum-assumption approximation rather than an exact conditional mean; it also values both vehicles at a common date, whereas differencing a fresh price against one posted up to a year earlier would mix price vintages and attribute the full effect of each revision to a single date. The construction yields 15 irregularly spaced observations spanning mid-2018 to spring 2024, with inter-observation intervals Δ_i ranging from several weeks to approximately one year. The exact-discrete OU transition density used in the log-likelihood accommodates unequal spacing through observation-specific conditional variances $v_i = \eta^2(1 - e^{-2\kappa\Delta_i})/(2\kappa)$; short-interval observations receive proportionally less weight, so the irregular sampling does not introduce bias. Because interpolation smooths the latent path, increments over intervals with an interpolated endpoint understate its variability; the estimates are therefore quasi-maximum-likelihood, and $\hat{\eta}$ should be read as a conservative (lower) estimate of latent cost volatility. For the same reason the standard errors below, which come from the numerical Hessian of the log-likelihood, are not consistent under this misspecification and should be read as optimistic; a sandwich estimator would widen them.[5] Figure 16 distinguishes dates where both models have freshly posted prices from those where one price is interpolated.

We estimate the OU process

$$dC(t) = \kappa(\theta - C(t))\,dt + \eta\,dW(t)$$

by maximum likelihood, imposing $\theta = 0$ as the long-run mean. The parity constraint reflects industry forecasts of price convergence by the late 2020s (BloombergNEF,

[5] A last-observation-carried-forward construction, which uses posted prices only, yields an essentially unchanged $\hat{\kappa} \approx 0.25$ but a higher $\hat{\eta} \approx \$3{,}600/\sqrt{\text{year}}$, consistent with attributing each full price revision to a single date. Given the modest and near-additive effect of η on the optimal boundary (Section 4.5), this alternative construction would raise adoption thresholds by a few percent without altering the waiting prescription.

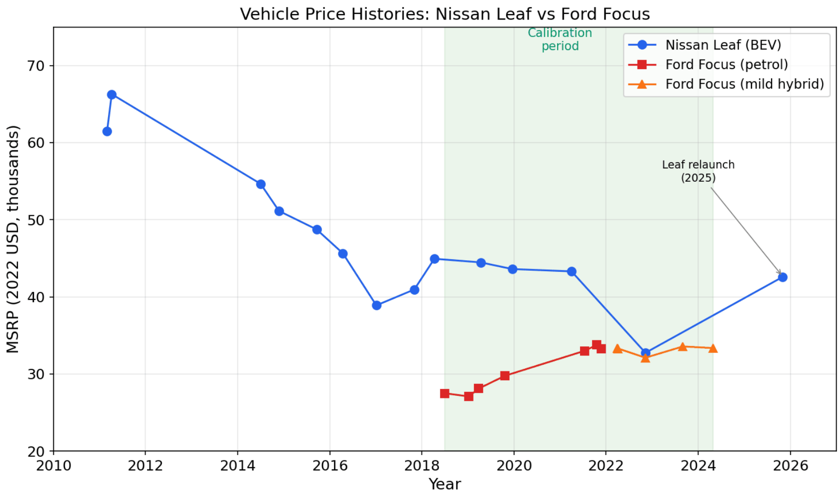

Figure 15: Price histories for Nissan Leaf and Ford Focus in constant 2022 USD, derived from On-The-Road (OTR) list prices collected from Comcar (`https://comcar.co.uk`). Minimum list prices are used for each model to maintain comparability across years; where multiple concurrent listings share a revision date, their minimum prices are averaged. The shaded region indicates the calibration period (2018–2024).

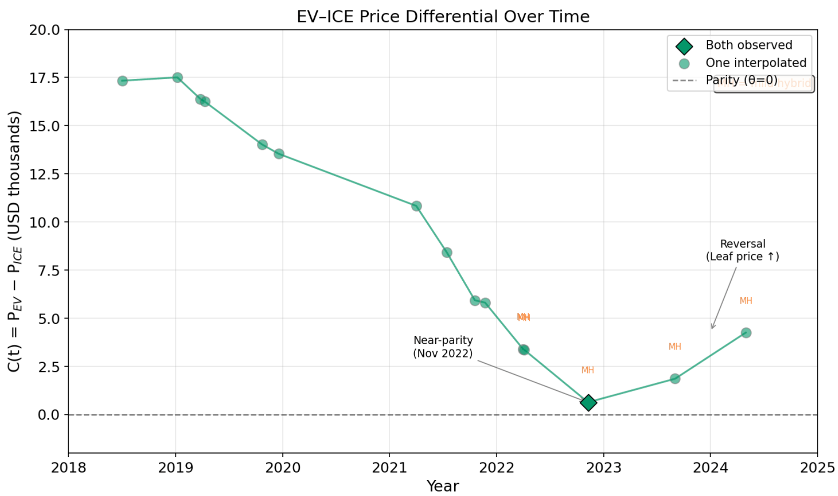

Figure 16: EV–ICE price differential $C(t)$ over time. Non-monotonic dynamics motivate the OU specification.

2025). The OU framework accommodates alternative equilibria should the data suggest otherwise.

For the OU process with observations $\{C_{t_i}\}_{i=1}^n$ at times $\{t_i\}$, the transition density is Gaussian:

$$C_{t_{i+1}} \mid C_{t_i} \sim \mathcal{N}\left(C_{t_i} e^{-\kappa \Delta_i}, \frac{\eta^2}{2\kappa}(1 - e^{-2\kappa\Delta_i})\right),$$

where $\Delta_i = t_{i+1} - t_i$. The log-likelihood follows:

$$\ell(\kappa, \eta) = -\frac{n-1}{2}\log(2\pi) - \frac{1}{2}\sum_{i=1}^{n-1}\log v_i - \frac{1}{2}\sum_{i=1}^{n-1}\frac{(C_{t_{i+1}} - C_{t_i}e^{-\kappa\Delta_i})^2}{v_i},$$

where $v_i = \eta^2(1 - e^{-2\kappa\Delta_i})/(2\kappa)$ is the conditional variance. We maximise numerically over $(\kappa, \eta) \in (0, \infty)^2$.

The calibration yields (asymptotic standard errors from the numerical Hessian of the negative log-likelihood):

$$\hat{\kappa} = 0.25 \text{ per year} \quad (\text{s.e.} \approx 0.09),$$
$$\hat{\eta} = \$2{,}320/\sqrt{\text{year}} \quad (\text{s.e.} \approx \$440).$$

The mean-reversion rate implies a half-life of $\ln(2)/\kappa \approx 2.8$ years. The volatility parameter implies a stationary standard deviation of $\eta/\sqrt{2\kappa} \approx \$3{,}300$; the observed swing from \$640 to \$4,300 represents a one-standard-deviation fluctuation.

Table 7 benchmarks the estimated 2.8-year half-life against learning-curve studies reviewed in Subsection 1.2.3. Despite methodological differences, all estimates fall within 2.5–6.0 years. The time-series approach captures total convergence (both EV cost reductions and ICE price increases) while learning curves typically isolate supply-side cost dynamics. This distinction likely explains why the time-series estimate lies at the faster end of the range.

Study	Segment(Region)	Half-life [years]
This study	C (UK)	2.8
Goetzel and Hasanuzzaman (2022)	C/D (Germany)	3.1–6.0*
Goetzel and Hasanuzzaman (2022)	E/F (Germany)	4.0
Weiss et al. (2012)	All (Global)	4.3
Weiss et al. (2019)	All (Germany)	2.6

*Range reflects uncertainty in production growth rate assumptions.

Table 7: Half-life estimates for EV–ICE price convergence across studies. This study uses OU time-series maximum-likelihood estimation; other estimates use learning-curve regressions.

Figure 17 overlays the calibrated model on observed data. All observations fall within the 95% confidence bands. The post-2022 widening to \$4,300 represents a one-standard-deviation event under the model. Figure 18 displays the evolution of the probability distribution for $C(t)$; by 2032, substantial probability mass lies below zero.

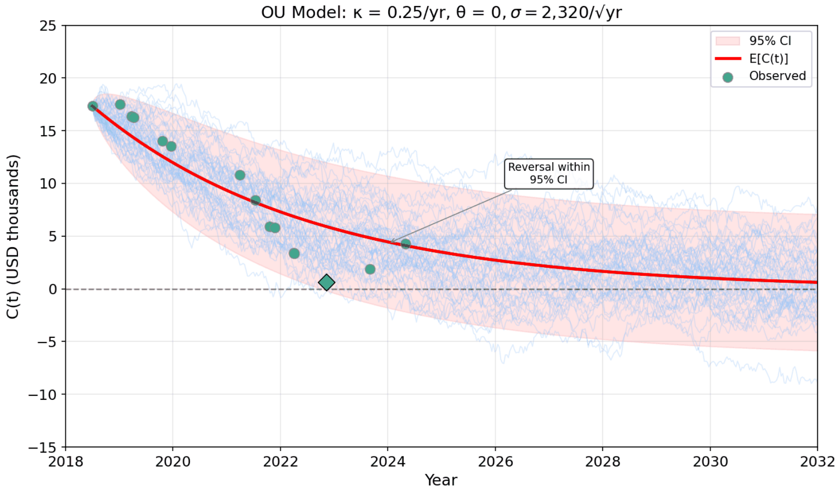

Figure 17: Calibrated OU model with expected path and 95% confidence interval. All observations fall within the confidence bands.

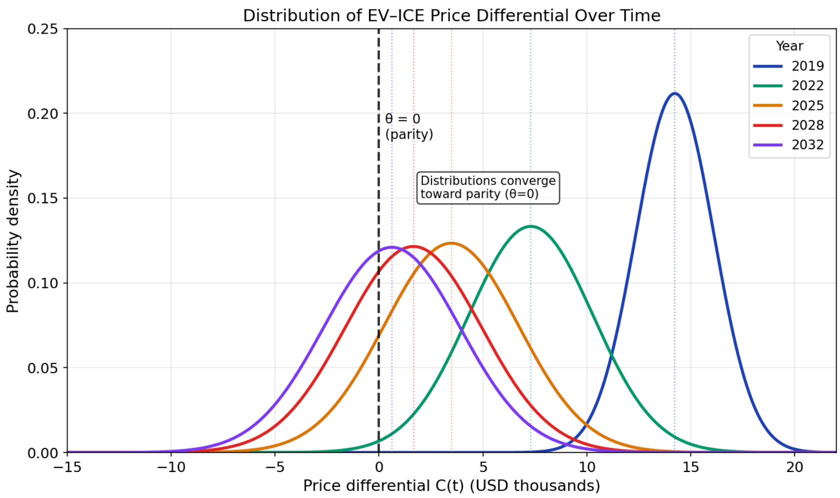

Figure 18: Evolution of the probability distribution for the price differential. By 2032, substantial probability mass lies below zero.

The calibration warrants appropriate caution. The fifteen observations support two-parameter estimation with meaningful precision, but standard errors are substantial. The Leaf–Focus comparison reflects dynamics specific to these models rather than segment averages, and the 2018–2024 sample period includes extraordinary supply chain volatility that likely inflates η relative to steady-state conditions. The methodology is portable to other vehicle pairs or segment-level indices, but the single-pair analysis demonstrates feasibility rather than providing definitive estimates.

This section has estimated the cost dynamics parameters: $\kappa = 0.25$ per year (half-life: 2.8 years), $\eta = \$2{,}320/\sqrt{\text{year}}$, and $\theta = 0$. Validation against learning-curve evidence confirms consistency across methodologies. These parameters are combined with the unit cost advantage estimates from Chapter 2 in Section 4.2, where the cross-market parameter synthesis and its justification are discussed.

4.2 Case Study: Application to the Delivery Fleet Case

Chapter 2 analysed the delivery fleet's adoption decision across five scenarios (A-E) under constant and jump-cost specifications. These scenarios varied by location (San Francisco, Seattle, Miami) and policy expectations (tax credit removal, technological breakthroughs). The present section extends this analysis by examining three additional scenarios (F-H) that incorporate stochastic cost dynamics calibrated from UK vehicle price data in Section 4.1.

In the constant-cost framework of Chapter 2, the optimal threshold was $x^{**}(C_0)$, a linear function of the fixed cost. The jump-cost extension yields a linear adoption boundary $d^*(C)$ whose intercept shifts with jump risk, and the deterministic-decline threshold of Section 3.2 depends on the current benefit level through the Lambert-W argument. For stochastic mean-reverting costs, the optimal stopping boundary becomes a two-dimensional curve $d = b(c)$ in the (C, D) state space, and the initial state $(C_0, D_0) = (5{,}800,\ 1.45)$ determines the fleet's distance from this boundary.

Table 8 reports the complete parameter set, integrating cost advantage parameters (α, σ, D_0, λ) from Subsection 2.4.2 in Chapter 2 with cost process parameters (κ, θ, η) from Section 4.1. The parameters are drawn from different empirical contexts, reflecting both data availability constraints and the economic specificity of the underlying mechanisms. Benefit parameters (α, σ, D_0, λ) depend on local energy prices, taxation, and usage intensity and are calibrated from US energy market data and US delivery company operations. Cost process parameters (κ, θ, η) reflect global technological forces (learning-by-doing in battery production and platform standardisation) and are estimated from specific UK EV–ICE price convergence patterns. This parameter synthesis is adequate for demonstrating the framework, though each application should calibrate using data appropriate to its specific context.

The parameter combination warrants explicit justification. Cost dynamics parameters (κ, θ, η) are calibrated from UK Leaf–Focus price differentials over 2018–2024, while the analysis applies to a US delivery company with initial cost differential $C_0 = \$5,800$ derived from US market Leaf versus Sentra pricing (net of Federal Tax Credit). This cross-market and cross-model parameter transfer assumes that purchase price convergence dynamics are determined primarily by global technology

Parameter	Value	Source
unit cost advantage process parameters (fixed across all models)		
α	0.34 \$/100mi per year	US energy prices (2015–2023)[a]
σ	2.42 \$/100mi per $\sqrt{\text{year}}$	US energy prices (2015–2023)[a]
D_0	1.45 \$/100mi	SF energy prices, Dec 2024[b]
λ	480 hundred-miles/year	Delivery company
Fixed cost process parameters (Scenarios F–H)		
κ	0.25 per year	UK EV–ICE price data[c]
θ	0 \$	Long-run parity assumption[c]
η	2,320 \$/$\sqrt{\text{year}}$	UK EV–ICE price data[c]
C_0	5,800 \$	US EV–ICE price difference[d]
Structural and scenario-specific		
r	3% per year	Standard assumption
ρ	-0.3 (baseline)	Economic reasoning[e]
ν	$\{1, \frac{1}{8}\}$ per year	Economic reasoning for scenarios[f]
d_e	$\{-3{,}325, 7{,}500\}$ \$	Economic reasoning for scenarios[f]
T	5 years	Planning horizon assumption
L	10 years	Vehicle operational lifetime assumption[g]

Table 8: Model parameters used in the comparison and their empirical sources. [a]Historical US oil and electricity prices (2015–2023); see Subsection 2.4.2 in Chapter 2. [b]Manufacturer efficiency specifications and San Francisco energy prices (December 2024, BLS). [c]UK EV–ICE price convergence analysis; see Section 4.1. [d]Leaf vs Sentra price difference (net of Federal Tax Credit); see Chapter 2. [e]Baseline correlation; sensitivity to $\rho \in \{-0.3, 0, +0.3\}$ examined in Section 4.5. [f]Scenario C: projected EV–ICE price parity by 2029. Scenario E: 50% probability of federal tax credit removal within 4 years; see Chapter 2. [g]Derived from average vehicle age in the U.K.; see Scenario H, Section 4.2.

factors rather than model-specific or regional features. Battery cost reductions from manufacturing scale and learning-by-doing operate at the industry level, affecting all C-segment vehicles similarly across developed markets. Regional differences enter through the initial premium C_0, which reflects market-specific pricing and incentives, while the convergence rate κ and volatility η capture segment-level dynamics transferable across markets and vehicle pairs within the same segment. In contrast, benefit parameters are inherently location- and usage-specific: fuel cost advantages depend on local energy prices, while utilisation intensity λ reflects the delivery company's operational characteristics. The calibration thus separates global cost convergence dynamics (transferable across C-segment vehicles and markets) from local benefit dynamics (application-specific). This separation is the appropriate structure for a framework demonstration: the mechanisms are identified from the best available data for each component, and the methodology is portable to any context where decision-specific recalibration is feasible.

We examine three additional cases:

- Scenario F: Fleet replacement assuming deterministic cost decline. The fixed cost follows the deterministic mean trajectory $C(t) = C_0 e^{-\kappa t}$, the conditional expectation of the Ornstein–Uhlenbeck process with $\theta = 0$, abstracting from stochastic volatility to isolate the effect of expected cost convergence on adoption timing.
- Scenario G: Fleet replacement under stochastic mean-reverting costs over an infinite planning horizon. The fixed cost follows the full Ornstein–Uhlenbeck process, capturing persistent price volatility driven by supply chain disruptions and policy uncertainty alongside mean reversion toward long-run parity.
- Scenario H: Fleet replacement under stochastic mean-reverting costs with finite planning horizon $T = 5$ years. The planning horizon reflects typical commercial fleet replacement cycles, and the operational lifetime $L = 10$ years is consistent with the average vehicle age of 9–10 years recorded in the UK by 2024 (SMMT motor parc data; UK Department for Transport vehicle licensing statistics[6]). The model incorporates terminal conditions reflecting these depreciation schedules and institutional constraints that prevent indefinite deferral of replacement decisions.

Each scenario is analysed to characterise the optimal stopping boundary and isolate the incremental effects of cost stochasticity, mean reversion, and horizon truncation relative to the constant-cost benchmark from Chapter 2.

4.3 Cross-Scenario Comparison and Model Specification Effects

The optimal adoption threshold values in Table 9 span more than a fivefold range (from 4.68 \$/100mi under anticipated subsidy withdrawal to 26.91 \$/100mi under finite horizon) reflecting different economic trade-offs across model specifications and scenario expectations. This section quantifies the contribution of four modelling dimensions to adoption timing: anticipated jumps in the adoption cost, expected cost convergence, cost volatility, and planning horizon. By comparing models that differ in only one dimension at a time, we isolate the economic mechanisms driving threshold variation and assess the robustness of adoption decisions to model specification.

Figure 19 visualises the optimal adoption boundaries across all model specifications. The constant-cost benchmark exhibits low cost sensitivity, as the fixed

[6] `https://www.gov.uk/government/statistics/vehicle-licensing-statistics-2024/`

Scenario	Description	Threshold ($/100mi)	Decision
Chapter 2: Constant and Jump Costs			
A	San Francisco, constant cost	6.08	Wait
B	Seattle, constant cost	6.08	Adopt
C	Seattle, negative jump ($d_e = -3{,}325$)	11.66	Wait
D	Miami, constant cost	6.08	Wait
E	Miami, positive jump ($d_e = 7{,}500$)	4.68	Adopt
Current Chapter: Stochastic Mean-Reverting Costs			
F	Deterministic decline ($\eta = 0$)	7.31	Wait
G	Stochastic OU, infinite horizon	7.93	Wait
H	Stochastic OU, finite horizon (T=5)	26.91	Wait

Table 9: Optimal thresholds and adoption decisions across all scenarios. All scenarios share identical usage benefit process parameters ($\alpha = 0.34$ \$/100mi/year, $\sigma = 2.42$ \$/100mi/$\sqrt{\text{year}}$, $\lambda = 480$ hundred-miles/year, $r = 3\%$/year) and current fixed cost ($C_0 = 5{,}800$ \$). Scenarios A-E from Chapter 2 vary initial benefit D_0 by location (San Francisco: 1.45, Seattle: 7.94, Miami: 4.80 \$/100mi); for the jump scenarios the reported thresholds are the adoption boundaries $d^*(C_0)$ evaluated at C_0: 11.66 \$/100mi in Scenario C, from Theorem 2.11, and 4.68 \$/100mi in Scenario E, from Proposition B.9, where the cost jump is upward. Scenario F uses deterministic cost decline ($\kappa = 0.25$/year, $\theta = 0$, $\eta = 0$); Scenarios G–H analyse stochastic mean-reverting costs ($\kappa = 0.25$/year, $\theta = 0$, $\eta = 2{,}320$ \$/$\sqrt{\text{year}}$) at the baseline state $D_0 = 1.45$ \$/100mi. Scenarios A–G use perpetual payoff; Scenario H uses planning horizon $T = 5$ years, vehicle operating lifetime $L = 10$ years, and annuity payoff.

cost enters only as a parameter in the perpetuity valuation. The deterministic decline boundary exhibits moderate positive slope, reflecting the incentive to wait for more favourable cost conditions. The stochastic infinite-horizon boundary lies above the deterministic case, with cost uncertainty creating additional option value that increases with the current cost level. The finite-horizon boundary shows the steepest positive slope, reflecting the annuity penalty: limited operational lifetime amplifies the importance of the upfront cost differential. The current state $(C_0, D_0) = (5{,}800, 1.45)$ lies in the continuation region across all specifications, prescribing waiting.

4.3.1 Jumps versus Constant Cost: Policy Anticipation

Comparing the jump boundaries of Scenarios C and E with the constant-cost benchmark (Scenarios A, B, D, $x^{**} = 6.08$ \$/100mi) isolates the effect of anticipated discrete shifts in the adoption cost. All three specifications share identical benefit dynamics and the same current cost $C_0 = 5{,}800$, differing only in whether the decision-maker attaches a Poisson hazard to a future change in C. This is the only

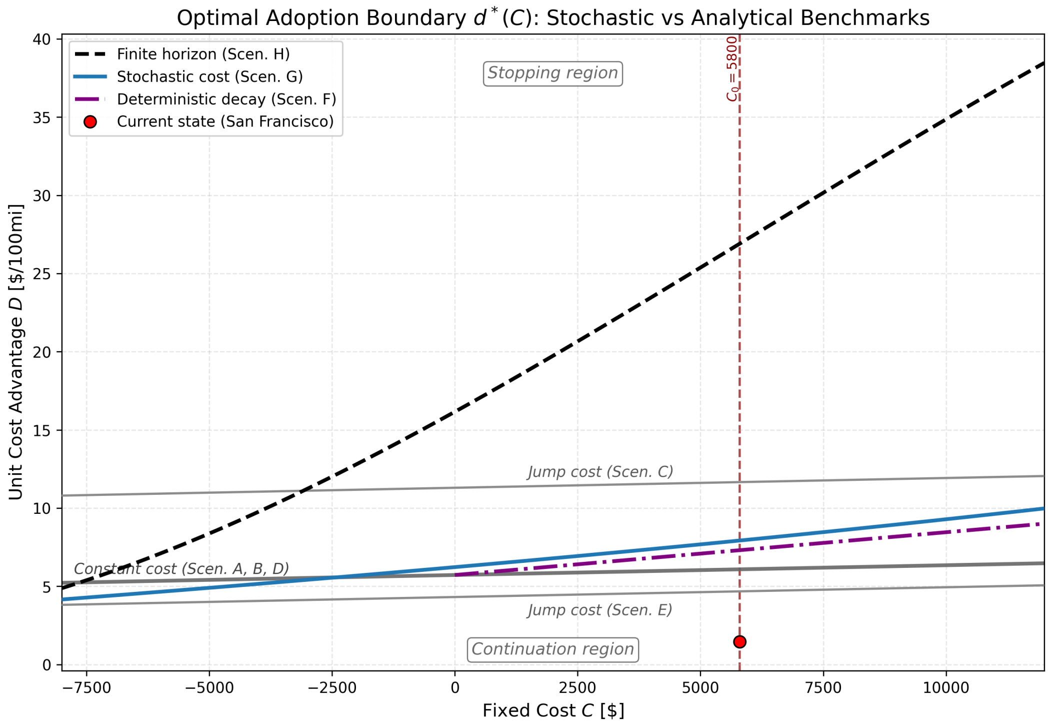

Figure 19: Comparison of optimal adoption boundaries across model and scenario specifications. The constant-cost benchmark (solid grey line, Scenarios A, B and D) exhibits a shallow positive slope. The jump-cost boundaries (thin grey lines, Scenarios C and E) flank the constant-cost benchmark: Scenario C ($d^*(C_0) = 11.66$) lies above, reflecting anticipated cost reductions that strengthen the incentive to wait, while Scenario E ($d^*(C_0) = 4.68$) lies below, reflecting the threat of subsidy removal that accelerates adoption. The deterministic cost decline boundary (purple dash-dot, Scenario F) exhibits a steeper positive slope, reflecting the incentive to wait for more favourable cost conditions. The stochastic infinite-horizon boundary (solid blue, Scenario G) lies above the deterministic case, with cost uncertainty creating additional option value. The finite-horizon boundary (dashed black, Scenario H) shows the steepest positive slope, reflecting the annuity penalty where limited operational lifetime amplifies cost sensitivity. The current state $(C_0, D_0) = (5{,}800, 1.45)$ is marked with a red dot.

one of the four dimensions whose effect is signed: anticipated cost reductions (Scenario C, a stream of technological breakthroughs with mean $|d_e| = 3{,}325$) raise the boundary to $d^*(C_0) = 11.66$ \$/100mi, a 92% increase, while the anticipated withdrawal of the \$7,500 federal credit (Scenario E) lowers it to 4.68 \$/100mi, a 23% decrease.

The two effects differ in magnitude (+5.58 against −1.40 \$/100mi). The individual jump is larger in Scenario E (\$7,500 against a mean of \$3,325), but it is anticipated only once every eight years, so the expected rate of cost change is \$3,325 per year in Scenario C against \$937.50 per year in Scenario E. The boundary responds to this product $\nu|d_e|$ rather than to the size of a single jump, and the fourfold difference in the two responses largely reflects the 3.5-fold difference in it. The remainder is directional, as Subsection 2.3.3 explains: a realised cost cut $\Delta C < 0$ lowers the hurdle on D by $(r/\lambda)|\Delta C|$ at the instant it occurs, so an anticipated reduction is worth somewhat more than an anticipated increase of the same expected annual size, because it can be acted on once it materialises, whereas an increase can only be pre-empted. Both effects are nonetheless large enough to reverse decisions: Scenario C moves Seattle ($D_0 = 7.94$) from adopt to wait, and Scenario E moves Miami ($D_0 = 4.80$) from wait to adopt.

4.3.2 Deterministic Decline versus Constant Cost: Convergence

The contrast between deterministic decline (Scenario F, $d^*(C_0) = 7.31$ \$/100mi) and the constant-cost benchmark (Scenarios A, B, D, $x^{**} = 6.08$ \$/100mi) isolates the effect of anticipated cost convergence on adoption timing. Both specifications eliminate cost volatility ($\eta = 0$), differing only in whether the decision-maker anticipates future cost reductions.

When costs decline deterministically along $C(t) = C_0 e^{-\kappa t}$ with half-life 2.8 years, the predictable cost reduction creates an incentive to delay adoption until prices fall. However, waiting forgoes operating-cost savings during the deferral period. With the calibrated parameters, the incentive to wait for lower costs outweighs the cost of foregone benefits, producing a 20% higher threshold relative to the constant-cost case.

At the current state $(C_0, D_0) = (5{,}800, 1.45)$, Scenario F prescribes waiting with a gap of 5.86 \$/100mi to the threshold. The positive drift in operating benefits ($\alpha = 0.34$ \$/100mi per year) and the anticipated cost decline of approximately \$1,283 within one year (calculated as $C_0(1 - e^{-\kappa})$) both create incentives to defer, making the case for waiting stronger than under constant costs. In contrast, Scenario A prescribes waiting for a different reason: the 4.63 \$/100mi gap reflects the value of preserving flexibility under benefit uncertainty when no cost relief is anticipated.

4.3.3 Deterministic versus Stochastic Decline: Volatility

Comparing the deterministic decline case (Scenario F, $d^*(C_0) = 7.31$ \$/100mi) with stochastic mean-reversion (Scenario G, $d^* = 7.93$ \$/100mi) isolates the pure effect of cost volatility. Both specifications share identical expected cost trajectories converging toward parity at rate $\kappa = 0.25$ per year, differing only in whether this path is certain or stochastic (η = \$2,320/$\sqrt{\text{year}}$). The 8.5% threshold increase quantifies the pure volatility effect:[7] uncertainty creates asymmetric payoffs where adopting at temporarily elevated costs locks in an unfavourable price, while waiting preserves flexibility to time adoption at favourable draws. Mean-reversion amplifies this effect since temporary cost spikes are transitory, making premature adoption particularly costly.

The option value created by cost uncertainty is quantitatively significant. At the current state, the option value under Scenario G is \$12,926 (Richardson-extrapolated 400×600 production pair; \$12,951 on the fine grid alone), representing 6.5% of the immediate exercise value. This translates to a threshold shift equivalent to requiring an additional \$296 in annual operating-cost savings (the unrounded threshold gap of 0.618 \$/100mi times λ) to trigger adoption.

4.3.4 Infinite versus Finite Horizon: The Annuity Penalty

Comparing infinite-horizon stochastic mean-reversion (Scenario G, $d^* = 7.93$ \$/100mi) with its finite-horizon counterpart (Scenario H, $d^* = 26.91$ \$/100mi) isolates the effect of planning horizon on adoption timing. Both specifications share identical cost and benefit dynamics, differing only in the horizon over which benefits accrue.

The 239% threshold increase quantifies the annuity penalty. This effect dominates the other structural choices: the finite-infinite gap (18.98 \$/100mi) is more than ten times the constant-deterministic and stochastic-deterministic gaps combined (1.23 and 0.62 \$/100mi), and it exceeds even the jump-anticipation gap of the calibrated policy scenario (5.58 \$/100mi) by a factor of three. With only 10 years to amortise the \$5,800 upfront premium, the required annual operating-cost advantage must be substantially higher. At the current benefit level $D_0 = 1.45$ \$/100mi, immediate adoption would generate annual savings of \$696, with a 10-year annuity present value of \$6,013, barely exceeding the upfront cost. However, delaying positions the firm to benefit from both cost convergence and benefit drift: expected benefits reach

[7] The Scenario F boundary is the semi-analytical rule of Section 3.2. Its level is set by asking whether adopting now beats waiting for a fixed benefit level; a firm that follows the rule does better than that, because its hurdle falls as the cost falls, so the rule places adoption slightly too early. Solving the deterministic problem exactly, by running the numerical scheme of Section 3.4 at $\eta = 0$, puts the boundary at 7.43 \$/100mi at C_0; the gap closes as the cost falls, and the value lost by following the semi-analytical rule from the baseline state is \$8. The 0.62 \$/100mi gap between Scenarios F and G therefore contains about 0.12 \$/100mi of this approximation, and the effect of volatility alone is closer to 0.50 \$/100mi.

3.15 $/100mi in 5 years while expected costs decline to $1,662, justifying the high threshold.

The steep positive slope of the finite-horizon boundary reflects amplified cost sensitivity: each additional dollar of upfront premium requires proportionally more annual savings to break even within the constrained payoff window, whereas the infinite-horizon boundary is dampened by long-run mean-reversion and extended payoff periods.

4.3.5 Summary: Model Specification Effects

Four modelling dimensions drive threshold variation across specifications, in descending order of magnitude: (1) planning horizon (18.98 $/100mi difference), (2) anticipated cost jumps (up to 5.58 $/100mi, and signed: +5.58 under anticipated reductions, −1.40 under anticipated withdrawal), (3) expected cost convergence (1.23 $/100mi difference), and (4) cost volatility (0.62 $/100mi difference). These effects are not merely technical refinements but reflect distinct economic mechanisms with different policy levers.

The horizon effect dominates because it directly determines the payoff window over which operating savings accrue. No amount of cost reduction or volatility elimination can overcome an insufficient amortisation period if annual benefits are modest. This finding suggests that adoption barriers for short-horizon decision-makers (fleet operators with 5–7 year replacement cycles, commercial lessees, or households with uncertain tenure) require different policy responses than barriers for long-horizon adopters.

Anticipating cost decline raises the threshold above the constant-cost case: knowing costs will fall, the decision-maker has an incentive to defer adoption and pay less. This creates a short-term paradox for policy. R&D support and learning curve acceleration may temporarily delay adoption by strengthening the case for waiting, even as they ultimately reduce the upfront cost barrier. The form of the commitment matters: deterministic cost guarantees with fixed timelines are preferable to vague innovation announcements, as the latter add volatility without resolving uncertainty, reinforcing the option value of waiting. Transparent, credible, and time-bound cost reduction commitments therefore minimise the delay effect while preserving the long-run adoption incentive.

Three of the four are levels of a single hierarchy: planning horizon, expected convergence and cost volatility all raise the hurdle in the same direction, each is measured against the constant-cost benchmark at the same state, and each is fixed by the decision environment, so the three can be ranked against one another. Anticipated jumps are orthogonal to that ranking rather than a fourth level of it: their contribution scales with the anticipated rate of cost change $\nu|d_e|$, which is a feature of the announced policy or technology change rather than of the decision environment, so the 5.58 \$/100mi recorded above is the size of the calibrated scenario and not a structural magnitude. A larger credit or a faster stream of breakthroughs would move the boundary further while leaving the horizon, convergence and volatility effects where they are. The effect is signed as well, and it registers when the hazard is perceived rather than when the cost itself moves. A jump term can therefore be superimposed on any of the three specifications without disturbing their order, and is worth carrying in any application, whatever horizon and cost dynamics are assumed. Its design implications are taken up in Section 5.3.

At the current state $(C_0, D_0) = (5{,}800, 1.45)$, all three specifications (F, G, H) prescribe waiting, but the economic rationale and policy sensitivity differ. Scenario F reflects the incentive to wait for anticipated cost convergence, where targeted cost-reduction commitments would most directly address the barrier. Scenario G reflects additional option value from cost uncertainty, requiring coordinated volatility reduction alongside cost convergence. Scenario H indicates a binding horizon constraint requiring structural extensions to the effective payoff window. Policymakers must diagnose which constraint binds before designing interventions.

4.4 Deterministic Cost Decline (Scenario F)

Scenario F replaces the constant-cost assumption with a deterministic exponential decline $C(t) = C_0 e^{-\kappa t}$, representing convergence toward cost parity at a rate of $\kappa = 0.25$ per year. While this fixed-cost trajectory is perfectly anticipated, the unit cost advantage process $D(t)$ remains stochastic, preserving downside uncertainty. The optimal decision rule arises from the interaction between the deterministic fixed cost decay, the unit cost advantage drift and volatility. Together, these dynamics generate a strict waiting premium: an elevated hurdle rate where deferring investment dominates immediate adoption in order to exploit the anticipated decline in the upfront cost differential. Figure 20 displays the dynamic switching threshold (the operational boundary (3.18)) as a function of the current cost state, alongside the constant-cost benchmark and the ex-ante theoretical target derived via the Lambert $\mathcal{W}$ function (3.16).

To illustrate the dynamic execution of this policy, the figure overlays two simulated stochastic paths of the firm's state $(C(t), D(t))$ in the phase plane. The horizontal axis represents cost, not time; because $C(t) = C_0 e^{-\kappa t}$ declines deterministically, the passage of calendar time drives each state leftward across the plot. Simultaneously, $D(t)$ fluctuates vertically according to its ABM dynamics. Each trajectory therefore traces a leftward-drifting, vertically oscillating path through the (C, D) plane. Adoption is triggered at the exact moment the path intersects the sloped operational boundary $d^*(C)$ from below. The theoretical target x^{**} is a static, path-dependent forecast: the minimum benefit level required to justify adoption, computed ex-ante from the initial state (C_0, D_0). In contrast, the operational boundary $d^*(C)$

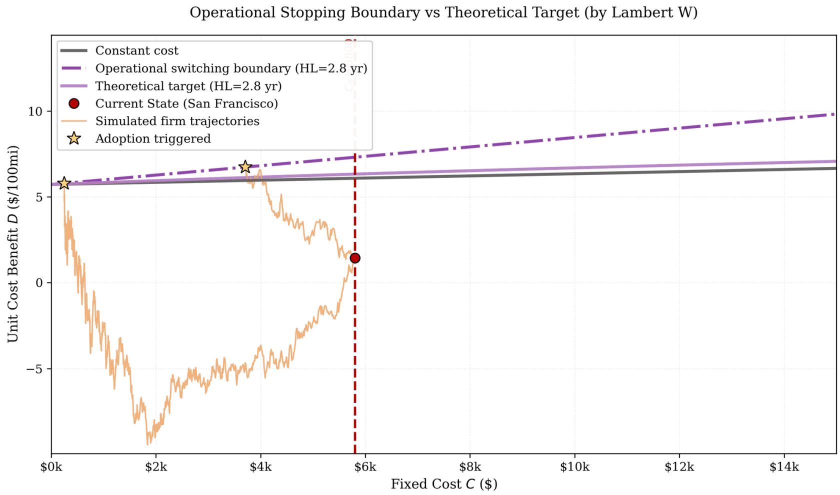

Figure 20: Operational switching boundary $d^*(C)$ and theoretical threshold x^{**} under deterministic exponential decline ($\kappa = 0.25$ per year, half-life 2.8 years), alongside the constant-cost benchmark. Two simulated firm trajectories show D evolving stochastically while C declines deterministically. Because the horizontal axis is cost rather than time, forward motion in calendar time corresponds to leftward drift in the phase plane. Adoption is triggered when the trajectory first crosses the operational boundary from below (yellow star). The current state $(C_0, D_0) = (5{,}800, 1.45)$ is marked in red.

serves as the true closed-loop decision rule, prescribing the state-dependent benefit level required for immediate execution at any continuously realised cost C_t. Because the current state $(C_0, D_0) = (5{,}800, 1.45)$ lies deep within the continuation region, well below both thresholds, the optimal policy definitively prescribes waiting.

The divergence between these curves reflects a simple economic distinction between acting today and planning for the future. Because the firm cannot lower its fixed cost C_t without waiting, the passage of time directly generates capital savings. The operational boundary $d^*(C)$ is steeply sloped because it governs immediate execution: when C_t is high today, the absolute dollar savings from deferring adoption by one year (κC_t) are substantial. To justify committing to the investment immediately, and permanently forfeiting those near-term cost reductions, the firm demands a commensurately high current benefit.

Conversely, the theoretical target x^{**} remains comparatively flat because it is a long-term forecast. It correctly anticipates that from the current state, a lengthy waiting period is required before $D(t)$ rises sufficiently to justify adoption. By the time that future date arrives, the cost differential will have largely decayed away, leaving little remaining convergence to forfeit. The expected adoption hurdle therefore aligns closely with the constant-cost benchmark.

In practice, these two constructs answer different managerial questions. The theoretical target x^{**} serves as a long-term forecasting tool, estimating the benefit level likely to trigger adoption in the future. Because this forecast looks far ahead to a state in which costs have largely converged, it yields a naturally low, stable expectation. The operational boundary $d^*(C)$, by contrast, serves as the live execution rule, determining whether current conditions justify irreversible commitment. Because costs are actively declining, adopting today means forfeiting imminent cost reductions; the operational boundary explicitly prices those forfeited savings into the adoption hurdle, ensuring the firm commits only when the current unit benefit is sufficient to offset the value of continued waiting.

A longer half-life implies slower cost convergence, which reduces the incentive to defer adoption at any given cost level. Figure 21 confirms this: lengthening the half-life from 2.8 to 4.6 years (κ from 0.25 to 0.15) produces a less steeply sloped operational boundary, shifting the adoption hurdle downward across the entire cost range. Both boundaries share the same intercept at $C = 0$, where no remaining convergence exists to motivate deferral regardless of speed. The deterministic case thus serves as a lower bound: introducing cost volatility $\eta > 0$ can only raise the threshold further by generating additional option value from timing flexibility, as quantified in Scenario G.

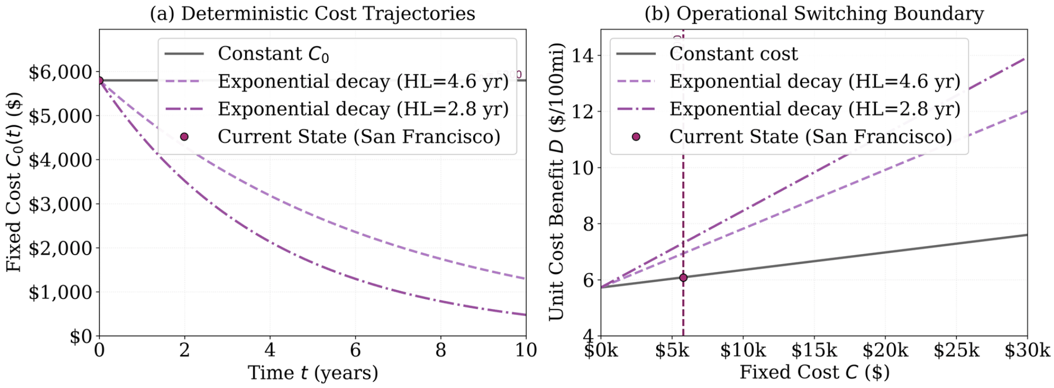

Figure 21: Sensitivity of cost trajectories and adoption boundaries to convergence rate κ. *(a)* Deterministic cost paths from the same initial state $C_0 = 5{,}800$ under half-life 2.8 years ($\kappa = 0.25$) and 4.6 years ($\kappa = 0.15$); the horizontal axis is calendar time and the vertical axis is the realised cost $C(t) = C_0 e^{-\kappa t}$. *(b)* Operational switching boundary $d^*(C)$ in the (C, D) phase plane: the minimum unit cost advantage required to justify immediate adoption at each cost level. Faster convergence (shorter half-life) produces a steeper boundary, reflecting the stronger incentive to defer when near-term cost reductions are substantial.

Differentiating the boundary (3.18) with respect to θ yields

$$\frac{\partial d^*}{\partial \theta} = \frac{r}{\lambda}\left(1 - \frac{\kappa_\kappa}{\kappa_0}\right) < 0 \quad \text{for } \kappa > 0, \tag{4.1}$$

with κ_0 and κ_κ as defined in (3.12); since $\kappa_\kappa > \kappa_0$ for $\kappa > 0$, the bracket is negative. The sign confirms that a lower long-run cost floor raises the individual adoption threshold: deferral offers access to deeper anticipated cost reductions, so the

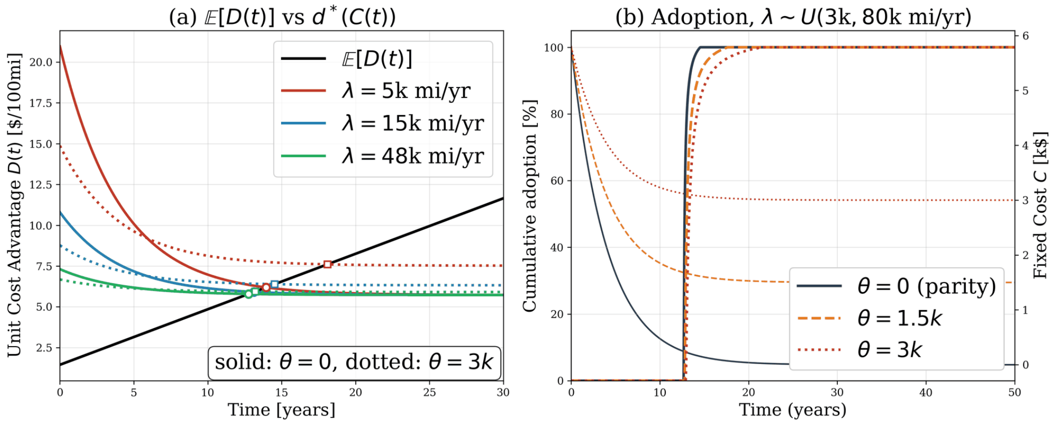

Figure 22: Schematic effect of the long-run cost level θ on adoption timing across a heterogeneous population. Panel (a): the expected unit cost advantage path $\mathbb{E}[D(t)] = D_0 + \alpha t$ (black) rises over time while the adoption boundary $d^*(C(t))$ falls as costs decline along the deterministic OU mean path. Boundaries are shown for three usage intensities; adoption occurs at the intersection (markers). Solid lines: $\theta = 0$; dotted lines: $\theta = \$3{,}000$. Panel (b): cumulative adoption share for a population with $\lambda \sim U(3{,}000,\ 80{,}000)$ miles/year under three cost floor scenarios. Unit cost advantage volatility is suppressed (D follows its expected path) to isolate the θ channel.

decision-maker demands a higher current unit cost advantage to justify immediate commitment. Figure 22(a) illustrates this for three representative usage intensities. Low-usage operators ($\lambda = 5{,}000$ miles/year) face a substantially higher boundary because the fixed adoption cost C weighs more heavily relative to their smaller annual benefit stream λD.

At the population level, however, a lower θ accelerates aggregate adoption. As $C(t)$ converges toward a lower floor, the boundary $d^*(C(t))$ eventually falls below the rising expected benefit for operators whose usage intensity would leave them permanently in the continuation region under a higher floor. Figure 22(b) shows the resulting adoption S-curves for λ drawn from a uniform distribution. Under $\theta = 0$, adoption begins earlier and completes faster than under $\theta = \$3{,}000$. The effect is moderate under the calibrated parameters because cost convergence ($\kappa = 0.25$, half-life 2.8 years) is fast relative to benefit growth ($\alpha = 0.34$ $/100mi per year): the benefit drift, not the cost floor, is the binding constraint on adoption timing.

4.5 Infinite Horizon Model Analysis (Scenario G)

This section presents the numerical results for Scenario G, the infinite-horizon model with stochastic mean-reverting costs. After characterising the value function structure and optimal boundary, we examine how the results depend on the calibrated parameters κ (mean reversion rate), η (cost volatility), and ρ (correlation), quantifying the robustness of the adoption decision to parameter uncertainty.

Figure 23 displays the numerical solution $V(c, d)$ obtained from the PSOR algorithm (Section 3.4) as a three-dimensional surface over the (c, d) state space. The surface exhibits two distinct regions separated by the optimal switching boundary (red curve). In the stopping region, where benefits are high relative to costs, the surface coincides with the exercise payoff plane $\hat{g}(c, d) = \lambda d/r + \lambda\alpha/r^2 - c$ of (3.4), appearing geometrically flat. This planar section corresponds to states where immediate adoption is optimal: the value function equals the intrinsic payoff because further delay would only reduce value through discounting. In the continuation region, where costs are elevated or benefits depressed, the surface curves upward above the exercise payoff plane, with the vertical distance representing the option value of waiting. The curved surface reflects how this option value varies with both state variables: it increases as states move deeper into unfavourable territory where mean reversion and volatility create greater value from strategic timing.

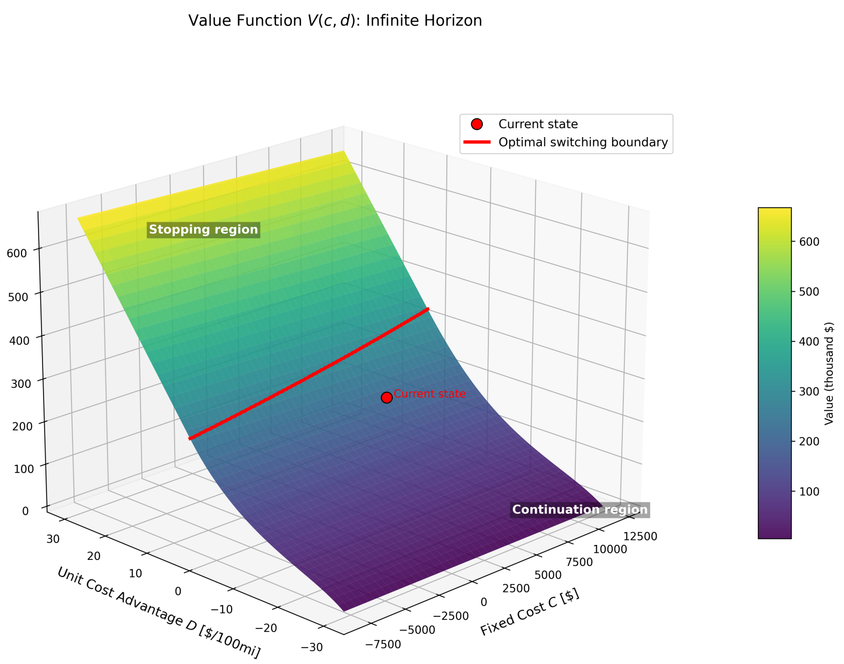

Figure 23: Three-dimensional visualisation of the value function $V(c, d)$ for Scenario G (infinite horizon, stochastic costs). The surface separates into a continuation region (curved) where $V > \hat{g}$ and a stopping region (planar) where $V = \hat{g}$. The red curve traces the optimal adoption boundary. The current state $(C_0, D_0) = (5{,}800, 1.45)$ is marked with a red dot in the continuation region.

The boundary between these regions, traced by the red curve, represents the optimal stopping boundary $d = b(c)$. At this boundary, the value function smoothly transitions from the curved continuation surface to the flat exercise plane, satisfying both the value-matching condition $V(c, b(c)) = \hat{g}(c, b(c))$ and the smooth-pasting condition $\partial V/\partial d = \partial \hat{g}/\partial d$ that characterise optimal stopping problems. The upward slope of the boundary reflects the economic intuition that higher costs require higher benefits to justify immediate adoption. At the current cost level $C_0 = \$5{,}800$, the optimal threshold is $d^*(C_0) = 7.93$ \$/100mi, substantially above the current benefit level $D_0 = 1.45$ \$/100mi.

Figure 24 projects the option value $V(c, d) - \hat{g}(c, d)$ onto the (c, d) plane with the boundary overlaid. The contour structure reveals several features of the timing problem. First, option value is identically zero in the stopping region above the boundary, as immediate exercise eliminates all timing flexibility. Second, option value increases monotonically as states move deeper into the continuation region, reaching approximately \$311,000 at the computational boundary. This extreme value reflects strong mean-reversion dynamics: when costs are highly elevated ($c > \$10{,}000$) or benefits severely depressed ($d < -20$ \$/100mi), anticipated favourable evolution makes waiting extremely valuable. Third, the contours are widely spaced just below the boundary and compress as the state moves away from it. This is smooth pasting seen geometrically: the option value vanishes at the boundary with zero slope, so for a firm already close to indifference a further improvement in the unit cost advantage changes the value of waiting very little. Deep in the continuation region the same improvement is worth a great deal more, and the level sets crowd together accordingly.

At the current state $(C_0, D_0) = (5{,}800, 1.45)$, the option value of \$12,926 represents 6.5% of the immediate exercise value $\hat{g}(C_0, D_0) = \$198{,}733$. While waiting is optimal, this modest ratio indicates that the incremental gain from timing the investment optimally, relative to adopting immediately, is small. The firm stands to benefit from deferral, but the value of that flexibility is limited compared to the perpetual unit cost advantage flow. Moderate shifts in either fuel prices or vehicle costs could bring the current state close to the adoption boundary within a few years.

The two-factor model introduces three parameters not present in the constant-cost framework: the cost volatility η, the mean-reversion rate κ, and the correlation ρ between cost and benefit innovations. While $\eta = 2{,}320$ \$/$\sqrt{\text{year}}$ and $\kappa = 0.25$ per year are calibrated from UK vehicle price data in Section 4.1, the correlation parameter $\rho = -0.3$ reflects economic reasoning rather than direct estimation. We now examine how the optimal adoption boundary responds to variation in η, κ and ρ, quantifying the robustness of the adoption decision to parameter uncertainty.

The mean-reversion rate κ governs the speed at which cost differentials converge toward long-run parity $\theta = 0$. Figure 25 displays the optimal boundary for three values: $\kappa = 0.15$ per year (slow convergence, half-life 4.6 years), the baseline $\kappa = 0.25$ per year (half-life 2.8 years), and $\kappa = 0.35$ per year (rapid convergence, half-life 2.0 years).

Faster mean reversion systematically raises the adoption threshold when costs are above parity. At the current cost level $C_0 = \$5{,}800$, increasing κ from 0.15 to 0.35

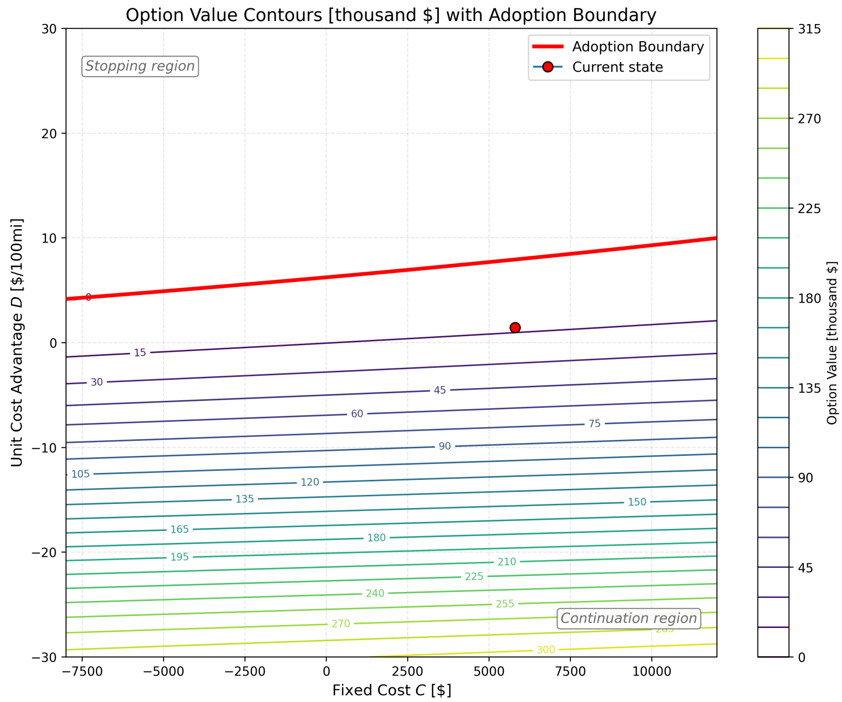

Figure 24: Contour plot of option value $V(c,d) - \hat{g}(c,d)$ [thousand \$] with optimal adoption boundary overlaid. Option value is zero in the stopping region and increases monotonically in the continuation region. Contours are widely spaced just below the boundary, where smooth pasting flattens the option value, and compress deeper in the continuation region. The current state exhibits option value of \$12,926.

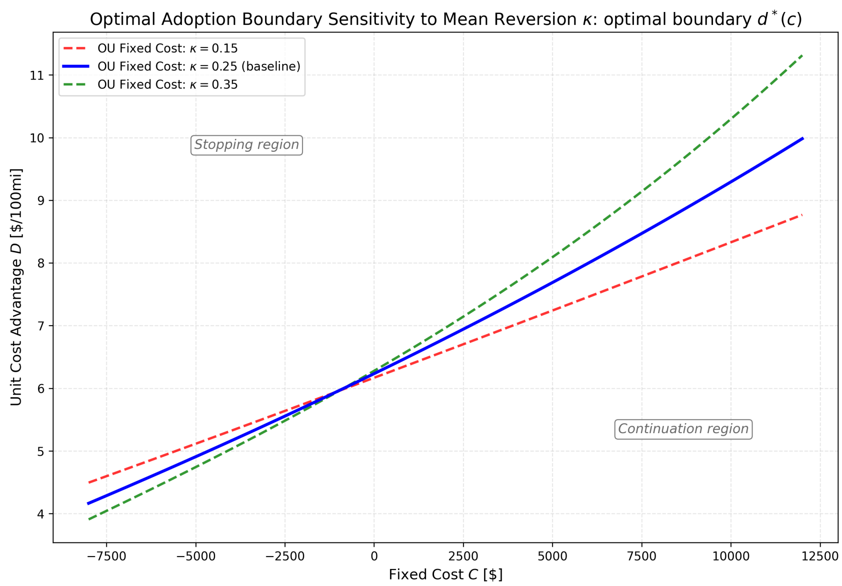

Figure 25: Sensitivity of the optimal adoption boundary to mean-reversion rate κ. The effect is non-monotonic: higher κ raises the threshold when costs are elevated, as stronger mean reversion increases the value of waiting for convergence, but lowers it when costs are sufficiently negative, where mean reversion erodes the realised cost advantage and reduces the incentive to wait.

raises the threshold from approximately 7.41 \$/100mi to 8.42 \$/100mi, an increase of 14%. This effect amplifies at higher cost levels: at $c = \$10{,}000$, the same parameter variation raises the threshold from approximately 8.33 \$/100mi to 10.30 \$/100mi, an increase of 24%. The economic mechanism is clear: when costs are elevated, faster mean reversion implies a stronger expected pull toward parity, increasing the value of waiting for that convergence and raising the benefit threshold required to justify immediate adoption. The direction reverses once costs fall sufficiently below parity, producing the boundary crossing visible in Figure 25: in the negative-C domain, faster mean reversion erodes the realised cost advantage and reduces the incentive to wait, placing the high-κ boundary below the low-κ boundary.

Near parity, the adoption decision becomes insensitive to κ because there is little cost convergence to anticipate. This convergence property confirms that the model correctly nests the constant-cost case as a limiting scenario when the cost state approaches its long-run target. Interestingly, although the spread across κ values narrows near parity, the dynamic boundaries do not converge at cost parity ($C = 0$) but intersect in the negative-C domain. This reflects the firm's opportunistic behaviour in the presence of volatility. At exact parity, the fixed costs of the two technologies are equal, but the firm retains an incentive to wait: because costs are volatile, the new technology may become temporarily cheaper than the incumbent ($C < 0$) due to transient market shocks or subsidies. This prospect keeps the stochastic boundary strictly elevated above the deterministic baseline at $C = 0$. The firm abandons this waiting behaviour only when the current state drops into sufficiently negative territory. Once the upfront cost advantage becomes large enough, the risk of losing it to mean-reversion outweighs the option value of waiting for yet more favourable conditions, and the boundaries converge.

The sensitivity to κ carries direct implications for parameter uncertainty. The calibrated value $\kappa = 0.25$ per year is estimated with a standard error of approximately 0.09 (Section 4.1), so the range $\kappa \in [0.15, 0.35]$ examined above spans roughly ± 1.1 standard errors around the point estimate, which is plausible sampling variation rather than an extreme stress test. Within this range the boundary moves materially: the benefit hurdle shifts by 14% at C_0 and by 24% at $c = \$10{,}000$. The location of the adoption boundary is therefore estimated only imprecisely, and for states in its vicinity, estimation error in κ alone could reverse the optimal decision. This uncertainty does not, however, overturn the decision at the San Francisco benchmark state (C_0, D_0), the location underlying Scenario A: a benefit level of 1.45 \$/100mi lies far below the boundary under every value in the range, so the recommendation to wait is robust for this case even though the precise boundary location is not.

The correlation parameter ρ captures the comovement between cost and benefit innovations. Figure 26 examines three cases: the baseline $\rho = -0.3$ (negative correlation), $\rho = 0$ (independence), and $\rho = +0.3$ (positive correlation). Correlation effects are substantial and monotonic: more negative correlation raises the threshold, making waiting more attractive. At the current cost level $C_0 = \$5{,}800$, the threshold varies from 7.19 \$/100mi under $\rho = +0.3$ to 7.93 \$/100mi under $\rho = -0.3$, a range of approximately 9%. At higher cost levels, the sensitivity amplifies: at $c = \$10{,}000$, the threshold ranges from approximately 8.58 to 9.29 \$/100mi (approximately 8% variation).

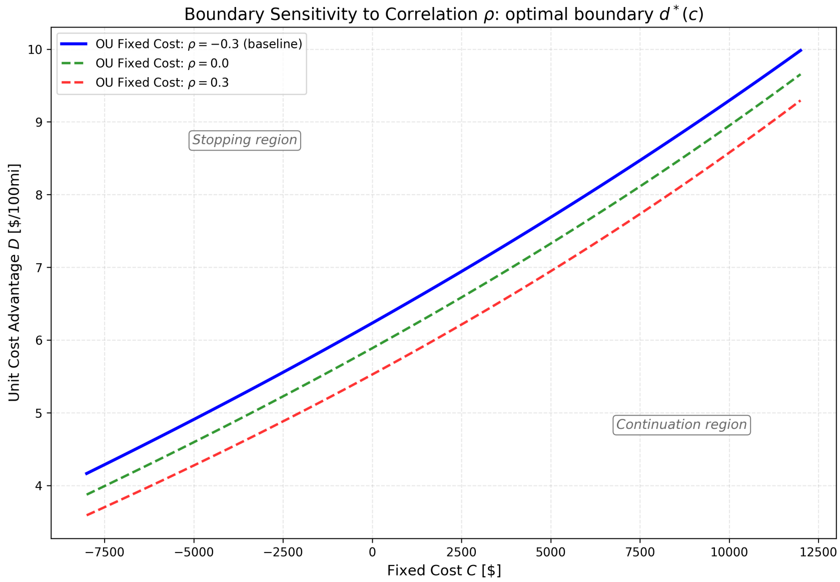

Figure 26: Sensitivity of the optimal adoption boundary to correlation ρ between cost and benefit innovations. Negative correlation elevates the threshold: because the exercise payoff is a difference, negatively correlated factors widen its spread, raising the option value of waiting.

The mechanism is the one set out in Subsection 3.1: what matters is the dispersion of the exercise payoff $\frac{\lambda}{r}D_t - C_t$, whose instantaneous variance (3.21) is decreasing in ρ. Under negative correlation a favourable benefit shock (a rise in D_t) tends to arrive alongside a fall in the adoption cost C_t, so the two reinforce each other in the difference and the payoff is more dispersed. The wider spread raises the value of holding the option, and the decision-maker becomes more selective: the threshold rises. Under positive correlation the two components move together, the difference is steadier, and the option is worth less, so adoption occurs at a lower unit cost advantage. Because the payoff is a difference rather than a sum, it is positive comovement that hedges here, not negative.

The assumed value $\rho = -0.3$ reflects the plausible tendency for fossil fuel price spikes to coincide with temporarily dropping battery and vehicle costs, as technological progress and demand shifts in energy markets may move these variables in opposite directions. This relationship is not structurally guaranteed, however, and could shift with changes in supply chain geography or energy market structure. The 8–9% threshold sensitivity across the correlation range $\rho \in [-0.3, +0.3]$ indicates non-trivial parameter uncertainty, though not sufficient to overturn the qualitative wait decision at the current state.

The cost volatility parameter η determines the magnitude of stochastic fluctuations in the fixed adoption cost (C, i.e. the purchase price differential). Figure 27 displays the optimal boundary for three values: $\eta = 1{,}500$ \$/$\sqrt{\text{year}}$ (low volatility), the calibrated $\eta = 2{,}320$ \$/$\sqrt{\text{year}}$ (baseline), and $\eta = 3{,}000$ \$/$\sqrt{\text{year}}$ (high volatility). The boundaries are nearly parallel, with higher volatility producing uniformly elevated thresholds across the entire cost range. At the current cost level $C_0 = \$5{,}800$, reducing η from 2,320 to 1,500 lowers the threshold by approximately 2.5%; increasing η to 3,000 raises it by approximately 2.4%.[8] The near-parallel nature of the boundaries and the small magnitude of the shifts indicate that volatility effects are approximately additive and quantitatively modest at the baseline calibration: higher uncertainty shifts the entire boundary upward by a roughly constant amount without altering its slope or curvature. This differs from the mean-reversion sensitivity, where boundaries fan out with cost level, and confirms that the timing decision is robust to moderate respecification of cost volatility.

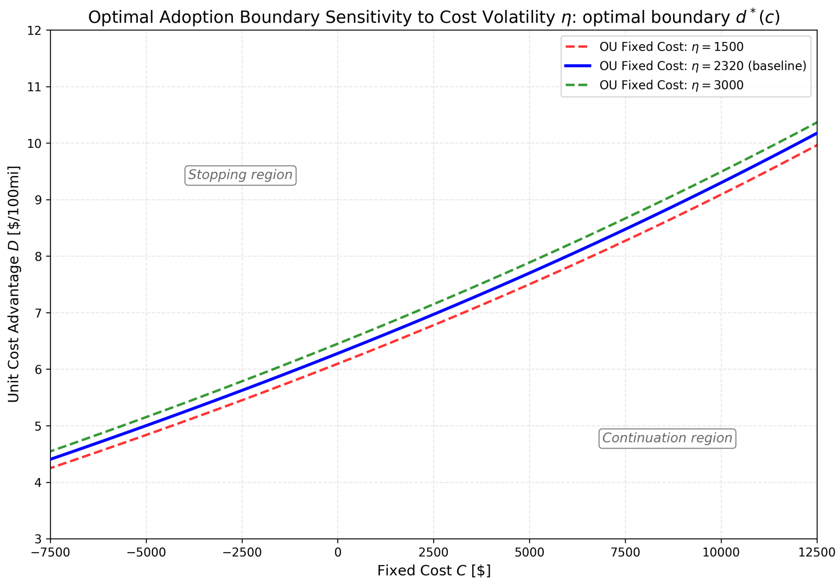

Figure 27: Sensitivity of the optimal adoption boundary to cost volatility η. Higher volatility (larger η) uniformly raises the threshold by increasing option value, making waiting more attractive to avoid adopting at temporarily unfavourable costs. The three boundaries remain close, indicating that volatility effects are roughly additive across cost levels.

[8] The η-sensitivity runs use a computational domain extended by 25% of the baseline range at each end, giving a domain 1.5 times wider overall. At higher cost volatility the artificial domain boundary conditions introduce a more significant bias in the extracted free boundary; widening the domain pushes that bias away from the economically relevant region. The wider domain produces a level shift relative to the baseline-domain threshold of 7.93 \$/100mi at C_0, so the relative shifts (about 2%) rather than the absolute levels are the meaningful quantity.

The economic mechanism is straightforward: higher cost volatility increases the probability of extreme price fluctuations in both directions. The asymmetry of the adoption decision (exercise is irreversible, but waiting preserves flexibility) implies that upside volatility risk (the possibility of adopting just before a sharp cost decline) outweighs downside volatility benefit (the possibility of failing to adopt during a transient period of low costs). This creates net option value that increases monotonically with η, making waiting more attractive as cost uncertainty rises.

The calibrated value $\eta = 2{,}320$ \$/$\sqrt{\text{year}}$ represents substantial uncertainty, corresponding to an annual standard deviation of approximately 35% of the initial cost differential. If future cost evolution exhibits lower volatility than the calibration period, perhaps due to supply chain stabilisation or manufacturing maturation, the threshold falls by approximately 2% relative to the calibrated case. Conversely, if policy uncertainty or technological disruption increases volatility, the threshold rises by approximately 2%. The modest magnitude of these effects (approximately 0.4 \$/100mi threshold variation across the plausible η range) indicates that the adoption decision is only weakly sensitive to volatility assumptions, and less so than to mean-reversion rate or correlation.

The sensitivity analyses produce threshold variations of 14–24% for κ, approximately 9% for ρ, and approximately 5% for η relative to the calibrated specification. The individual effects suggest that parameter uncertainty could shift thresholds non-trivially in either direction, but the qualitative conclusion that waiting is optimal at the current state $(C_0, D_0) = (5{,}800, 1.45)$ is robust across all parameter ranges examined. This robustness reflects the economic structure: elevated cost volatility combined with mean reversion creates substantial option value that delays adoption when unit cost advantages are modest relative to fixed adoption cost differentials.

4.6 Finite Horizon Analysis (Scenario H)

The finite-horizon model limits the planning window to $T = 5$ years with operational lifetime $L = 10$ years, reflecting realistic vehicle replacement cycles and fleet planning constraints. The optimal boundary becomes time-dependent: $d^*(C, t)$. At the decision date $t = 0$, the threshold is $d^*(C_0, 0) = 26.91$ \$/100mi, exceeding the infinite-horizon case by 239%. This dramatic elevation reflects the annuity penalty: because operational benefits accrue only over a finite window, the required annual operating-cost advantage must be substantially higher to amortise the upfront investment.

Figure 28 displays the boundary evolution as the planning horizon shrinks. From the decision date $t = 0$ (blue) through $t = 4$, the boundaries remain tightly clustered, reflecting that the option value of deferral is largely insensitive to the remaining horizon when sufficient time remains. The boundary then collapses sharply at $t = 5$, converging toward the zero-NPV line as the deadline eliminates all waiting value. At the terminal date, the adoption rule reduces to an immediate profitability comparison: vehicles must generate positive net present value upon adoption or be permanently rejected. The boundary exhibits high sensitivity to the cost state, with slope $\partial d^*/\partial C$ substantially larger than in the infinite-horizon case. This amplification reflects how horizon constraints magnify the importance of the upfront cost: with only 10 years to recover the investment, each dollar of additional cost premium translates more directly into higher benefit requirements, whereas infinite-horizon models allow extended payoff periods to dampen this relationship. At the current state $(C_0, D_0) = (5{,}800, 1.45)$, the 25.46 \$/100mi gap to the boundary indicates that adoption would destroy value unless operating conditions improve dramatically.

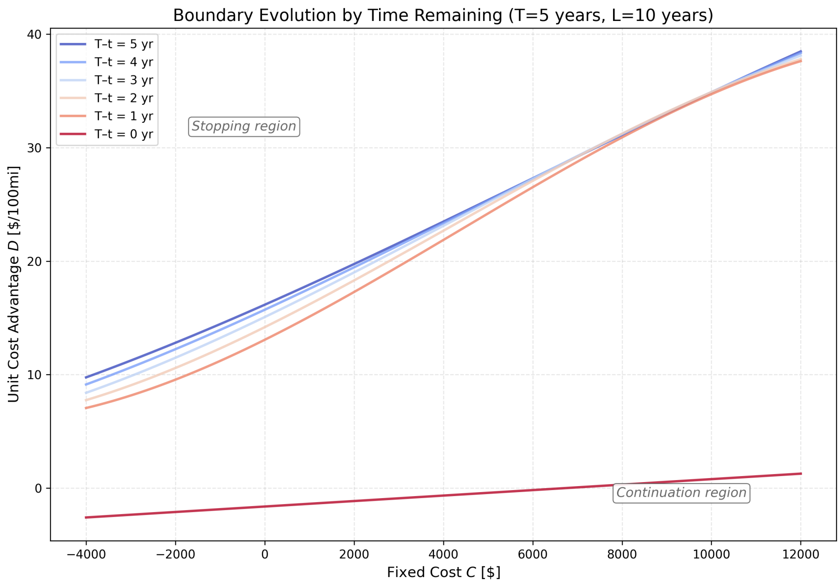

Figure 28: Time evolution of the optimal adoption boundary for Scenario H. Each curve corresponds to a different remaining decision window $T - t$. With the full horizon remaining ($T - t = 5.0$, blue), the boundary lies far above the zero-NPV line ($T - t = 0$, red), reflecting substantial option value from timing flexibility. As the remaining horizon shrinks, the boundary converges toward the intrinsic value threshold, with option value vanishing at expiry. The planning horizon is $T = 5$ years and the operational lifetime post-adoption is $L = 10$ years (Section 3.3 of Chapter 3).

Figure 29 examines how the adoption boundary at $t = 0$ responds to systematic variation in both temporal parameters. As T increases, the boundary shifts upward, reflecting the additional option value of a longer decision window, but with the operating life fixed at $L = 10$ the effect is small: at the baseline cost the boundary rises by less than 0.11 \$/100mi between $T = 5$ and $T = 30$, and beyond $T = 10$ years the curves are nearly indistinguishable, indicating that the stationary solution provides an adequate approximation once the planning window exceeds the mean-reversion half-life of costs a few times over. On the other hand, longer operating lives reduce the annuity penalty by distributing the upfront cost over a larger benefit window, shifting the boundary downward monotonically: at the baseline cost the boundary falls from 26.91 \$/100mi at $L = 10$ to 6.42 at $L = 200$. Convergence is slower than in the case of extending the planning horizon: at $L = 100$ the boundary still lies 1.96 \$/100mi above the infinite-lifetime proxy, and a gap of 0.18 \$/100mi remains even at $L = 200$. The two panels thus confirm that T and L exercise quantitatively as well as qualitatively distinct influences on adoption timing: the annuity factor scaling the benefit stream dominates under the calibrated cost dynamics, while the length of the decision window is second order once it covers a few cost half-lives.

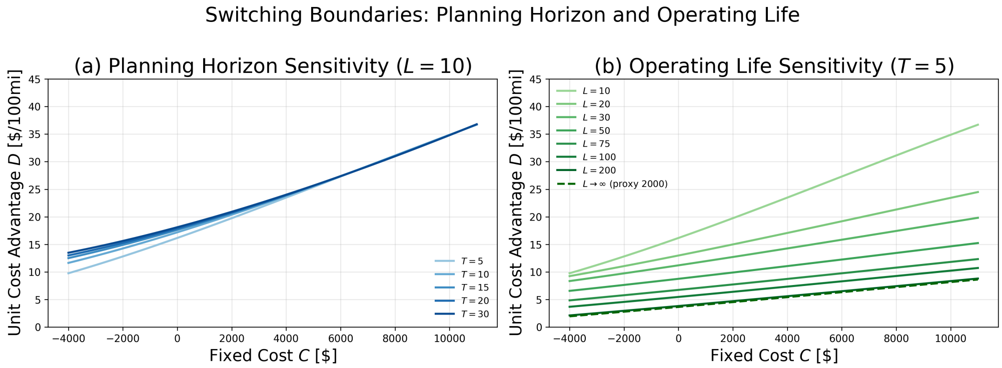

Figure 29: Sensitivity of the optimal adoption boundary at $t = 0$ to the planning horizon T and operating lifetime L. Panel (a): varying $T \in \{5, 10, 15, 20, 30\}$ years with $L = 10$ fixed; the boundaries shift upward only marginally as T grows and are nearly indistinguishable beyond $T = 10$. Panel (b): varying $L \in \{10, 20, 30, 50, 75, 100, 200\}$ years with $T = 5$ fixed, alongside the infinite-lifetime proxy ($L \to \infty$, dashed); longer operating lives reduce the annuity penalty and lower the threshold monotonically.

The finite-horizon solver introduces additional complexity relative to the infinite-horizon case: the solution must be evolved backward through time from a terminal condition, the payoff structure incorporates finite-life annuity factors, and the boundary conditions require careful treatment to prevent instabilities. This necessitates additional numerical validation. Figure 30 validates the finite difference solution against Monte Carlo simulation using the Longstaff-Schwartz algorithm.[9] The validation employs 100,000 simulated paths with time step $\Delta t = 0.01$ years, regressing continuation values on quadratic polynomial features of the state variables.[10]

[9]Standard Monte Carlo simulation is inherently forward-looking and cannot directly handle American-style optimal stopping problems, where the decision to exercise depends on the conditional expectation of future continuation values. The Longstaff-Schwartz algorithm (Longstaff and Schwartz, 2001) overcomes this by using least-squares regression to estimate these conditional expectations from the cross-section of simulated paths, enabling the necessary backward induction.

[10]A 500,000-path re-run with cubic features moves the Monte Carlo boundary at C_0 from 26.4 to 31.1 \$/100mi, against a replication spread of 25.6 to 26.4 over four quadratic runs at 100,000 paths (standard deviation 0.38). Boundary location is sensitive to the basis in a way the value function is not, which is the limitation of least-squares Monte Carlo for free-boundary problems rather than a discrepancy in the finite-difference solution.

The finite difference boundary (blue solid) and Monte Carlo estimates (red dots) show broad agreement across the cost range, though exact alignment is not achieved. This discrepancy is a well-known artefact of numerical optimal stopping. Because the continuation value and the immediate payoff surface are tangent at the boundary by the smooth-pasting condition, the net advantage of waiting is exceptionally flat in this local region. Minor methodological differences between the finite difference discretisation and the Longstaff-Schwartz algorithm's global polynomial basis functions therefore translate into noticeable horizontal variation in the plotted boundary. Furthermore, the two methods align most precisely in the central, economically relevant region of the state space, with discrepancies widening at the extreme boundaries. This behaviour is a direct consequence of the Longstaff-Schwartz methodology: the algorithm estimates continuation values via cross-sectional regression, making its accuracy dependent on the local density of simulated paths. Because simulations naturally concentrate within the empirically relevant parameter space, the regression benefits from high data density in the central region. Conversely, the extreme cost domain boundaries are data-poor, and the polynomial basis functions are forced to extrapolate, leading to expected numerical degradation. The zero-NPV line (grey dashed) represents the intrinsic value threshold; the finite difference boundary lies systematically above this line, capturing the residual option value from the finite operational period $[T, T + L] = [5, 15]$ years.

The finite-horizon analysis reveals that institutional constraints and vehicle depreciation schedules alter the adoption calculus. Policy interventions targeting finite-horizon decision-makers must operate at a larger scale than for perpetual investors: the threshold of 26.91 \$/100mi implies that direct cost subsidies would need to exceed \$10,000 per vehicle to trigger immediate adoption at current benefit levels. Alternative policy approaches that extend the effective operational lifetime, such as vehicle-to-grid programmes monetising battery storage value beyond transportation use or regulatory frameworks facilitating second-life applications, address the constraint driving the high threshold: the limited window to recover the upfront investment.

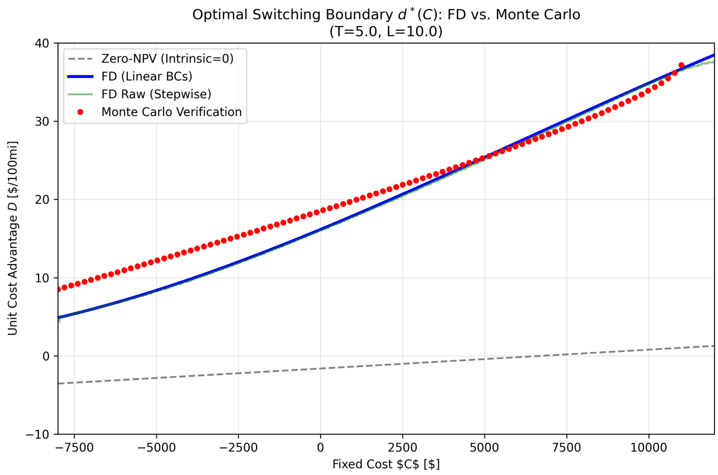

Figure 30: Validation of the finite-horizon boundary against Monte Carlo simulation using the Longstaff-Schwartz algorithm. The finite difference solution with linearised boundary conditions (FD Linear BCs, blue solid) is compared against a level-crossing extraction on the same solution (FD Raw, green), which is quantised to the benefit-direction grid spacing. Monte Carlo estimates (red dots based on 100,000 paths) are shown for cross-validation. The two methods align in the central, economically relevant region of the state space, with discrepancies widening toward the extreme boundaries where simulated paths are sparse. The zero-NPV line (grey dashed) represents the intrinsic value threshold where immediate adoption breaks even. Parameters: $T = 5$ years, $L = 10$ years.

Chapter 5

DISCUSSION AND CONCLUSIONS

This dissertation analyses the optimal timing of irreversible technology adoption when both unit cost advantages and fixed adoption costs evolve stochastically. Chapter 2 establishes the model specification and derives analytical solutions for constant and jump costs, drawing on the author's co-authored publication (Szabó et al., 2025). Chapter 3 extends the analysis by incorporating mean-reverting stochastic cost dynamics and develops a numerical solution methodology using finite difference methods for the resulting two-factor optimal stopping problem. Chapter 4 applies this framework in a calibrated case study of electric vehicle adoption by a high-utilisation delivery fleet, quantifying how cost volatility, mean reversion, and planning horizons interact to determine optimal timing.

5.1 Summary of Main Results

The model sequence yields a progression of results that together establish a diagnostic hierarchy for clean technology adoption decisions. Throughout, results are quantified using a calibrated case study of a high-utilisation delivery fleet replacing internal combustion vehicles with electric equivalents, with benefit dynamics estimated from US energy price data and cost dynamics from UK vehicle price series (Table 8; Subsection 2.4.2 and Section 4.1). In the baseline constant-cost model, with calibrated parameters reflecting US energy price dynamics ($\alpha = 0.34$ \$/100mi per year, $\sigma = 2.42$ \$/100mi per $\sqrt{\text{year}}$), the optimal threshold $x^{**} = 6.08$ \$/100mi is 4.34 times what a payback rule would require. Recovering C_0 over the vehicle's ten-year operating life, discounted at r and ignoring benefit growth, calls for a unit cost advantage of only 1.40 \$/100mi, $D_{\text{BE}} = C_0/(\lambda A_L)$ with $A_L = (1 - e^{-rL})/r$. Geographic variation in energy prices determines whether this threshold is crossed under the calibrated parameters: in Seattle, where cheap hydropower yields $D_0 = 7.94$ \$/100mi, immediate adoption is optimal despite uncertainty, while San Francisco and Miami, with lower initial benefits ($D_0 = 1.45$ and 4.80 \$/100mi respectively), remain in the waiting region.

Extending the model to incorporate discrete jumps in fixed adoption costs reveals a policy timing asymmetry. Anticipated cost reductions delay adoption while anticipated cost increases accelerate it. In the scenario in which decision makers expect future subsidies or technological breakthroughs, the closed-loop threshold rises from 6.08 to 11.66 \$/100mi, a 92% increase. Conversely, expecting subsidy expiration lowers the threshold to 4.68 \$/100mi, a 23% decrease. The 2025 expiration of the US federal EV tax credit under the One Big Beautiful Bill Act was followed by market dynamics consistent with this prediction, as discussed in Section 5.3.

The two-factor model with mean-reverting costs, calibrated from UK vehicle price data (convergence rate $\kappa = 0.25$ per year, volatility $\eta = \$2{,}320/\sqrt{\text{year}}$, long-run parity $\theta = 0$; see Section 4.1), separates convergence from volatility effects. In this calibrated example, the two mechanisms contribute unequally. When costs decline deterministically along $C(t) = C_0 e^{-\kappa t}$, the threshold rises to 7.31 \$/100mi, an increase of 1.23 \$/100mi over the constant-cost benchmark. Introducing cost volatility around the same expected path raises the threshold further to 7.93 \$/100mi, a further 0.62 \$/100mi. Expected convergence therefore contributes roughly twice the threshold increment of pure stochastic volatility in this setting, because the anticipation of lower future costs strengthens the case for waiting even in the absence of uncertainty. Correlation between cost and benefit shocks introduces further asymmetric effects: negative correlation, where unit-cost advantage increases from fuel price spikes coincide with temporary fixed cost reductions, raises the threshold by approximately 9%.

Taken together, the model sequence traces a progression from a stylised one-factor benchmark to a fully specified two-factor stochastic framework, with each extension revealing a qualitatively new mechanism rather than merely refining a quantitative estimate. The constant-cost model establishes that uncertainty alone is sufficient to generate large adoption delays even when the expected payoff is positive; the closed-loop jump-cost extension shows that the direction of anticipated policy change matters as much as its magnitude; the two-factor stochastic model demonstrates that cost uncertainty and unit cost advantage uncertainty interact through a curved free boundary that cannot be summarised by a single scalar threshold x^{**} but requires a state-dependent boundary $d = b(c)$ in the two-dimensional state space; and the finite-horizon extension reveals that amortisation constraints can dwarf all other effects by an order of magnitude. The policy implication of this progression is diagnostic: the appropriate intervention depends entirely on which mechanism dominates for the target population, and misidentifying the binding constraint might lead to proportionally misdirected policy.

Across the model sequence, a clear hierarchy of constraint magnitudes emerges for the calibrated case: planning horizons dominate expected cost convergence, which in turn dominates cost volatility. The relative magnitudes may differ under alternative calibrations, particularly those featuring higher cost volatility or slower mean reversion. Cross-model comparison motivates a distinction between two qualitatively different types of adoption barrier. The infinite-horizon threshold of 7.93 \$/100mi represents what we term 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. This is a behavioural barrier driven entirely by option value, and it is in principle reducible by eliminating cost and benefit volatility; accelerating convergence shortens the calendar wait to adoption but raises the state-contingent benefit hurdle, because a stronger pull toward parity increases the value of waiting at any given cost level above parity. The finite-horizon threshold of 26.91 \$/100mi represents the *structural barrier*: with only five years to decide and ten years to recover the upfront investment, the firm can only just amortise the EV premium at the calibrated current state, whatever its expectations about future fixed adoption costs. The annuity factors available over a finite lifetime capture only a fraction of the infinite-horizon value, so the threshold explodes irrespective of option value considerations. In terms of incremental contributions, the structural barrier (18.98 \$/100mi) is roughly fifteen times larger than the convergence increment (1.23 \$/100mi) and thirty times larger than the pure volatility increment (0.62 \$/100mi). This taxonomy provides a structural explanation for the longstanding energy efficiency gap (Jaffe and Stavins, 1994; Allcott and Greenstone, 2012): the gap might not be primarily strategic (firms miscalculating option value) but structural (firms lacking the operational horizon to justify the investment).

5.2 Theoretical and Methodological Contributions

The single-factor analytical framework in Chapter 2 yields a complete closed-form characterisation of the adoption problem. The optimal threshold x^{**} is derived explicitly as a function of model parameters, enabling direct comparative statics: the threshold is strictly increasing in unit cost advantage volatility σ and strictly decreasing in drift α and usage intensity λ, with each relationship derived analytically rather than inferred from simulation. The multiplicative entry of λ implies that high-utilisation operators face systematically lower adoption hurdles, providing theoretical grounding for the observation that commercial fleets adopt electric vehicles earlier than households. Beyond the threshold, the framework provides a closed-form expected waiting time expression. For the exponentially declining fixed adoption cost extension, the framework further yields a closed-form operational boundary $d^*(C)$ as a function of the current fixed adoption cost level: the practically relevant decision rule, telling a manager at any observed C_t exactly what unit cost advantage is required to trigger adoption immediately, rather than prescribing a single static target computed from initial conditions. This suite of results (threshold and waiting time for the constant-cost case, operational boundary for the exponentially decaying fixed cost extension) provides a directly implementable decision toolkit extending beyond the fixed-cost threshold characterisations standard in the real options literature (Dixit and Pindyck, 1994).

The extension to stochastic mean-reverting fixed adoption costs represents the central methodological contribution. Incorporating an Ornstein–Uhlenbeck cost process transforms the problem into a two-factor optimal stopping problem that no longer admits closed-form solutions. The economic significance of this goes beyond mathematical inconvenience: in the one-factor model, the analytical solution yields a threshold x^{**} that is independent of the current benefit level D_0 and depends on the fixed adoption cost C only as a parameter, so that a manager could in principle implement it as a single static target for any given cost level. Introducing cost volatility ($\eta > 0$) and correlation ($\rho \neq 0$) generates mixed derivative terms in the HJB equation that destroy this structure, and the optimal decision boundary becomes a curved manifold $d = b(c)$. This means the required benefit hurdle is a function of where the current EV–ICE price differential C_t sits relative to its long-run parity trend: when costs are far above parity the firm requires a higher benefit to adopt immediately, when costs are near or below parity the hurdle falls. Because the curved free boundary cannot be characterised analytically, a numerical approach is required. Finite difference methods are adopted for their ability to localise this boundary sharply; the resulting linear complementarity problem is solved using projected successive over-relaxation (PSOR) to accelerate convergence. The correlated dynamics introduce mixed derivative terms that, under standard discretisation schemes, can produce non-monotonic value functions or spurious exercise boundaries. Two discretisation choices address this: standard upwinding for the first-order drift terms, and a sign-oriented 7-point diagonal stencil for the mixed derivative that switches orientation based on the sign of ρ, with the goal of preserving M-matrix properties of the discrete operator. This approach contrasts with simulation-based alternatives such as Least-Squares Monte Carlo (Longstaff and Schwartz, 2001): while LSM is well suited to high-dimensional problems, finite difference methods provide the sharp, precisely localised free boundaries essential for characterising policy thresholds, whereas LSM boundaries exhibit noise from regression approximation particularly at the extremes of the state space where path density is low (Stentoft, 2004).

A distinct methodological contribution is the two-parameter finite-horizon framework separating the planning horizon T from the operational lifetime L. Many finite-horizon real options models collapse these into a single parameter, either by assuming the asset generates returns only until the decision deadline or by treating the decision and payoff windows as identical. The dissertation shows that this conflation is economically substantive: T governs the option value of strategic waiting while L determines the annuity factor scaling the benefit stream, and the two parameters exercise qualitatively different effects on the adoption boundary. Increasing T raises the boundary by expanding strategic flexibility; increasing L lowers it by extending the amortisation window. The calibrated results demonstrate that L effects dominate T effects under realistic vehicle replacement assumptions, establishing that operational lifetime, not decision urgency, is the primary driver of finite-horizon adoption reluctance. This finding could not emerge from a single-parameter horizon model. The finite-horizon formulation also requires a structurally different numerical scheme: whereas the infinite-horizon problem yields a time-invariant elliptic variational inequality solved by PSOR iteration to stationarity, the finite-horizon problem introduces explicit time dependence and is solved by backward induction from the terminal condition, marching the value surface $V(t, c, d)$ and the time-dependent free boundary $d = b(t, c)$ toward $t = 0$. The validation strategy is tailored accordingly: for the infinite-horizon case, grid convergence analysis confirms that discretisation errors remain small relative to the economic effects of interest; for the finite-horizon case, Monte Carlo simulation using the Longstaff-Schwartz algorithm provides an independent boundary cross-check where analytical benchmarks are unavailable.

The analytical characterisation of policy timing asymmetry is a further theoretical contribution. Extending the constant-cost framework to allow discrete jumps in fixed adoption costs reduces the problem exactly to one-dimensional Lévy optimal stopping and delivers the optimal closed-loop adoption boundary $d^*(C)$ in closed form. Under this rule, anticipated future subsidies raise adoption thresholds while anticipated subsidy expiration lowers them, creating use-it-or-lose-it incentives. This asymmetry implies that a subsidy announced with a firm expiry date shifts adoption in time rather than adding to it, unlike the same subsidy with an open-ended timeline: the announced expiry concentrates adoption into the pre-deadline window at the cost of a post-deadline demand collapse. The 2025 OBBBA episode is consistent with this pattern, with the qualifications set out in Section 5.3.

The dissertation also makes a methodological contribution in structural calibration under data constraints. The cross-market parameter synthesis (US energy prices for benefit dynamics, UK vehicle prices for cost dynamics) is not merely a practical compromise: it reflects a principled decoupling of global supply-side learning parameters from local demand-side utilisation parameters. EV–ICE price convergence (κ, θ, η) is driven by vehicle manufacturer production scale, platform standardisation, and global battery supply chains; the rate at which a Nissan Leaf's production premium converges toward ICE parity is likely to be similar across markets, even if local taxes, import duties, and dealer margins introduce some geographic variation. Usage-based advantages (D_0, α, σ), by contrast, are strictly local, determined by regional energy grids, fuel taxation, and usage intensity λ. By isolating these two parameter sets, the calibration methodology provides a portable template: the global cost dynamics can be transferred across markets while only the local benefit parameters need re-estimation. The cost convergence parameters (κ, η) were estimated from UK vehicle price data; the implied half-life of 2.8 years was cross-checked against learning-curve benchmarks from the broader literature, confirming that the

OU specification yields economically plausible convergence dynamics.

5.3 Economic and Policy Implications

The structural–strategic taxonomy established in Section 5.1 sharpens the diagnostic logic of policy design. Instruments that reduce unit cost advantage or fixed cost uncertainty lower the strategic barrier, and instruments that extend the effective operational lifetime address the structural barrier. Instruments that accelerate technology improvement or cost convergence are a separate case: they shorten calendar time to adoption but raise the state-contingent hurdle, so they are not substitutes for uncertainty-reduction policies. Because in this application the structural barrier is fifteen times larger than the convergence increment, policy portfolios weighted toward horizon-extension instruments are likely to be more cost-effective for horizon-constrained populations, though the optimal balance depends on the specific target population and market context.

5.3.1 Rethinking Direct Cost Subsidies

When finite operational lifetimes raise adoption thresholds by 239%, the capacity of direct purchase subsidies to shift adoption decisions is structurally limited: a $7,500 federal tax credit (representing roughly 20% of the base vehicle price) addresses only a fraction of the structural barrier, reducing the fixed adoption cost without extending the operational lifetime over which that cost is amortised. For most clean technology adoption decisions (exceptions include long-lived infrastructure such as nuclear power plants or hydroelectric dams, where operational lifetimes approximate the infinite-horizon benchmark), this amortisation constraint can be the binding barrier, and subsidies risk becoming transfers to those who would adopt regardless rather than incentives moving marginal adopters. Cost subsidies are not ineffective, but they are necessary and insufficient when the binding constraint is the annuity penalty from a short payoff window; cost reduction alone cannot substitute for horizon-extension mechanisms.

Policies accelerating expected cost convergence toward parity, through R&D subsidies or learning-curve programmes, have a more limited timing effect than is commonly assumed. The calibrated results show that anticipated deterministic cost decline raises the adoption threshold relative to the constant-cost case, because the prospect of lower future prices strengthens the case for waiting. Under the calibrated parameters, the incentive to wait for lower fixed adoption costs outweighs the foregone operating savings, producing a 20% threshold increase relative to the constant-cost case. R&D subsidies targeting long-run price parity lower θ and thereby reduce the structural component of the fixed adoption cost over time, but may delay near-term adoption relative to a scenario without anticipated convergence, since the prospect of a lower future C_t strengthens the option value of waiting.

Beyond the level of subsidies, their temporal design matters. The closed-loop jump-cost model identifies a specific failure mode of time-limited subsidies that is absent from standard NPV analysis: the policy cliff. When adopters anticipate a discrete future change in fixed adoption costs, the optimal stopping threshold shifts immediately upon announcement rather than at the moment of change. A subsidy with a firm expiry date is economically analogous to an anticipated positive jump in C_t at the deadline, which lowers the current threshold and pulls forward adoption from the post-expiry period into the pre-deadline window. The 2025 OBBBA episode is consistent with this mechanism operating across two stages, though only loosely: model launches, inventory positioning and the 2025 sales cycle moved the same way over the same window, and the data do not separate them. Prior to enactment, the elimination of the \$7,500 federal EV tax credit was a central policy commitment of the incoming administration and a core provision of the reconciliation bill from its introduction in early 2025; as the bill advanced through Congress with near-certain passage, adopters faced an anticipated jump with rising probability rather than a known future event, which already began pulling forward demand through early summer 2025. Upon enactment on 4 July 2025, the probability reached unity and the deadline was fixed at 30 September 2025, equivalent in the jump-cost framework to announcing a jump of size d_e = \$7,500 arriving with certainty at a known future time. The calibrated Scenario E, which spread a 50% chance of credit removal evenly over four years (intensity $\nu = 1/8$ per year, d_e = \$7,500), already lowered the threshold from 6.08 to 4.68 \$/100mi; the post-enactment scenario with unit probability represents the limiting case, implying a threshold drop larger than 1.40 \$/100mi and a correspondingly stronger adoption-accelerating effect. As the deadline approached, the option value of waiting collapsed toward zero and purchases that would have spread over one to two years concentrated into a single quarter, with US EV sales reaching 458,298 units in Q3 2025, a 22% increase year on year.[1] Once the credit expired, the pool of near-threshold adopters had been exhausted; Q4 2025 EV sales fell to 229,185 units, 43% below the same quarter of 2024 and 50% below Q3.[2] The mechanism is not confined to the United States. Germany's abrupt termination of the Umweltbonus purchase grant in December 2023 was followed by a 27% year-on-year fall in battery-electric registrations in 2024, the worst-performing

[1] Atlas EV Hub, "Q4 and Full-Year 2025 EV Sales Data Insights," `https://www.atlasevhub.com/weekly-digest/q4-2025-sales-data-insights/`, accessed February 2026; Autobody News, "Q4 and Full-Year 2025 Vehicle Sales: GM Leads U.S. Market, Hybrids Surge as EV Sales Collapse," `https://www.autobodynews.com/news/q4-and-full-year-2025-vehicle-sales-gm-leads-u-s-market-hybrids-surge-as-ev-sales-collapse`, January 2026.

[2] ING Research, "US Electric Vehicle Transition Delayed but Not Derailed by Subsidy Exit," `https://think.ing.com/articles/subsidy-exit-delays-the-us-ev-transition-but-wont-derail-it/`, 2025.

powertrain segment of the year.[3] Recent Hungarian and EU purchase-grant schemes follow the same time-limited, budget-capped design that the framework identifies as cliff-prone, suggesting that the delay-and-spike cycle documented for the US and Germany is a general property of deadline-based subsidy architecture rather than a market-specific anomaly. Volume-based or algorithmic phase-outs, where the subsidy declines by a fixed amount for each increment of cumulative sales, would distribute the adjustment across successive cohorts of adopters rather than concentrating it at a single deadline, reducing the magnitude of each individual cliff effect at the cost of introducing multiple smaller discontinuities.

5.3.2 Horizon-Extension and Volatility-Reduction Mechanisms

Given the dominance of operating horizon effects, policies extending effective operational lifetimes merit priority alongside direct cost measures. Developing robust secondary markets for used electric vehicles extends the first owner's effective payoff window. When residual values are uncertain or depreciation is rapid, households rationally also shorten their planning horizons because they cannot reliably monetise remaining vehicle value upon sale. The battery is the component with the most rapid value depreciation and the primary driver of the EV price premium; its degradation risk is therefore the central source of residual value uncertainty in secondary markets. Certification programmes for battery health, standardised warranty transfers, and transparent performance metrics reduce these information asymmetries, supporting higher residual values. The UK battery health certification initiatives and Japan's Battery-as-a-Service models (Weiller, 2014; Shi and Hu, 2022), which separate ownership of the vehicle from ownership of the battery, provide institutional templates: they shift long-term degradation risk from the consumer to the manufacturer and extend the first owner's effective planning horizon without direct fiscal transfers (Agrawal and Bellos, 2017).

Alternative policy approaches address the constraint driving the high adoption threshold: the limited window to recover the fixed adoption cost. Vehicle-to-grid (V2G) integration, beyond its contributions to grid stability and energy security (Aguilar Lopez et al., 2024), may generate additional revenue for EV owners by discharging electricity back to the grid during peak demand periods, potentially increasing the effective annual return on the adoption investment. Battery second-life applications create additional value after the primary transportation lifetime: a battery degraded below automotive thresholds may retain substantial capacity for stationary grid storage, extending the effective payoff window beyond vehicle retirement. Regulatory frameworks facilitating such second-life applications, or accounting standards that properly recognise end-of-life residual values, effectively

[3] Kraftfahrt-Bundesamt registration statistics; see also Autovista24, "BEV slump drives down German new-car market," `https://autovista24.autovistagroup.com/news/bev-slump-drives-down-german-new-car-market/`, 2025.

extend the asset's economic payoff window without requiring the vehicle to remain on the road. For high-utilisation fleets, both mechanisms are particularly relevant: high annual mileage generates operating savings rapidly but also accelerates physical depreciation, necessitating shorter replacement cycles and imposing a strictly finite primary operating horizon; V2G revenues and second-life battery value can partially offset this compressed payback window.

Fixed adoption cost volatility, while secondary to horizon effects in the calibrated setting, raises the threshold by 0.62 \$/100mi relative to the deterministic decline case, representing a distinct policy target. The model demonstrates that uncertainty about trajectories, rather than the trend itself, is what drives option value: deterministic cost decline generates a smaller threshold increment than stochastic decline with identical expected path, because resolved uncertainty eliminates the asymmetric payoffs that make waiting valuable. Benefit volatility operates through a parallel but distinct mechanism: in the one-factor analytical model, higher volatility σ of the unit cost advantage process raises the optimal threshold x^{**} directly, as greater uncertainty over future operating savings increases the option value of waiting even when the expected drift is unchanged (Szabó et al., 2025). Credible long-term technology roadmaps and regulatory certainty reduce option value of waiting more effectively than vague commitments, even when they do not change expected cost levels, a conclusion consistent with empirical evidence that policy uncertainty suppresses green investment independently of cost fundamentals (Wang et al., 2023). California's Advanced Clean Cars II regulation, mandating specific EV sales shares through 2035, and the European Union's phase-out timeline for internal combustion engine sales both address demand-side and supply-side volatility simultaneously: by committing to known future market structures, they remove regulatory uncertainty as a source of additional option value of waiting. Financial instruments such as price guarantees or floor contracts that cap downside risk in the unit cost advantage, for instance by guaranteeing a minimum fuel price differential or electricity tariff discount for fleet operators, address the benefit volatility barrier directly by limiting the losses from early adoption if energy price differentials narrow unexpectedly.

5.3.3 Targeting and Distributional Considerations

The analytical result that optimal thresholds decrease monotonically with usage intensity λ provides theoretical grounding for differentiated subsidy design. High-utilisation operators face lower adoption barriers because operating savings accrue faster; mathematically, they face a steeper loss function for delaying adoption because they forgo substantially larger daily operating savings, making the upfront premium easier to justify. A delivery company driving 48,000 miles per year faces a threshold 23% lower than a household driving 8,000 miles, holding other parameters constant. Empirically, however, Davis (2019) finds that early EV adopters tend to drive fewer miles than average ICE vehicle owners, the opposite of the model's prediction; this divergence likely reflects factors outside the model parameter calibration, such as range anxiety, limited charging infrastructure, and the suitability of EVs primarily as second vehicles for shorter trips. Uniform subsidies therefore misallocate resources: they transfer disproportionate value to high-intensity users who would adopt at lower incentive levels, while providing insufficient support to marginal low-intensity adopters facing higher barriers. This alignment between emissions impact and financial incentive (high-mileage users both benefit most from adoption and face the lowest model-implied threshold) provides a rationale for proposals targeting high-mileage gasoline superusers (Metz and London, 2021; Aemmer et al., 2023): a high-mileage switch displaces the most fuel, and since these users face the smallest gap to their adoption threshold, the incentive required to trigger their switch is also the smallest, so subsidies concentrated on them displace more fuel per dollar than uniform schemes. Differentiated subsidy structures could improve cost-effectiveness through mileage-based phase-outs, per-mile payments targeting high-usage populations while minimising deadweight transfers, or usage-conditional incentives contingent on maximum annual mileage commitments. Commercial fleets with centralised telematics and transparent usage data face minimal verification costs, suggesting differentiated subsidies for commercial segments could be implemented immediately while household programmes maintain simpler structures.

Geographic heterogeneity in energy prices introduces a parallel targeting dimension. The model shows that adoption readiness varies dramatically by location: Seattle, with $D_0 = 7.94$ \$/100mi, crosses the adoption threshold without any subsidy, while San Francisco, with $D_0 = 1.45$ \$/100mi, requires substantially larger interventions. Uniform national subsidies generate deadweight transfers in high-benefit regions while remaining insufficient in low-benefit regions. Systematically tying incentive levels to local energy price differentials and usage intensity would improve cost-effectiveness. In jurisdictions with carbon pricing, the unit cost advantage of electric vehicles increases with the carbon price applied to gasoline, though this benefit is itself subject to policy uncertainty: credible long-term carbon price floors reduce this volatility while increasing expected benefits, operating through both the variance and the drift of the unit cost advantage process to accelerate adoption.

The structural–strategic taxonomy also clarifies how the framework extends beyond electric vehicles. The methodological contribution is diagnostic and not technology-specific: identify the dominant barrier type for the target population, then direct policy toward the corresponding instrument class. For residential heat pumps, the operating advantage fluctuates with gas and electricity prices while installation costs decline through learning-by-doing; an analogous planning horizon constraint likely dominates, driven by homeowner mobility and equipment replacement cycles, though this conjecture requires empirical calibration to confirm. Mechanisms extending effective equipment lifespans, such as portable heat pump modules or standardised installations facilitating resale value transfer, would address the structural barrier more directly than installation rebates. For residential solar, transferable leases and property-attached financing such as Property Assessed Clean Energy (PACE) loans serve the same function. For industrial equipment, the hierarchy may differ: lower price volatility but longer physical lifetimes compress the volatility increment while reducing the annuity penalty, so whether horizon or volatility effects dominate requires context-specific calibration rather than direct transfer of the EV parameters.

5.3.4 Forecasting Bias in Integrated Assessment Models

A recurring theme across model specifications is that standard net present value approaches systematically underestimate adoption barriers and therefore overestimate diffusion speeds. The gap between the NPV break-even point (1.40 \$/100mi, $D_{\mathrm{BE}} = C_0/(\lambda A_L)$ over a 10-year operating life) and the optimal adoption threshold (6.08 \$/100mi in the baseline) is a factor of 4.34, and therefore of a different order from the corrections such models usually contemplate. Large-scale integrated assessment models used to project decarbonisation pathways, including simulation-based energy-economy models of the kind surveyed by Mercure et al. (2018), typically represent adoption through deterministic cost-comparison or discrete-choice rules that abstract from option value and cost volatility. This abstraction generates a structural bias.

By assuming certainty about cost and benefit trajectories, standard IAM formulations trigger adoption at the simple break-even condition, ignoring the option value of waiting that pushes the true threshold higher. A strict deterministic NPV rule parameterised with these same inputs would predict immediate, complete adoption in cities like Miami ($\frac{\lambda D_0}{r} = \frac{480\times4.80}{0.03} = \$76{,}800 \gg C_0 = \$5{,}800$), whereas the option value framework correctly prescribes continued waiting since $D_0 = 4.80 < x^{**} = 6.08$ \$/100mi. For marginal adopters with benefits near the break-even level, the deterministic logic of an IAM structurally overpredicts the speed of adoption. By failing to account for the value of preserving flexibility under uncertainty, these models systematically underestimate the true economic hurdle rate required to trigger irreversible investment.

The magnitude of this effect is not marginal, and its consequences for model design are instructive. IAM practitioners are aware that deterministic NPV rules predict unrealistically sharp adoption cliffs: the moment costs cross break-even, the model switches the entire fleet instantaneously. To prevent this, modellers frequently impose implicit discount rates of 20–30% or introduce logit-share friction parameters to slow predicted diffusion. These adjustments are calibrated to match observed behaviour but lack micro-founded justification. The present real options analysis provides that justification: high effective hurdle rates are not arbitrary behavioural friction, but the rational pricing of option value under uncertainty. Carbon price trajectories, subsidy phase-out schedules, and infrastructure investment timelines derived from IAM simulations should be evaluated accordingly, as policy interventions scaled merely to reach the NPV break-even threshold will consistently fall short of triggering adoption for decision-makers facing irreversible investment decisions under uncertainty.

5.4 Limitations and Directions for Future Research

The model's structural assumptions introduce several boundaries to the analysis. The optimal stopping framework assumes irreversible adoption, ruling out switching back to conventional technologies if market conditions deteriorate or upgrading to next-generation vehicles sequentially. The direction of that approximation is known: any recoverable fraction of the outlay lowers the option value of waiting and therefore the adoption threshold, so the thresholds reported here are upper bounds with respect to this assumption. Future research incorporating compound American options could evaluate whether early adoption is valuable as a stepping stone or whether waiting for next-generation vehicles dominates. The single-decision-maker framework abstracts from social learning, network effects, and infrastructure complementarities. When early adopters generate positive externalities through charging network density or information diffusion, the privately optimal threshold exceeds the socially optimal threshold, providing an additional rationale for adoption subsidies beyond the mechanisms analysed here.

The stochastic process specifications impose further constraints. The unit cost advantage process is modelled as arithmetic Brownian motion and the cost process as an Ornstein-Uhlenbeck process; both choices prioritise analytical tractability over long-run forecast accuracy. Arithmetic Brownian motion allows negative realisations of the unit cost advantage, which are economically meaningful (ICE vehicles may become cheaper to operate than EVs under adverse energy price conditions) but implies unbounded variance over long horizons. Alternative specifications such as geometric Brownian motion would prevent negative values but impose proportional volatility scaling that may be less appropriate for price differentials that can change sign. The mean-reversion specification for costs imposes a stationary long-run equilibrium at $\theta = 0$, embedding strong assumptions about competitive convergence that abstract from persistent cost advantages driven by regulatory asymmetries or lock-in effects. The continuous-time diffusion processes further assume smooth dynamics that may inadequately capture discrete shocks characteristic of technology transitions, though the jump-cost extension in Chapter 2 only partially addresses this for policy shocks: its stationary Poisson specification places the jump at an exponentially distributed time rather than at a known date, so it captures the pull-forward mechanism rather than pricing a fixed deadline. The volatility parameters σ and η are treated as exogenous constants throughout, which is a reasonable approximation for EV adoption where fossil fuel and electricity prices are driven by macro energy markets largely independent of fleet electrification rates at current adoption levels.

The empirical calibration faces several constraints that limit the generalisability of specific threshold values. The cost process estimation relies on fifteen observations from the UK Leaf–Focus price series, which supports two-parameter estimation but with substantial standard errors. The correlation ρ between cost and benefit shocks requires joint high-frequency observations rarely available in practice; the present calibration assigns it on economic grounds in the absence of joint data. The cross-market parameter synthesis, combining US energy prices for benefit dynamics with UK vehicle prices for cost dynamics, reflects data availability constraints but introduces potential misspecification. Because slight variations in drift and volatility parameters can shift the qualitative ranking between waiting and adopting, the specific hierarchy of constraint magnitudes may differ under alternative calibrations, and the single vehicle-pair application limits generalisability across segments with different cost structures and usage patterns.

Two directions offer the most productive extensions. The first is empirical: cross-country and cross-segment replication would test whether the hierarchy of constraints reflects technology fundamentals or market-specific factors. More ambitiously, fitting the model to observed adoption patterns through structural estimation could recover population distributions of planning horizons and risk preferences. A heterogeneous-agent version of the framework, where adopters differ in λ, T, and L, would generate realistic diffusion curves without imposing representative-agent assumptions, and would identify which interventions move which segments of the population, providing a micro-founded basis for targeting policy rather than relying on aggregate calibration. Developing datasets with joint observations of unit cost advantages and fixed adoption costs, potentially from individual-level data combining energy consumption records with vehicle purchase prices, would enable simultaneous estimation of the full parameter vector including the correlation structure.

The second direction is theoretical. The present framework treats the policy environment as exogenous, but in practice policy responds to adoption dynamics: the OBBBA episode illustrates this clearly, as the credit expiration was itself a policy choice responding to prior adoption patterns, and the resulting demand collapse fed back into automaker investment decisions within weeks of the deadline. Endogenising this feedback, where subsidy persistence depends on aggregate adoption rates and adoption rates respond to subsidy expectations, would require game-theoretic or mean-field extensions. A related extension concerns infrastructure complementarities: charging network density affects individual adoption thresholds, which in turn determine network density, creating a coordination problem that individual optimal stopping cannot resolve, and future research could identify whether policy sequencing, infrastructure before subsidies or vice versa, materially affects transition speed. In broader applications of the framework beyond EVs, incorporating time-varying or endogenous volatility would also become relevant: Slameršak et al. (2022) document that the upfront energy requirements of scaling low-carbon infrastructure can create a transitional demand spike that amplifies market volatility precisely when adoption is accelerating, introducing a self-reinforcing mechanism that a fixed-σ framework cannot capture.

The central finding is that planning horizon constraints generate threshold increases an order of magnitude larger than cost volatility or expected convergence under the calibrated parameters, establishing a diagnostic hierarchy that inverts the implicit priority ordering of most current EV policy. Purchase subsidies address the smallest quantified barrier; secondary market development, vehicle-to-grid integration, and regulatory lifetime extension address the largest. Whether this ordering is robust to alternative technologies and market contexts is an empirical question, but the framework provides the structure within which it can be answered.

Appendix A

APPENDIX: HISTORICAL DATA SOURCES

This appendix provides a detailed list of data sources referenced in this report, including historical data for inflation and location-based oil prices and electricity prices for the US and EU.

- **US Inflation**
  Source: World Bank, World Development Indicators
  Description: Annual inflation rates (consumer prices) available from 1960 to 2023.
  Link: `https://databank.worldbank.org/source/world-development-indicators/Series/FP.CPI.TOTL.ZG#`
  Download date: October 4, 2024

- **EUR Inflation**
  Source: Eurostat
  Description: Harmonised Index of Consumer Prices (HICP) for the European Union (27 countries), available from 2011 to 2023.
  Link: `https://ec.europa.eu/eurostat/databrowser/view/tec00118/default/table?lang=en&category=t_prc.t_prc_hicp`
  Download date: October 4, 2024

- **US Oil Prices**
  Source: U.S. Energy Information Administration (EIA)
  Description: Weekly oil prices, including monthly average for the US, with data available from 2002 to 2024.
  Link: `https://www.eia.gov/dnav/pet/hist/leafhandler.ashx?f=m&n=pet&s=emm_epm0_pte_nus_dpg`
  Download date: October 4, 2024

- **EUR Oil Prices**
  Source: European Commission, Energy Directorate
  Description: Monthly opening prices for oil in Europe based on weekly observations, including taxes. Data available from 2005 to 2024.
  Link: `https://energy.ec.europa.eu/document/download/906e60ca-8b6a-44e7-8589-652854d2fd3f_en?filename=Weekly_Oil_Bulletin_Prices_History_maticni_4web.xlsx`
  Download date: October 4, 2024

- **US Electricity Prices**
  Source: U.S. Energy Information Administration (EIA)
  Description: Monthly electricity price data, available from 2002 to June 2024.
  Link: `https://www.eia.gov/electricity/data/browser/#/topic/7?geo=g&agg=0,1&endsec=vg`
  Download date: October 4, 2024

- **EUR Electricity Prices**
  Source: Ember
  Description: Monthly wholesale electricity prices for European countries.
  Link: `https://ember-energy.org/data/european-wholesale-electricity-price-data/`
  Download date: October 4, 2024

- **US Gasoline Prices**
  Source: American Automobile Association (AAA)
  Description: Current gasoline prices for the US, updated as of October 4, 2024.
  Link: `https://gasprices.aaa.com/state-gas-price-averages/`
  Download date: October 4, 2024

- **US Regional Energy Prices**
  Source: U.S. Bureau of Labor Statistics (BLS)
  Description: Regional gasoline and electricity prices
  Link: `https://www.bls.gov/regions/midwest/data/averageenergyprices_selectedareas_table.htm`
  Download date: March 4, 2025

- **UK Vehicle List Prices**
  Source: Comcar UK (comcar.co.uk)
  Description: Historical On-The-Road (OTR) list prices for UK vehicles, organised by manufacturer, model, UK registration plate year, and fuel type. Data includes minimum list prices in GBP, availability dates, and CO2 emissions information. Prices are available from March 2011 to November 2025

(projected), with observations for Nissan Leaf (EV) and Ford Focus (ICE) models.
Link: `https://comcar.co.uk/companycar/tax/history/`
Download date: December 2025

- **US Consumer Price Index (CPI)**
  Source: Federal Reserve Economic Data (FRED), Federal Reserve Bank of St. Louis
  Description: Consumer Price Index for All Urban Consumers: All Items (CPI-AUCSL). Monthly index measuring inflation in the United States, with base period 1982-84 = 100. Data available from 1947 to present, aggregated to monthly observations. Used to deflate nominal USD car prices to constant 2022 USD.
  Link: `https://fred.stlouisfed.org/series/CPIAUCSL`
  API Series ID: CPIAUCSL
  Download date: December 2025

- **GBP-USD Exchange Rate**
  Source: Federal Reserve Economic Data (FRED), Federal Reserve Bank of St. Louis
  Description: U.S. / U.K. Foreign Exchange Rate (DEXUSUK). Daily exchange rates between British Pounds and U.S. Dollars, aggregated to monthly averages for alignment with price observations. The rate represents how many USD one GBP is worth. Data available from 1971 to present. Used to convert nominal GBP prices to nominal USD prices.
  Link: `https://fred.stlouisfed.org/series/DEXUSUK`
  API Series ID: DEXUSUK
  Download date: December 2025

Appendix B

MATHEMATICAL FOUNDATIONS AND SUPPLEMENTARY PROOFS

B.1 Proof of Theorem 2.7

In general, the density function of the first hitting time of a standard Wiener process with an arbitrary drift c ($X(t) = ct + W(t)$, $W(0) = x$) can be expressed as follows (Borodin and Salminen (2015),p.295):

$$f_{T_a}(t) = \frac{|a - x| e^{\frac{-(a-x-ct)^2}{2t}}}{\sqrt{2\pi t^3}} \tag{B.1}$$

where

$$T_a = \inf_{t \geq 0} \{X(t) \geq a\} \tag{B.2}$$

To express the density function for the process with arbitrary volatility ($D(t) = D_0 + \alpha t + \sigma W(t)$, $W(0) = 0$), we conduct a change of variables.

$$D_{x^*} = \inf_{t \geq 0} \{D(t) \geq x^*\} = \inf_{t \geq 0} \{\frac{\alpha}{\sigma} t + W(t) \geq \frac{x^* - D_0}{\sigma}\} \tag{B.3}$$

The density function of D_{x^*} is given by

$$f_{D_{x^*}}(t) = \frac{|\frac{x^* - D_0}{\sigma}| e^{\frac{-(\frac{x^*-D_0}{\sigma} - \frac{\alpha}{\sigma} t)^2}{2t}}}{\sqrt{2\pi t^3}}. \tag{B.4}$$

Note that in the case of positive drifts ($\alpha > 0$) and $D_0 < x^*$, $f_{D_{x^*}}$ is the density function of an IG $\left(\frac{(x^* - D_0)}{\alpha}, \left(\frac{(x^* - D_0)}{\sigma}\right)^2\right)$ distributed random variable.

By applying the properties of IG variables, we obtain

$$\mathbb{E}_{D_0}(D_{x^*}) = \frac{x^* - D_0}{\alpha}. \tag{B.5}$$

This means that for positive drifts ($\alpha > 0$) and $D_0 < x^{**}$, we can substitute Equation (2.19) to obtain

$$\mathbb{E}_{D_0}(D_{x^{**}}) = \frac{1}{\alpha}\Big(\frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\alpha}{r} + \frac{Cr}{\lambda} - D_0\Big). \tag{B.6}$$

On the other hand, in the case of non-positive drifts ($\alpha \leq 0$) and $D_0 < x^{**}$, the expected waiting time is infinite. For $\alpha < 0$ the threshold is reached with probability $\mathbb{P}_{D_0}(D_{x^{**}} < \infty) < 1$; for $\alpha = 0$ it is reached almost surely, but the first-passage time has a heavy tail with no finite mean. In both cases,

$$\mathbb{E}_{D_0}(D_{x^{**}}) = \infty. \tag{B.7}$$

B.2 Proof of Theorem 2.9

As set out in Subsection 2.2.5, Proposition 5.14 of Dayanik and Karatzas (2003) reduces optimality of τ^* to showing that Γ is unbounded above. This follows from the objects constructed there. Write $\kappa_0 = \big(\sqrt{\alpha^2 + 2r\sigma^2} - \alpha\big)/\sigma^2$, so that (2.16) evaluated at the threshold x^{**} of (2.19) reads $V_{D_{x^{**}}}(x) = \hat{g}(x^{**})e^{\kappa_0(x - x^{**})}$ for $x < x^{**}$, and set

$$W(x) = \begin{cases} \hat{g}(x^{**})\, e^{\kappa_0(x - x^{**})}, & x < x^{**}, \\ \hat{g}(x), & x \geq x^{**}. \end{cases}$$

The first-order condition that determines x^{**} in Theorem 2.6 is $\kappa_0 \hat{g}(x^{**}) = \lambda/r$, which is exactly the statement that W is continuously differentiable at x^{**}. Since $\hat{g}$ is affine and the exponential branch is strictly convex and tangent to it at x^{**}, $W \geq \hat{g}$ on $\mathbb{R}$. On $(-\infty, x^{**})$, $(\mathcal{L}_D - r)W = 0$ by construction, and on (x^{**}, ∞)

$$(\mathcal{L}_D - r)\hat{g}(x) = \alpha\frac{\lambda}{r} - r\hat{g}(x) = rC - \lambda x < 0,$$

because $x > x^{**} > Cr/\lambda$, the last inequality following from $1/\kappa_0 > \alpha/r$, which holds for all $r, \sigma > 0$. Hence W is an r-superharmonic majorant of $\hat{g}$, so $e^{-rt}W(D_t)$ is a supermartingale and optional stopping gives $\mathbb{E}_x[e^{-r\tau}\hat{g}(D_\tau)] \leq W(x)$ for every $\tau \in \mathcal{T}$, that is $V \leq W$. The first-passage rule at level x^{**} attains W, so $V = W$ and in particular $V = \hat{g}$ on $[x^{**}, \infty)$. Therefore Γ is unbounded above and no such ℓ exists. Hence τ^* is an optimal stopping time.

B.3 Proof of Theorem 2.14

Assuming that the jump process $Y(t)$ is independent of the diffusion $D(t)$, that $\nu\,\mathbb{E}|\xi| < \infty$ for the jump sizes ξ, and that $x^* \geq D_0$, rearranging Equation (2.53) gives

$$
\begin{aligned}
V_{D_{x^*},C} =& \mathbb{E}_{D_0,C_0}\left\{\mathbb{E}_{D_0,C_0}\left[-e^{-rD_{x^*}}\,C(D_{x^*}) + \int_{D_{x^*}}^{\infty} e^{-rt}\lambda\,D(t)\,dt\,\middle|\,D_{x^*}=s\right]\right\}\\
=& \int_0^{\infty} f_{D_{x^*}}(s)\,\mathbb{E}_{D_0,C_0}\Bigg[-e^{-rs}\,C(s)\\
&+\int_s^{\infty} e^{-rt}\lambda\left(x^* + \alpha\,(t-s) + \sigma\,\hat{W}(t-s)\right)dt\Bigg]ds\\
=& \int_0^{\infty} f_{D_{x^*}}(s)\Bigg[-e^{-rs}(C_0 + \mathbb{E}_{D_0,C_0}Y(s)) +\\
&\mathbb{E}_{D_0,C_0}\Big(\int_s^{\infty} e^{-rt}\lambda\left(x^* + \alpha\,(t-s) + \sigma\,\hat{W}(t-s)\right)dt\Big)\Bigg]ds,
\end{aligned}
\tag{B.8}
$$

where $f_{D_{x^*}}(s)$ is the density of the hitting time D_{x^*}, and, by the strong Markov property, $\{\hat{W}(u)\}_{u\geq 0}$ is a standard Brownian motion started at zero and independent of $\mathcal{F}_{D_{x^*}}$. The conditioning on $D_{x^*} = s$ leaves the law of $Y(s)$ unchanged because Y is independent of D. Note that at the hitting time we have $D(s) = x^*$. For $s \geq 0$ we have

$$\mathbb{E}_{D_0,C_0}[Y(s)] = \nu\,d_e\,s, \tag{B.9}$$

where ν is the intensity of the Poisson process and d_e is the average jump size.

Next, using the arguments in Section 2.2, we have

$$\mathbb{E}_{D_0,C_0}\int_s^{\infty} e^{-rt}\lambda\left(x^* + \alpha\,(t-s) + \sigma\,\hat{W}(t-s)\right)dt = e^{-rs}\Big(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2}\Big)$$

where the Brownian term vanishes by Fubini's theorem and $\mathbb{E}\,\hat{W}(u) = 0$ for every $u \geq 0$.

Equation (B.8) can be written as

$$V_{D_{x^*},C} = \int_0^{\infty} f_{D_{x^*}}(s)e^{-rs}\left[\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C_0 - \nu\,d_e\,s\right]ds. \tag{B.10}$$

Next, we define the Laplace transform of the stopping time D_{x^*} as

$$\mathcal{L}(r) = \mathbb{E}_{D_0,C_0}(e^{-rD_{x^*}}) = \int_0^{\infty} e^{-rs}f_{D_{x^*}}(s)\,ds. \tag{B.11}$$

Noting that

$$\frac{d}{dr}\mathcal{L}(r) = -\int_0^{\infty} s\,e^{-rs}f_{D_{x^*}}(s)\,ds,\ ^{1}$$

Equation (B.10) becomes

$$V_{D_{x^*},C} = \left[\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C_0\right]\mathcal{L}(r) + \nu\, d_e \frac{d}{dr}\mathcal{L}(r). \tag{B.12}$$

By the standard first-passage Laplace transform for arithmetic Brownian motion (Borodin and Salminen, 2015, p. 295), as in (2.15), we arrive at

$$\mathcal{L}(r) = \mathbb{E}(e^{-rD_{x^*}}) = \exp\left\{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma}(x^* - D_0) - |x^* - D_0|\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right)\right\}. \tag{B.13}$$

Since $x^* \geq D_0$ by hypothesis, $|x^* - D_0| = x^* - D_0$. Differentiating with respect to r, we obtain

$$\frac{d}{dr}\mathcal{L}(r) = \mathcal{L}(r)\cdot\frac{d}{dr}\left[-\frac{x^* - D_0}{\sigma}\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right] = -\mathcal{L}(r)\,\frac{x^* - D_0}{\sigma}\,\frac{1}{\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}.$$

Substituting this result into Equation (B.12), the value function becomes

$$V_{D_{x^*},C} = \left[\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - C_0 - \frac{\nu\, d_e\,(x^* - D_0)}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}\right] \times \exp\left\{\frac{1}{\sigma}\left(\frac{\alpha}{\sigma}(x^* - D_0) - (x^* - D_0)\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\right)\right\}. \tag{B.14}$$

The total value is subsequently modified through the application of the Laplace transform $\mathcal{L}(r)$, which accounts for both the discounting effects and the uncertainty related to the stopping time D_{x^*}.

B.4 Proof of Theorem 2.15

Proof. The formula (B.14) is valid on the branch $x^* \geq D_0$, and the supremum is accordingly taken over $x^* \geq D_0$; thresholds $x^* < D_0$ are hit immediately and yield the immediate-stopping value. On this branch the value function in Equation (B.14) can be rewritten as

$$\hat{V}_{D_{x^*},C} = \left[\tilde{A}\,x^* + \tilde{B}\right] e^{G(x^*+H)}, \tag{B.15}$$

where

$$\tilde{A} = \frac{\lambda}{r} - \frac{\nu\, d_e}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}},$$

[1]The interchange of differentiation and integration is justified by the Leibniz integral rule. Since $f_{D_{x^*}}(s)$ is a density function, the integrals involved converge and the integrand $e^{-rs} f_{D_{x^*}}(s)$ is differentiable with respect to r. Moreover, its derivative is uniformly dominated by an integrable function, ensuring the validity of differentiating under the integral sign.

$$\tilde{B} = \frac{\lambda\alpha}{r^2} - C_0 + \frac{\nu\, d_e\, D_0}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}},$$

and, as before, we set

$$G = \frac{1}{\sigma}\Big(\frac{\alpha}{\sigma} - \sqrt{\frac{\alpha^2}{\sigma^2} + 2r}\Big), \qquad H = -D_0.$$

To find the optimal threshold we differentiate $\hat{V}_{D_{x^*},C}$ with respect to x^*. Noting that the exponential function is always positive, the first derivative is given by

$$\frac{d\hat{V}_{D_{x^*},C}}{dx^*} = \Big(\tilde{A} + G\big[\tilde{A}\,x^* + \tilde{B}\big]\Big)\, e^{G(x^*+H)}. \tag{B.16}$$

Setting $\frac{d\hat{V}_{D_{x^*},C}}{dx^*} = 0$ is equivalent to requiring

$$\tilde{A} + G\big[\tilde{A}\,x^* + \tilde{B}\big] = 0. \tag{B.17}$$

Solving for x^{***} gives

$$\begin{aligned} \tilde{A}\,x^{***} + \tilde{B} &= -\frac{\tilde{A}}{G}, \\ x^{***} &= -\frac{1}{G} - \frac{\tilde{B}}{\tilde{A}}. \end{aligned} \tag{B.18}$$

The second-order derivative should be negative at x^{***}, which reads as

$$\frac{d^2\hat{V}_{D_{x^*},C}}{d^2x^*}(x^{***}) < 0 \tag{B.19}$$

In order to show this property we argue along the lines of the proof of Theorem 2.6. Note that for $r > 0$: $G < 0$ holds for any parameter combination. As $\tilde{A} > 0$, hence $G\tilde{A} < 0$, and as a result, $(\tilde{A}Gx^* + \tilde{A} + \tilde{B}G)$ is a decreasing function of x^*.

Therefore, for $x^* < x^{***}$, we have that $(\tilde{A}Gx^* + \tilde{A} + \tilde{B}G)$ is positive, and for $x^* > x^{***}$, we have that $(\tilde{A}Gx^* + \tilde{A} + \tilde{B}G)$ is negative. Along with the positive property of $e^{G(x^{***}+H)}$, we can conclude that $\frac{d\hat{V}_{D_{x^*},C}}{dx^*}$ changes from positive to negative at x^{***}, meaning that $\frac{d^2\hat{V}_{D_{x^*},C}}{d^2x^*}(x^{***}) < 0$ and x^{***} is a local maximum of $\hat{V}_{D_{x^*},C}$. As $\hat{V}_{D_{x^*},C}$ has a unique stationary point, x^{***} is also a global maximum.

Substituting back the definitions of G, $\tilde{A}$, and $\tilde{B}$, we obtain

$$x^{***} = \frac{\sigma^2}{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha} - \frac{\dfrac{\lambda\alpha}{r^2} - C_0 + \dfrac{\nu\, d_e\, D_0}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}}{\dfrac{\lambda}{r} - \dfrac{\nu\, d_e}{\sigma\sqrt{\frac{\alpha^2}{\sigma^2} + 2r}}}. \tag{B.20}$$

On the other hand, $\hat{V}_{D_{x^*},C}$ is strictly decreasing for $x^* > x^{***}$. Hence, when $x^{***} \leq D_0$, every admissible threshold $x^* \geq D_0$ lies on the decreasing branch, delaying the stopping decision would only decrease the expected payoff, and it is optimal to stop immediately at D_0. □

B.5 Sketch of the proof assumptions for the boundary conjecture

This appendix outlines, at a conceptual level, the analytical steps and assumptions that would be required to establish a formal proof of the monotone boundary conjecture stated in Subsection 3.1.1. No complete proof is attempted here; the discussion is intended to clarify the relationship between our conjectured structure and existing results in the theory of optimal stopping for multidimensional diffusions.

Given a continuous strong Markov process (C_t, D_t) and a measurable reward function $\hat{g}(c, d)$ of at most polynomial growth, standard results (see, e.g., Peskir and Shiryaev (2006, Ch. III)) guarantee the existence of a value function

$$V(c, d) = \sup_{\tau \geq 0} \mathbb{E}_{c,d}[e^{-r\tau} \hat{g}(C_\tau, D_\tau)],$$

and an optimal stopping time of first-entry type,

$$\tau^\star = \inf\{t \geq 0 : (C_t, D_t) \in \mathcal{S}\},$$

where $\mathcal{S} = \{(c, d) : V(c, d) = \hat{g}(c, d)\}$ is a closed stopping region. This part of the argument requires only the continuity and nondegeneracy of the diffusion coefficients of (C_t, D_t).

To obtain a threshold-type structure $\mathcal{S} = \{(c, d) : d \geq b(c)\}$, stronger assumptions are required:

(i) The reward $\hat{g}(c, d)$ must be increasing in d and decreasing in c.

(ii) The infinitesimal generator $\mathcal{L}$ of (C_t, D_t) must preserve these monotonicities (comparison property).

(iii) The diffusion must be nondegenerate and possess smooth transition densities.

(iv) The boundary $\partial\mathcal{S}$ must be regular in the sense of stochastic potential theory (no flat portions or cusps).

Under these conditions, the value function V inherits the monotonicity of $\hat{g}$, which implies that for each c, the set $\{d : (c, d) \in \mathcal{S}\}$ is an interval $[b(c), \infty)$ for some nondecreasing function $b(\cdot)$.

Further smoothness results (e.g., optimal regularity of V in $C^{1,2}$) would follow from elliptic PDE theory applied to the variational inequality

$$\max\{(\mathcal{L} - r)V,\ \hat{g} - V\} = 0.$$

Existence of smooth fit ($\nabla V = \nabla \hat{g}$ on $\partial\mathcal{S}$) requires additional conditions, including nondegeneracy of the diffusion and bounded growth of $\hat{g}$.

The OU–ABM process used in this thesis satisfies the Markov and continuity conditions but not all the regularity assumptions required for a full proof. Consequently, the monotone-boundary structure is treated as a *conjecture*, supported by the economic monotonicity of $\hat{g}(c,d)$ and by numerical evidence presented in Subsection 3.1.1. Establishing a rigorous proof would require extending known one-dimensional methods (e.g., the comparison principle or excessive-function approach) to this two-dimensional diffusion, which remains an open problem in the literature.

B.6 Optimal Stopping Foundations

This section collects the standard optimal stopping results underlying the variational inequality formulation in Chapter 3. The definitions and theorems follow Peskir and Shiryaev (2006) and Oksendal (2013).

Definition B.1 (Admissible Stopping Time)**.** A random variable $\tau : \Omega \to [0,T]$ is a *stopping time* with respect to $(\mathcal{F}_t)_{t\geq 0}$ if $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \in [0,T]$. The set of all such stopping times starting from t is denoted $\mathcal{T}_{t,T}$.

Definition B.2 (General Optimal Stopping Problem)**.** For a payoff function $\Pi : [0,T] \times E \to \mathbb{R}$, the value function is

$$V(t,x) = \sup_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}\left[e^{-r(\tau-t)}\Pi(\tau, X_\tau) \mid X_t = x\right], \tag{B.21}$$

where $r > 0$ is the discount rate and X_t satisfies standard integrability conditions.

The global optimisation decomposes into local decisions via the *Dynamic Programming Principle* (DPP).

Theorem B.3 (Dynamic Programming Principle)**.** *For any $\sigma \in \mathcal{T}_{t,T}$, the value function satisfies:*

$$V(t,x) = \sup_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}\left[e^{-r(\tau\wedge\sigma-t)}V(\tau \wedge \sigma, X_{\tau\wedge\sigma}) \mid X_t = x\right]. \tag{B.22}$$

For an infinitesimal increment $\Delta t > 0$, the decision-maker chooses between immediate exercise and continuation:

$$V(t,x) = \max\left\{\Pi(t,x), \mathbb{E}\left[e^{-r\Delta t}V(t+\Delta t, X_{t+\Delta t}) \mid X_t = x\right]\right\}. \tag{B.23}$$

Theorem B.4 (General Verification)**.** *Let $v(t,x) \in C^{1,2}([0,T] \times E)$ satisfy*

$$\min\left\{rv - \frac{\partial v}{\partial t} - \mathcal{L}v,\ v - \Pi\right\} = 0, \quad \forall (t,x) \in [0,T] \times E, \tag{B.24}$$

where $\mathcal{L}$ denotes the infinitesimal generator of the Markov process X (Oksendal, 2013). For the two-factor diffusion of Chapter 3 this is exactly the operator of equation (3.6)*. Then $v(t,x) = V(t,x)$ and the optimal stopping time is $\tau^* = \inf\{s \in [t,T] : v(s,X_s) = \Pi(s,X_s)\}$.*

B.6.1 Excess Value Representation for the Two-Factor Model

To analyse the structure of the value function, we separate the perpetual benefit that would accrue under immediate adoption from the residual value associated with the option to wait. We therefore define the *excess value*

$$u(c,d) := V(c,d) - \varphi(d),$$

where $\varphi(d)$ denotes the expected present value of the perpetual benefit stream generated by the arithmetic Brownian motion D_t:

$$\varphi(d) = \mathbb{E}\left[\int_0^\infty e^{-rs}\,\lambda D_s\,ds \,\middle|\, D_0 = d\right] = \lambda\left(\frac{d}{r} + \frac{\alpha}{r^2}\right), \tag{B.25}$$

Substituting this definition into the value function (3.5), we obtain:

$$u(c,d) = \sup_{\tau \geq 0}\ \mathbb{E}_{c,d}\left[-e^{-r\tau}C_\tau - \int_0^\tau e^{-rs}\,\lambda D_s\,ds\right], \tag{B.26}$$

where $u(c,d)$ represents the additional value the firm derives from optimally timing its adoption decision, above the value of adopting immediately.

To develop an explicit expression for (B.26), consider the discounted cost process $f(t, C_t) = e^{-rt}C_t$. Applying Ito's lemma (see, e.g., Oksendal (2013)) to f under the OU dynamics (3.2) yields

$$d\big(e^{-rt}C_t\big) = e^{-rt}\left[\kappa(\theta - C_t) - rC_t\right]dt + e^{-rt}\eta\,dW_t^C.$$

Integrating from 0 to a stopping time τ and taking expectations gives

$$\mathbb{E}_{c,d}[e^{-r\tau}C_\tau] = c \ + \ \mathbb{E}_{c,d}\left[\int_0^\tau e^{-rs}\big(\kappa\theta - (\kappa + r)C_s\big)\,ds\right], \tag{B.27}$$

where the stochastic integral $\int_0^\tau e^{-rs}\eta\,dW_s^C$ has zero mean by the optional stopping theorem for square-integrable martingales (see, e.g., Oksendal (2013)).[2]

Substituting (B.27) into (B.26) yields:

$$u(c,d) = -c \ + \ \sup_{\tau \geq 0}\ \mathbb{E}_{c,d}\left[\int_0^\tau e^{-rs}\,\big((\kappa + r)C_s - \kappa\theta - \lambda D_s\big)\,ds\right], \tag{B.28}$$

This representation (B.28) isolates the driver of the timing decision: the integrand $(\kappa + r)C_s - \kappa\theta - \lambda D_s$ gives the instantaneous balance between the expected change in discounted cost and the forgone flow of usage-based benefits.

The decomposition $V = \varphi(d) + u(c,d)$ clarifies the economic structure of the problem. In the numerical implementation developed in Subsection 3.4.2, however, we work *directly* with the value function $V(c,d)$ and the immediate-adoption payoff

$$\hat{g}(c,d) = \frac{\lambda d}{r} + \frac{\lambda\alpha}{r^2} - c,$$

treating the stopping region as $\mathcal{S} = \{(c,d) : V(c,d) = \hat{g}(c,d)\}$ and the continuation region as its complement.

[2] The process $M_t := \int_0^t e^{-rs}\eta\,dW_s^C$ is square-integrable because $\int_0^\infty e^{-2rs}\eta^2\,ds = \eta^2/(2r) < \infty$. For admissible stopping times with $\mathbb{E}\big[\int_0^\tau e^{-2rs}\,ds\big] < \infty$ (ensured by exponential discounting), we have $\mathbb{E}_{c,d}[M_\tau] = 0$; cf. Oksendal (2013, Thm. 4.6.1) or Karatzas and Shreve (2014, Prop. 3.2.10).

B.7 Finite-horizon immediate-stop benchmark

For $dD_t = \alpha\, dt + \sigma\, dW_t$ and horizon T, the discounted benefit of stopping immediately at t is

$$g_T(d,t) \;=\; -C \;+\; \lambda\left[\frac{1-e^{-r(T-t)}}{r}\,d \;+\; \frac{1-e^{-r(T-t)}\big(r(T-t)+1\big)}{r^2}\,\alpha\right],$$

which converges to $g_\infty(d) = -C + \lambda\big(d/r + \alpha/r^2\big)$ as $T \to \infty$. We use g_∞ in the stationary two-factor model and set $u := V - g_\infty$ for the homogeneous VI and boundary conditions.

B.8 Derivations for Deterministic–Cost Optimal Stopping

This appendix provides the analytical details underlying the deterministic–cost results in Section 3.2. Throughout, the benefit process evolves according to the arithmetic Brownian motion

$$dD_t = \alpha\, dt + \sigma\, dW_t, \qquad D_0 = d,$$

and the firm exercises the option when D_t hits the level x^*:

$$\tau_{x^*} := \inf\{t \geq 0 : D_t \geq x^*\}.$$

B.8.1 Laplace Transform of the First Passage Time

Let τ_{x^*} be as above, and assume $x^* > d$.[3] For a drifted Brownian motion,

$$D_t = d + \alpha t + \sigma W_t,$$

the Laplace transform of τ_{x^*} is given by (see Borodin and Salminen, 2015, p. 295):

$$\mathbb{E}_d[e^{-q\tau_{x^*}}] = \exp\left[-\frac{\sqrt{\alpha^2 + 2q\sigma^2} - \alpha}{\sigma^2}(x^* - d)\right], \qquad q > 0. \tag{B.29}$$

We define the abbreviations used in Section 3.2:

$$\kappa_0 := \frac{\sqrt{\alpha^2 + 2r\sigma^2} - \alpha}{\sigma^2}, \qquad \kappa_q := \frac{\sqrt{\alpha^2 + 2(r+q)\sigma^2} - \alpha}{\sigma^2}.$$

Then:

$$\mathbb{E}_d[e^{-r\tau_{x^*}}] = e^{-\kappa_0(x^*-d)}, \qquad \mathbb{E}_d[e^{-(r+q)\tau_{x^*}}] = e^{-\kappa_q(x^*-d)}.$$

[3] For $x^* \leq d$, the stopping time is zero and the value follows from immediate exercise.

B.8.2 Expected Calendar Time to Adoption

While the Laplace transform in Subsection B.8.1 characterises discounted expectations, it is also informative to consider the expected calendar time until adoption.

Let

$$\tau_{x^*} := \inf\{t \geq 0 : D_t \geq x^*\}, \qquad D_t = D_0 + \alpha t + \sigma W_t,$$

with $x^* > D_0$.

For $\alpha > 0$, the first-passage time τ_{x^*} follows an inverse Gaussian distribution, implying

$$\mathbb{E}_{D_0}(\tau_{x^*}) = \frac{x^* - D_0}{\alpha}. \tag{B.30}$$

This expression has a natural economic interpretation: the expected waiting time equals the deterministic drift time required to close the gap between the current benefit level D_0 and the adoption threshold x^*.

If $\alpha < 0$, the diffusion drifts downward and the boundary x^* is reached with probability less than one, so that

$$\mathbb{P}_{D_0}(\tau_{x^*} < \infty) < 1, \qquad \mathbb{E}_{D_0}(\tau_{x^*}) = \infty.$$

In the driftless case $\alpha = 0$, the boundary is reached almost surely, but the first-passage time has a heavy tail and $\mathbb{E}_{D_0}(\tau_{x^*}) = \infty$ nonetheless.

Consequently, meaningful expected waiting times arise only when the benefit process exhibits positive long-run growth.

B.8.3 Derivation of Finite Horizon Payoff Function

This appendix details the integration of the expected profit stream for the finite-horizon adoption problem. We consider a firm that adopts the technology at time τ, utilising it for a fixed operational lifespan L. The benefit process $D(t)$ follows an Arithmetic Brownian Motion (ABM):

$$dD_t = \alpha dt + \sigma dZ_t, \tag{B.31}$$

which implies that the expected value at future time $s \geq \tau$, conditional on information at τ, is given by:

$$\mathbb{E}_\tau[D(s)] = D_\tau + \alpha(s - \tau). \tag{B.32}$$

The total expected discounted profit Π_T is the integral of the benefit stream over the interval $[\tau, \tau + L]$, discounted back to the adoption time τ, minus the adoption cost C:

$$\Pi_T(D_\tau, C_\tau) = \lambda \int_\tau^{\tau+L} e^{-r(s-\tau)} \mathbb{E}_\tau[D(s)]\, ds - C_\tau. \tag{B.33}$$

Substituting the expectation $\mathbb{E}_\tau[D(s)]$ and applying the change of variable $u = s - \tau$ (where $du = ds$, with limits 0 to L):

$$\begin{aligned}\Pi_T &= \lambda \int_0^L e^{-ru}(D_\tau + \alpha u)\, du - C_\tau \\ &= \lambda D_\tau \underbrace{\int_0^L e^{-ru}\, du}_{I_1} + \lambda\alpha \underbrace{\int_0^L u e^{-ru}\, du}_{I_2} - C_\tau.\end{aligned} \tag{B.34}$$

Integral I_1 (Constant Benefit Component). This is a standard exponential integral:

$$I_1 = \left[-\frac{1}{r}e^{-ru}\right]_0^L = -\frac{1}{r}(e^{-rL} - 1) = \frac{1}{r}(1 - e^{-rL}). \tag{B.35}$$

Integral I_2 (Growth Component). We solve this using integration by parts. Let $v = u$ and $dw = e^{-ru}du$. Then $dv = du$ and $w = -\frac{1}{r}e^{-ru}$.

$$\begin{aligned}I_2 &= \int_0^L u e^{-ru}\, du = \left[-\frac{u}{r}e^{-ru}\right]_0^L - \int_0^L \left(-\frac{1}{r}e^{-ru}\right) du \\ &= \left(-\frac{L}{r}e^{-rL} - 0\right) + \frac{1}{r}\int_0^L e^{-ru}\, du \\ &= -\frac{L}{r}e^{-rL} + \frac{1}{r}\left[-\frac{1}{r}e^{-ru}\right]_0^L \\ &= -\frac{L}{r}e^{-rL} + \frac{1}{r^2}(1 - e^{-rL}).\end{aligned} \tag{B.36}$$

Rearranging terms to group by common factors:

$$I_2 = \frac{1}{r^2}\left(1 - e^{-rL} - rLe^{-rL}\right). \tag{B.37}$$

Final Expression. Substituting I_1 and I_2 back into the profit equation yields the final analytical form for the payoff function:

$$\Pi_T(D, C, \tau, L) = \frac{\lambda D}{r}\left(1 - e^{-rL}\right) + \frac{\lambda\alpha}{r^2}\left(1 - e^{-rL} - rLe^{-rL}\right) - C. \tag{B.38}$$

This formula captures both the capitalised value of the current benefit level (first term) and the value of the expected drift over the useful life of the asset (second term).

B.8.4 Expected Value Under a General Deterministic Cost Path

If the firm switches at x^*, it pays $C(\tau_{x^*})$ and obtains the perpetual benefit stream:

$$\varphi(x^*) = \mathbb{E}\left[\int_0^\infty e^{-rt}\lambda D_t\,dt \,\middle|\, D_0 = x^*\right] = \frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2}. \tag{B.39}$$

Hence the expected value of switching at level x^* is:

$$V(d;x^*) = \varphi(x^*)\,\mathbb{E}_d[e^{-r\tau_{x^*}}] \;-\; \mathbb{E}_d[e^{-r\tau_{x^*}}C(\tau_{x^*})]\,. \tag{B.40}$$

Closed-form solutions require $C(\tau_{x^*})$ to take a form compatible with the Laplace transform of τ_{x^*}. This happens exactly when:

$$C(t) = C_1 e^{-\beta t} + C_2, \tag{B.41}$$

since then:

$$\mathbb{E}_d[e^{-r\tau_{x^*}}C(\tau_{x^*})] = C_1\,\mathbb{E}_d[e^{-(r+\beta)\tau_{x^*}}] + C_2\,\mathbb{E}_d[e^{-r\tau_{x^*}}].$$

B.8.5 Deterministic OU–Mean Cost

The deterministic mean path of the OU process solves:

$$\frac{dC(t)}{dt} = \kappa(\theta - C(t)), \qquad C(0) = C_0,$$

with solution

$$C(t) = C_0 e^{-\kappa t} + \theta(1 - e^{-\kappa t}) = \theta + (C_0 - \theta)\,e^{-\kappa t}. \tag{B.42}$$

Thus in (B.41):

$$C_1 = C_0 - \theta, \qquad C_2 = \theta, \qquad \beta = \kappa.$$

Substituting into (B.40):

$$V(d;x^*) = \left(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2}\right) e^{-\kappa_0(x^*-d)} - (C_0 - \theta)\,e^{-\kappa_\kappa(x^*-d)} - \theta e^{-\kappa_0(x^*-d)}.$$

Grouping terms yields:

$$V(d;x^*) = \left(\frac{\lambda x^*}{r} + \frac{\lambda\alpha}{r^2} - \theta\right) e^{-\kappa_0(x^*-d)} \;-\; (C_0 - \theta)\,e^{-\kappa_\kappa(x^*-d)}, \tag{B.43}$$

matching equation (3.11) in the chapter.

The optimal switching threshold. To maximise $V(d;x^*)$, differentiate (B.43) with respect to x^*. Let:

$$A := \frac{\lambda}{r}, \qquad B := \frac{\lambda\alpha}{r^2} - \theta.$$

Then:

$$V(d;x^*) = (Ax^* + B)e^{-\kappa_0 X} - (C_0 - \theta)e^{-\kappa_\kappa X}, \qquad X := x^* - d.$$

Differentiate:

$$\frac{dV}{dx^*} = Ae^{-\kappa_0 X} - \kappa_0(Ax^* + B)e^{-\kappa_0 X} - (C_0 - \theta)(-\kappa_\kappa)e^{-\kappa_\kappa X}.$$

Set equal to zero:

$$A(\kappa_0 x^* + \kappa_0 B/A - 1)e^{-\kappa_0 X} = (C_0 - \theta)\kappa_\kappa e^{-\kappa_\kappa X}. \tag{B.44}$$

Using $B/A = \frac{\alpha}{r} - \frac{r\theta}{\lambda}$ and rearranging:

$$e^{-(\kappa_\kappa - \kappa_0)(x^{**} - d)} = \frac{\frac{\lambda}{r}\left(\kappa_0 x^{**} + \frac{\alpha}{r}\kappa_0 - 1\right) - \kappa_0\theta}{(C_0 - \theta)\kappa_\kappa}. \tag{B.45}$$

Thus the optimal threshold x^{**} solves:

$$e^{-(\kappa_\kappa - \kappa_0)(x^{**} - d)} = \frac{\frac{\lambda}{r}\left(\kappa_0 x^{**} - 1\right) + \frac{\lambda\alpha}{r^2}\kappa_0 - \kappa_0\theta}{(C_0 - \theta)\kappa_\kappa}. \tag{B.46}$$

This matches equation (3.13).

Limiting behaviour. In the limit $\kappa \to 0$ (no cost decline), $\kappa_\kappa \to \kappa_0$, the exponential factor tends to one, and the first-order condition reduces to the constant-cost threshold $x^{**} = 1/\kappa_0 - \alpha/r + rC_0/\lambda$. For $\kappa > 0$ the equation is transcendental; its exact solution is derived next.

B.8.6 Explicit solution of the threshold via the Lambert $\mathcal{W}$ function

This subsection derives the closed form (3.16) quoted in Section 3.2. The first-order condition (3.13) is transcendental, the threshold appearing both linearly and in an exponent, but it can be put in the canonical form $we^w = \zeta$ solved by the Lambert $\mathcal{W}$ function, the inverse of $w \mapsto we^w$ (Corless et al., 1996).

Write the first-order condition (3.13) with the exponential on the left,

$$(C_0 - \theta)\,\kappa_\kappa\, e^{-\Delta\kappa\,(x^{**} - d)} = \frac{\lambda}{r}\big(\kappa_0 x^{**} - 1\big) + \frac{\lambda\alpha}{r^2}\kappa_0 - \kappa_0\theta, \qquad \Delta\kappa := \kappa_\kappa - \kappa_0. \tag{B.47}$$

The right-hand side is linear in x^{**} and factors as

$$\frac{\lambda\kappa_0}{r}\big(x^{**} - \tilde{x}\big), \qquad \tilde{x} := x_{\text{base}} + \frac{r\theta}{\lambda}, \qquad x_{\text{base}} = \frac{1}{\kappa_0} - \frac{\alpha}{r}, \tag{B.48}$$

as expanding the bracket confirms: $\frac{\lambda\kappa_0}{r}x^{**} - \frac{\lambda\kappa_0}{r}\big(\frac{1}{\kappa_0} - \frac{\alpha}{r}\big) - \kappa_0\theta$ is the right-hand side of (B.47). The threshold therefore enters only through its distance from $\tilde{x}$, which suggests the substitution

$$w := \Delta\kappa\,\big(x^{**} - \tilde{x}\big), \qquad \text{so that} \qquad x^{**} - \tilde{x} = \frac{w}{\Delta\kappa}. \tag{B.49}$$

The substitution removes $\tilde{x}$ from the linear factor, leaving w alone on the right-hand side. Splitting the exponent of (B.47) around $\tilde{x}$,

$$-\Delta\kappa\,(x^{**} - d) = -\Delta\kappa\,(x^{**} - \tilde{x}) - \Delta\kappa\,(\tilde{x} - d) = -w + \Delta\kappa\,(d - \tilde{x}),$$

so that $e^{-\Delta\kappa(x^{**}-d)} = e^{-w}e^{\Delta\kappa(d-\tilde{x})}$. Substituting this and (B.48) into (B.47),

$$(C_0 - \theta)\,\kappa_\kappa\, e^{\Delta\kappa(d-\tilde{x})}\, e^{-w} = \frac{\lambda\kappa_0}{r}\cdot\frac{w}{\Delta\kappa}.$$

Multiplying both sides by $e^{w}\, r\,\Delta\kappa/(\lambda\kappa_0)$ isolates the canonical form,

$$w\,e^{w} = \zeta, \qquad \zeta := \frac{(C_0 - \theta)\,\kappa_\kappa\,\Delta\kappa\, r}{\lambda\kappa_0}\, e^{\Delta\kappa\,(d-\tilde{x})}, \tag{B.50}$$

whence $w = \mathcal{W}(\zeta)$ and, undoing (B.49),

$$x^{**} = \tilde{x} + \frac{1}{\Delta\kappa}\,\mathcal{W}\!\left(\frac{(C_0 - \theta)\,\kappa_\kappa\,\Delta\kappa\, r}{\lambda\kappa_0}\, e^{\Delta\kappa\,(d-\tilde{x})}\right). \tag{B.51}$$

Setting $\theta = 0$ gives $\tilde{x} = x_{\text{base}}$ and reduces (B.51) to (3.16), the pure exponential-decay case quoted in the main text. The derivation shows that the closed form is not special to $\theta = 0$: it holds for the general deterministic mean path, with the cost premium $C_0 - \theta$ in place of C_0 and the base threshold displaced by $r\theta/\lambda$.

Two remarks on the branch of $\mathcal{W}$. For $\kappa > 0$ we have $\kappa_\kappa > \kappa_0$, so $\Delta\kappa > 0$, and whenever the initial cost exceeds its long-run target, $C_0 > \theta$, the argument ζ in (B.50) is strictly positive. On $\zeta > 0$ the equation $we^{w} = \zeta$ has exactly one real solution, the principal branch $\mathcal{W}_0$, so the threshold is unambiguous and no branch selection is required. Since $\mathcal{W}_0(\zeta) > 0$ there, $x^{**} > \tilde{x}$: the anticipated cost decline raises the target above the level that would be optimal against a cost frozen at θ.

B.9 Optimal Stopping under Jump-Diffusion: The Single-Factor Reduction

This section provides the constructive proof underlying the closed-loop results of Subsection 2.3.1 (Theorem 2.11), resolving the suboptimality of the precommitted rule identified in Remark 2.16. Throughout, the transformed state process is

$$Z(t) = \frac{\lambda}{r}D(t) + \frac{\lambda\alpha}{r^2} - C(t)$$

of (2.37), the immediate-adoption payoff at the prevailing state: the capitalised unit cost advantage, inclusive of its expected further growth, less the fixed adoption cost. Because the investment cost $C(t)$ is subject to a compound Poisson process representing discrete technological breakthroughs, $Z(t)$ follows a one-dimensional jump-diffusion process:

$$dZ(t) = \mu_Z dt + \sigma_Z dW(t) + dY(t) \tag{B.52}$$

where $\mu_Z = \frac{\lambda\alpha}{r}$, $\sigma_Z = \frac{\lambda\sigma}{r}$, and $Y(t)$ is a compound Poisson process with arrival rate ν, independent of the Brownian motion W as assumed in Section 2.3. That independence makes Z a Lévy process and its exponent (2.39) the sum of the exponents of the two components. Following the framework established by Kou (2002), the upward jump sizes j (corresponding to sudden reductions in $C(t)$) are independent and identically distributed, independent of the arrival times, and exponentially distributed with probability density function $f(j) = \gamma e^{-\gamma j}$ for $j > 0$.

By Proposition 2.13, which extends Theorem 1 of Mordecki (2002a) to the filtration of (C, D), the globally optimal strategy is to adopt the technology at the first passage time $\tau^* = \inf\{t \geq 0 : Z(t) \geq z^*\}$.

The following lemma records the root structure of the characteristic equation used in Subsection 2.3.1 and throughout this section. It adapts Lemma 2.1 of Kou and Wang (2003), stated there for double-exponential jumps and a quartic equation, to the present one-sided case, in which the equation is a cubic. The adaptation is not merely notational: the root count there relies on both jump rates being finite, and clearing the single remaining denominator here leaves a cubic rather than a quartic.

Lemma B.5 (Root structure of the characteristic equation)**.** *Let $\sigma_Z > 0$, $\nu > 0$, $\gamma > 0$, and let*

$$\psi_Z(\beta) = \mu_Z\beta + \tfrac{1}{2}\sigma_Z^2\beta^2 + \nu\left(\frac{\gamma}{\gamma-\beta} - 1\right), \qquad \beta \neq \gamma.$$

For every $r > 0$ the equation $\psi_Z(\beta) = r$ has exactly two positive roots β_1 and β_2, and they satisfy $0 < \beta_1 < \gamma < \beta_2$.

Proof. Consider first the interval $(0, \gamma)$. Differentiating twice gives

$$\psi_Z''(\beta) = \sigma_Z^2 + \frac{2\nu\gamma}{(\gamma-\beta)^3} > 0 \qquad \text{for } 0 < \beta < \gamma,$$

so ψ_Z is strictly convex there. Moreover $\psi_Z(0) = 0 < r$ and $\psi_Z(\beta) \to +\infty$ as $\beta \uparrow \gamma$, because the jump term diverges while the quadratic part stays bounded. By continuity a root exists in $(0, \gamma)$. By strict convexity the sublevel set $\{\beta \in [0, \gamma) : \psi_Z(\beta) \le r\}$ is an interval containing 0, and ψ_Z exceeds r beyond its right endpoint. The right endpoint is therefore the unique root $\beta_1 \in (0, \gamma)$.

Consider next the interval (γ, ∞). As $\beta \downarrow \gamma$ the jump term tends to $-\infty$, so $\psi_Z(\beta) \to -\infty$. As $\beta \to \infty$ the quadratic term dominates the jump term, which converges to $-\nu$, so $\psi_Z(\beta) \to +\infty$. By continuity at least one root exists in (γ, ∞).

To count the roots exactly, multiply $\psi_Z(\beta) = r$ through by $\gamma - \beta$, which is nonzero away from γ. The resulting equation is

$$q(\beta) := (\gamma - \beta)\Big(\tfrac{1}{2}\sigma_Z^2\beta^2 + \mu_Z\beta - (\nu + r)\Big) + \nu\gamma = 0,$$

a cubic polynomial with at most three real roots, and every root of $\psi_Z(\beta) = r$ is a root of q. Now $q(0) = -\gamma(\nu + r) + \nu\gamma = -r\gamma < 0$, while $q(\beta) \to +\infty$ as $\beta \to -\infty$, because both factors of the product term tend to $+\infty$. Hence q has a root in $(-\infty, 0)$. Together with the roots already located in $(0, \gamma)$ and (γ, ∞), this accounts for all three roots of the cubic. Consequently β_1 is the only root in $(0, \gamma)$ and β_2 is the only root in (γ, ∞), which establishes $0 < \beta_1 < \gamma < \beta_2$. Finally $\beta = \gamma$ itself is not a solution, since $q(\gamma) = \nu\gamma > 0$. □

B.9.1 Constructive verification of the optimal threshold

Remark B.6 (Regularity at the trigger). The verification below applies Itô's formula to a candidate value function that is C^2 in the continuation region but only continuous, not C^1, at the trigger when the trigger is not chosen optimally. In that case Itô's formula does not apply directly, and optional stopping is justified by a mollification and localisation argument, exactly as in the proof of Theorem 3.1 of Kou and Wang (2003). At the optimal trigger z^* smooth pasting restores C^1 regularity, and the argument below proceeds without modification.

Theorem 2.11 quotes the trigger from Mordecki (2002a). The classical route of pasting conditions reaches the same level by an argument that relies on neither that paper nor Kou and Wang (2003), and is recorded here as an independent derivation.

Proposition B.7 (Verification by pasting conditions). *Let β_1 and β_2, with $0 < \beta_1 < \gamma < \beta_2$, be the two positive roots of the characteristic equation*

$$r = \mu_Z\beta + \frac{1}{2}\sigma_Z^2\beta^2 + \nu\left(\frac{\gamma}{\gamma - \beta} - 1\right) \tag{B.53}$$

supplied by Lemma B.5. The trigger determined by value matching, smooth pasting, and the integro-differential pasting condition is

$$z^* = \frac{1}{\beta_1} + \frac{1}{\beta_2} - \frac{1}{\gamma}, \tag{B.54}$$

which coincides with the level $z^* = p/\beta_1 + (1-p)/\beta_2$ *of Theorem 2.11.*

Proof. The argument proceeds in three steps. The first constructs the Bellman integro-differential equation (IDE) on the continuation region and factors out the characteristic equation, whose roots are fixed by the exogenous parameters alone. The second imposes the boundary conditions: value matching and smooth pasting, which the continuous diffusion component requires, and a third integro-differential pasting condition accounting for jumps that cross the boundary without hitting it. The third solves the resulting system for the trigger and verifies the agreement with (2.41).

Step 1: The Integro-Differential Equation and General Solution
In the continuation region ($z < z^*$), the value function $V(z)$ must satisfy the Bellman equation $rV = \mathcal{L}_Z V$, where $\mathcal{L}_Z$ is the infinitesimal generator of the jump diffusion Z. The differential terms of $\mathcal{L}_Z$ form the diffusion generator and the integral term is the compensator of the compound Poisson component (see, e.g., Øksendal and Sulem, 2007). Written out, this is the integro-differential equation (IDE)

$$rV(z) = \mu_Z V'(z) + \frac{1}{2}\sigma_Z^2 V''(z) + \nu \int_0^\infty [V(z+j) - V(z)]\gamma e^{-\gamma j} dj \tag{B.55}$$

Assuming a solution of the form $e^{\beta z}$ yields the characteristic equation stated in Equation (B.53). Multiplying by $(\gamma - \beta)$ to clear the denominator yields a cubic polynomial in β. The negative root is discarded, so that the option value vanishes as $z \to -\infty$, leaving the two positive roots β_1 and β_2 of Lemma B.5. Because r, μ_Z, σ_Z, ν, and γ are fully exogenous parameters, these roots act as static structural constants, defining the shape of the value function independently of the firm's optimal stopping decision. By the principle of superposition, the general solution for the option value is:

$$V(z) = B_1 e^{\beta_1 z} + B_2 e^{\beta_2 z} \tag{B.56}$$

Step 2: Boundary Conditions and the Expected Overshoot
Because the state process contains a continuous diffusion component ($\sigma_Z > 0$), the boundary z^* is strictly regular for the stopping region (Mordecki, 2002a). Therefore, the continuous value-matching and smooth-pasting conditions apply:

$$B_1 e^{\beta_1 z^*} + B_2 e^{\beta_2 z^*} = z^* \tag{B.57}$$

$$B_1 \beta_1 e^{\beta_1 z^*} + B_2 \beta_2 e^{\beta_2 z^*} = 1 \tag{B.58}$$

However, the state variable can jump over the boundary z^* without hitting it exactly. To ensure the IDE holds for all $z < z^*$, the piecewise definition of $V(z)$ is substituted back into the integral term of Equation (B.55) and the integral partitioned at the jump distance that crosses the boundary ($j = z^* - z$):

$$I = \int_0^{z^*-z} [B_1 e^{\beta_1(z+j)} + B_2 e^{\beta_2(z+j)}]\gamma e^{-\gamma j} dj + \int_{z^*-z}^{\infty} [z + j]\gamma e^{-\gamma j} dj - V(z) \quad \text{(B.59)}$$

Evaluating the continuation integral yields terms proportional to $e^{\beta_1 z}$ and $e^{\beta_2 z}$, alongside residual terms evaluated at the upper limit:

$$I_{cont} = \sum_{i=1}^{2} B_i \left(\frac{\gamma}{\gamma - \beta_i}\right) e^{\beta_i z} - \left[\sum_{i=1}^{2} B_i \left(\frac{\gamma}{\gamma - \beta_i}\right) e^{\beta_i z^*}\right] e^{-\gamma(z^*-z)} \quad \text{(B.60)}$$

Evaluating the overshoot integral utilising the memoryless property of the exponential distribution yields:

$$I_{over} = \left[z^* + \frac{1}{\gamma}\right] e^{-\gamma(z^*-z)} \quad \text{(B.61)}$$

When substituted back into the IDE, the $e^{\beta_1 z}$ and $e^{\beta_2 z}$ terms perfectly cancel the differential operators because β_1 and β_2 are roots of the characteristic equation. The IDE holds identically for all $z < z^*$ if and only if the algebraic coefficient of the remaining $e^{-\gamma(z^*-z)}$ residual vanishes. Setting this expression to zero establishes the integro-differential pasting condition (Kou, 2002):

$$B_1 \left(\frac{\gamma}{\gamma - \beta_1}\right) e^{\beta_1 z^*} + B_2 \left(\frac{\gamma}{\gamma - \beta_2}\right) e^{\beta_2 z^*} = z^* + \frac{1}{\gamma} \quad \text{(B.62)}$$

Step 3: Solving for the Endogenous Decision Variables

Treating the roots β_1 and β_2 as known constants, the non-linear system defined by Equations (B.57), (B.58), and (B.62) contains exactly three endogenous unknowns: the option value multipliers B_1 and B_2, and the optimal stopping threshold z^*. In the variables $u_i := B_i e^{\beta_i z^*}$ the three conditions read

$$u_1 + u_2 = z^*, \quad \text{(B.63)}$$

$$\beta_1 u_1 + \beta_2 u_2 = 1, \quad \text{(B.64)}$$

$$\frac{\gamma}{\gamma - \beta_1} u_1 + \frac{\gamma}{\gamma - \beta_2} u_2 = z^* + \frac{1}{\gamma}. \quad \text{(B.65)}$$

Subtracting (B.63) from (B.65) and using $\gamma/(\gamma - \beta) - 1 = \beta/(\gamma - \beta)$ removes z^* from the pasting condition,

$$\frac{\beta_1}{\gamma - \beta_1} u_1 + \frac{\beta_2}{\gamma - \beta_2} u_2 = \frac{1}{\gamma}. \quad \text{(B.66)}$$

For a given z^*, the value-matching and smooth-pasting conditions (B.63) and (B.64) are linear in u_1 and u_2, with solution

$$u_1 = \frac{\beta_2 z^* - 1}{\beta_2 - \beta_1}, \qquad u_2 = \frac{1 - \beta_1 z^*}{\beta_2 - \beta_1}. \quad \text{(B.67)}$$

Substituting (B.67) into (B.66) and multiplying through by $\beta_2 - \beta_1$,

$$\begin{aligned}\frac{\beta_2 - \beta_1}{\gamma} &= \frac{\beta_1(\beta_2 z^* - 1)}{\gamma - \beta_1} + \frac{\beta_2(1 - \beta_1 z^*)}{\gamma - \beta_2} \\ &= \frac{\beta_1\beta_2(\beta_1 - \beta_2)\, z^* + \gamma(\beta_2 - \beta_1)}{(\gamma - \beta_1)(\gamma - \beta_2)},\end{aligned}$$

the second line placing the two fractions over a common denominator and collecting the terms in z^* separately from the constants. Cancelling the factor $\beta_2 - \beta_1$ and cross-multiplying,

$$\begin{aligned}\gamma\big(\gamma - \beta_1\beta_2 z^*\big) &= (\gamma - \beta_1)(\gamma - \beta_2) \\ \gamma^2 - \gamma\beta_1\beta_2 z^* &= \gamma^2 - \gamma(\beta_1 + \beta_2) + \beta_1\beta_2 \\ z^* &= \frac{\gamma(\beta_1 + \beta_2) - \beta_1\beta_2}{\gamma\beta_1\beta_2} = \frac{1}{\beta_1} + \frac{1}{\beta_2} - \frac{1}{\gamma},\end{aligned} \tag{B.68}$$

which is (B.54). The threshold is thus a function of the environmental roots alone. Mapping z^* back through the definition $Z(t) = \frac{\lambda}{r}D(t) + \frac{\lambda\alpha}{r^2} - C(t)$ and isolating $D(t)$ in $Z(t) \geq z^*$ returns the boundary $d^*(C(t))$ of (2.42), the constant $\frac{\lambda\alpha}{r^2}$ contributing the term $-\frac{\alpha}{r}$ after multiplication by $\frac{r}{\lambda}$.

It remains to check the agreement with Theorem 2.11, and this is the computation (2.44), which turns the mixture $p/\beta_1 + (1-p)/\beta_2$ of the two roots into $1/\beta_1 + 1/\beta_2 - 1/\gamma$. The pasting solution and the expected-supremum representation of Mordecki (2002a) therefore coincide, as they must. □

B.9.2 Directional effect of jump risk on the threshold

This subsection proves the two directional comparative statics stated in prose in Subsection 2.3.1. Both work through the representation (2.52) of the optimal level z^* as the integral of the discount factor Λ, so a change in the jump specification affects z^* only through its effect on Λ. Throughout, $z_0^* = 1/\beta_0$ denotes the no-jump threshold and β_0 the positive root of the exponent $\psi_{Z,0}(\beta) = \mu_Z\beta + \frac{1}{2}\sigma_Z^2\beta^2$ without jumps.

Proposition B.8 (Cost-reduction risk raises the threshold)**.** $z^* > z_0^*$, *so the adoption boundary* $d^*(c)$ *lies uniformly above its no-jump counterpart, shifted upward in parallel by* $(r/\lambda)(z^* - z_0^*) > 0$.

Proof. Write $Z_t^\nu = Z_t^0 + J_t$, where Z^0 is the Brownian component with drift and J is the independent compound Poisson process collecting the upward jumps. Since J has nonnegative increments, $Z_t^\nu \geq Z_t^0$ for all t along every path. For a distance $h > 0$, let τ_h^ν and τ_h^0 denote the times at which Z^ν and Z^0 first gain the distance h. The pathwise ordering gives $\tau_h^\nu \leq \tau_h^0$, and therefore $\Lambda^\nu(h) = \mathbb{E}[e^{-r\tau_h^\nu}] \geq \mathbb{E}[e^{-r\tau_h^0}] = \Lambda^0(h)$.

The inequality is strict for every $h > 0$. Consider the event that the first jump of J arrives before τ_h^0 and has size exceeding h. This event has positive probability, because the first arrival time is exponential with rate $\nu > 0$, the jump size exceeds h with probability $e^{-\gamma h} > 0$, $\tau_h^0 > 0$ almost surely, and the three ingredients are independent. On this event $\tau_h^\nu < \tau_h^0$ strictly, so $\Lambda^\nu(h) > \Lambda^0(h)$. Integrating the strict inequality over $h \in (0, \infty)$ in (2.52) gives $z^* > z_0^*$. The parallel shift of the boundary follows from (2.42), which is linear in z^* with coefficient r/λ. □

Proposition B.9 (Cost-increase risk lowers the threshold)**.** *Suppose instead that the jumps of the fixed cost are upward, so that Z has downward jumps of size $\xi > 0$ arriving at rate $\nu > 0$, where ξ may follow any distribution on $(0, \infty)$. Then $z^* = 1/\Phi_Z(r) < z_0^*$, where $\Phi_Z(r)$ is the unique positive root of the Lévy exponent equation $\psi_Z(\beta) = r$.*

Proof. With downward jumps the Lévy exponent becomes

$$\psi_Z(\beta) = \mu_Z\beta + \tfrac{1}{2}\sigma_Z^2\beta^2 + \nu\big(\mathbb{E}[e^{-\beta\xi}] - 1\big), \tag{B.69}$$

finite for every $\beta \geq 0$ because $\xi > 0$ makes $e^{-\beta\xi} \leq 1$, so no moment condition on ξ is required. The first two terms are strictly convex since $\sigma_Z > 0$, and $\beta \mapsto \mathbb{E}[e^{-\beta\xi}]$ is convex, being an average of the convex functions $\beta \mapsto e^{-\beta x}$; hence ψ_Z is strictly convex on $[0, \infty)$. It satisfies $\psi_Z(0) = 0$, and since the jump term is confined to $[-\nu, 0]$ while the quadratic term diverges, $\psi_Z(\beta) \to \infty$ as $\beta \to \infty$. The equation $\psi_Z(\beta) = r$ therefore has a unique positive root $\Phi_Z(r)$ for every $r > 0$.

Two specifications are used later. Exponential jumps $\xi \sim \text{Exp}(\gamma)$ give $\mathbb{E}[e^{-\beta\xi}] = \gamma/(\gamma + \beta)$, the mirror image of the exponent (2.39) of the cost-reduction case. A deterministic jump of fixed size $d_e > 0$, the specification calibrated in Scenario E of Subsection 2.4.1, gives $\mathbb{E}[e^{-\beta\xi}] = e^{-\beta d_e}$. Neither the argument above nor the one below distinguishes between the two.

The process Z has no upward jumps, so it crosses any level from below continuously and there is no overshoot. The discount factor of a gain of distance h is a single exponential, $\Lambda(h) = e^{-\Phi_Z(r)h}$. To verify this in detail, note that $M_t = \exp\big(\Phi_Z(r)(Z_t - Z_0) - rt\big)$ is a martingale, because $\mathbb{E}[e^{\Phi_Z(r)(Z_t - Z_0)}] = e^{t\psi_Z(\Phi_Z(r))} = e^{rt}$. Up to the passage time τ_h the process $Z_t - Z_0$ never exceeds h, so $0 \leq M_{t\wedge\tau_h} \leq e^{\Phi_Z(r)h}$ and the stopped martingale is bounded. Optional stopping therefore applies and gives $\mathbb{E}[e^{\Phi_Z(r)h - r\tau_h}\mathbf{1}\{\tau_h < \infty\}] = 1$, that is $\Lambda(h) = e^{-\Phi_Z(r)h}$. Substituting into (2.52) yields $z^* = \int_0^\infty e^{-\Phi_Z(r)u}\,\mathrm{d}u = 1/\Phi_Z(r)$.

For the comparison with the no-jump threshold, subtract the no-jump exponent from (B.69) to obtain, for every $\beta > 0$,

$$\psi_Z(\beta) - \psi_{Z,0}(\beta) = -\nu\big(1 - \mathbb{E}[e^{-\beta\xi}]\big) < 0,$$

the inequality because $\xi > 0$ almost surely gives $e^{-\beta\xi} < 1$ pathwise, hence $\mathbb{E}[e^{-\beta\xi}] < 1$. (For exponential jumps the difference is $-\nu\beta/(\gamma+\beta)$, and for a deterministic jump $-\nu(1 - e^{-\beta d_e})$.) So $\psi_Z(\beta) < \psi_{Z,0}(\beta)$ for every $\beta > 0$. Evaluating at the no-jump root β_0 gives $\psi_Z(\beta_0) < \psi_{Z,0}(\beta_0) = r$. Since ψ_Z is convex with $\psi_Z(0) = 0$, the set

$\{\beta \geq 0 : \psi_Z(\beta) \leq r\}$ is the interval $[0, \Phi_Z(r)]$, and $\psi_Z(\beta_0) < r$ places β_0 strictly inside it. Hence $\Phi_Z(r) > \beta_0$ and $z^* = 1/\Phi_Z(r) < 1/\beta_0 = z_0^*$. □

B.9.3 Analytical Comparison of Adoption Boundaries

To isolate the structural evolution of the investment timing, Table B.1 presents the explicit operational thresholds required to trigger technology adoption under the three modelling frameworks.

Modelling Framework	**Adoption Threshold (Required Unit Cost Advantage)**
Constant Cost	$x^{**} = \frac{r}{\lambda} C_0 + \frac{1}{\beta_D} - \frac{\alpha}{r}$
Precommitted Benchmark	$x^{***} = \frac{1}{\beta_D} - \frac{\frac{\lambda\alpha}{r^2} - C_0 + \epsilon_1}{\frac{\lambda}{r} - \epsilon_2}$
Closed-Loop Optimum	$d^*(C(t)) = \frac{r}{\lambda} C(t) + \frac{r}{\lambda}\left(\frac{1}{\beta_1} + \frac{1}{\beta_2} - \frac{1}{\gamma}\right) - \frac{\alpha}{r}$

Table B.1: Comparison of the adoption thresholds. β_D is the positive root of the characteristic equation of the benefit diffusion, $\frac{1}{2}\sigma^2\beta^2 + \alpha\beta = r$, so that $1/\beta_D = \sigma^2/(\sqrt{\alpha^2 + 2r\sigma^2} - \alpha)$. It is the counterpart in the D coordinate of the no-jump root β_0 of (2.47), the two being related by $\beta_0 = (r/\lambda)\beta_D$. The abbreviations $\epsilon_1 = JD_0$ and $\epsilon_2 = J$ with $J = \nu d_e/(\sigma\sqrt{\alpha^2/\sigma^2 + 2r})$ collect the jump adjustment of the precommitted rule. β_1 and β_2 are the two positive roots $0 < \beta_1 < \gamma < \beta_2$ of the cubic characteristic equation (2.45) of the transformed jump-diffusion process Z, whose third root is negative, and γ is the parameter of the exponentially distributed cost reductions.

The economic content of Table B.1 is discussed in Subsections 2.3.1 and 2.3.3. Only the structural distinction is recorded here. The constant-cost threshold x^{**} and the precommitted threshold x^{***} are static triggers fixed at time zero, the latter anchored to the initial cost state C_0 and the initial advantage D_0. The closed-loop boundary $d^*(C(t))$ is a function of the prevailing cost: a realised cost reduction shifts the hurdle on D down by $(r/\lambda)|\Delta C|$ at the instant it occurs, while the intercept term $\frac{1}{\beta_1} + \frac{1}{\beta_2} - \frac{1}{\gamma}$ prices the option to wait for further jumps, including the expected overshoot of the boundary.

Appendix C

NUMERICAL IMPLEMENTATION DETAILS

This chapter tabulates numerical parameters referenced in Section 3.4 and Subsection 3.4.2. Prose derivations and algorithmic motivation appear there. The tables provide a self-contained record of the specifications used in the computations.[1]

C.1 Grid Specifications

Table C.1: Computational domain bounds. The domain spans approximately two to five standard deviations from the baseline state $(C_0, D_0) = (5{,}800,\ 1.45)$.

State variable	Lower	Upper	Units
Fixed Cost C	−8,000	12,000	\$/vehicle
Unit-cost advantage D (baseline)	−30	30	\$/100 mi
Unit-cost advantage D (Scenario H)	−30	60	\$/100 mi

C.2 Positive Coefficients and the Grid Aspect Ratio

Convergence of PSOR to the unique discrete solution, and convergence of that solution to the viscosity solution of the variational inequality, both rest on the discrete operator being monotone, which for the seven-point stencil of Subsection 3.4.2 means that every off-diagonal weight of $(\mathcal{L}_h - r)$ is non-negative. Collecting the contributions of the central second differences, the upwind drift terms and the mixed derivative, the four axial weights at an interior node are

$$a_E = \frac{\eta^2}{2\Delta C^2} + \frac{\max\{\mu_C, 0\}}{\Delta C} - \frac{|\rho|\eta\sigma}{2\Delta C \Delta D}, \qquad a_W = \frac{\eta^2}{2\Delta C^2} + \frac{\max\{-\mu_C, 0\}}{\Delta C} - \frac{|\rho|\eta\sigma}{2\Delta C \Delta D},$$

[1]AI-generated content: the descriptive text of this appendix (Claude Opus and Claude Fable, Anthropic). The tables are generated by the computational code.

Table C.2: Grid dimensions and uniform step sizes; $\Delta C = (C_{\max} - C_{\min})/N_C$ and $\Delta D = (D_{\max} - D_{\min})/N_D$. The multigrid warm-start interpolates a coarse solution onto the fine grid by bilinear interpolation before starting the fine-grid PSOR. Convergence-study grids use no warm-start or Richardson extrapolation.

Purpose	N_C	N_D	ΔC (\$/veh)	ΔD (\$/100 mi)
Convergence study (no warm-start, no Richardson)				
50×50	50	50	400	1.200
75×75	75	75	267	0.800
100×100	100	100	200	0.600
150×150	150	150	133	0.400
200×200	200	200	100	0.300
Baseline production (Richardson pair, $D \in [-30, 30]$)				
Coarse	200	300	100	0.200
Fine	400	600	50	0.100
Scenario H production (Richardson pair, $D \in [-30, 60]$)				
Coarse	200	300	100	0.300
Fine	400	600	50	0.150
Sensitivity analyses (ρ, κ): single grid, no Richardson				
Fine	400	600	50	0.100
η sensitivity, extended domain ($C \in [-13{,}000, 17{,}000]$, $D \in [-45, 45]$)				
Fine	400	600	75	0.150

$$a_N = \frac{\sigma^2}{2\Delta D^2} + \frac{\alpha}{\Delta D} - \frac{|\rho|\eta\sigma}{2\Delta C\Delta D}, \qquad a_S = \frac{\sigma^2}{2\Delta D^2} - \frac{|\rho|\eta\sigma}{2\Delta C\Delta D},$$

and the two diagonal weights are $|\rho|\eta\sigma/(2\Delta C\Delta D) > 0$. Discarding the drift terms, which are non-negative and only help, the binding requirement in the benefit direction is $\sigma^2/(2\Delta D^2) \geq |\rho|\eta\sigma/(2\Delta C\Delta D)$, that is

$$\frac{\Delta C}{\Delta D} \geq \frac{|\rho|\,\eta}{\sigma}, \tag{C.1}$$

with the analogous condition $\Delta C/\Delta D \leq \eta/(|\rho|\sigma)$ in the cost direction. At the baseline calibration ($\eta = \$2{,}320$, $\sigma = 2.42$, $|\rho| = 0.3$) the admissible window is $287.6 \leq \Delta C/\Delta D \leq 3{,}195$.

The production grids are laid out inside that window, at $\Delta C/\Delta D = 500$ on both members of the extrapolation pair, which satisfies (C.1) across the whole sensitivity range, including the largest cost volatility examined. A ratio below the lower bound turns the two benefit-direction weights a_N and a_S negative against a diagonal of order 2,100, so the operator is not an M-matrix and the convergence guarantee of Barles and Souganidis (1991) does not apply. Table C.3 records the check for every grid on which a reported result is computed. The condition holds on all of them, so the convergence result applies to every solution reported in the dissertation.

Sensitivity of the thresholds to the condition. How much the condition matters for the reported numbers can be measured directly. Solving to a maximum nodal change of 10^{-4} at $\Delta D = 0.30$, where (C.1) is violated, and at $\Delta D = 0.15$, where it holds, with the extraction method held fixed so that only the grid changes, gives

Table C.3: Positive-coefficient check. The south weight a_S is the binding one at this calibration. The requirement rises with η, from 287.6 at the baseline to 371.9 at η = 3,000; the baseline and extended-domain production ratios of 500 clear both, and the Scenario H and convergence-study grids run at 333.3, which clears the requirement at the baseline η.

Grid	$\Delta C/\Delta D$	a_S	Condition (C.1)
Coarse 200 × 300, ΔD = 0.20	500.0	+31.10	satisfied
Fine 400 × 600, ΔD = 0.10, η = 2,320	500.0	+124.39	satisfied
η runs, 400 × 600 extended domain, η = 3,000	500.0	+33.34	satisfied
Scenario H fine 400 × 600, ΔD = 0.15	333.3	+17.85	satisfied
Convergence study, 50 × 50 to 200 × 200	333.3	+0.28 to +4.46	satisfied

Table C.4: Optimal threshold $d^*(C_0)$ in \$/100mi under a doubling of the benefit resolution. The $\rho = 0$ row carries no mixed derivative and satisfies (C.1) at both spacings.

	ΔD = 0.30	ΔD = 0.15	Shift
ρ = +0.3	7.269	7.206	−0.063
ρ = 0 (no cross term)	7.624	7.581	−0.043
ρ = −0.3 (baseline)	7.963	7.937	−0.026
Correlation range $d^*(-0.3) - d^*(+0.3)$	0.694	0.731	+0.037

That the $\rho = 0$ case, which has no mixed derivative and is monotone at both spacings, shifts by more than the baseline does identifies the movement as ordinary discretisation error in the benefit direction rather than as an artefact of the lost monotonicity, and the correlation effect itself changes by about five per cent. Satisfying (C.1) therefore secures the theoretical guarantee without materially changing the reported thresholds; the choice that does move them is the boundary extraction of Subsection 3.5.2.

The lagged cross-diagonal. The mixed derivative contributes $+2|a_{CD}|$ to the diagonal of the discrete operator, but the implementation leaves that contribution on the right-hand side of the Gauss–Seidel update rather than folding it into the divisor. The fixed point is unchanged, since the term appears on both sides of the same equation, but the splitting differs: at $\Delta D = 0.30$ the absolute row sum of the iteration matrix at the baseline is 1.04 once the lagged centre term $2|a_{CD}|$ is counted (the six neighbour weights alone contribute 0.99), so the splitting carries no contraction guarantee where two axial weights are negative, and folding the term into the divisor would lower the row sum to 0.94 without restoring the sign structure. On the production grids every weight is positive, so no reported result depends on the choice.

Verification of the discrete operator. The implementation is checked by applying the assembled stencil to the test functions 1, c, d, c^2, d^2 and cd, whose derivatives are known exactly. The upwind drift terms reproduce $\mu_C V_c$ and αV_d in both the upwind and the downwind branch, up to the $O(h)$ artificial diffusion that upwinding introduces (1.4% of $\eta^2/2$ and 0.6% of $\sigma^2/2$ at the baseline on the production grid), and the mixed stencil (3.45) reproduces $\rho\eta\sigma V_{cd}$ for both signs of ρ, returning zero on the five test functions with no cross derivative.

C.3 Grid Convergence

Reference state: $(C_0, D_0) = (5{,}800,\ 1.45)$. The production fine grid (400×600, multigrid warm-start) yields $V_{\text{fine}} = \$211{,}684$; Richardson extrapolation against the 200×300 coarse solve gives $\hat{V} = \$211{,}659$, which is the value reported in Chapter 4. The convergence study below uses square grids without multigrid warm-start or Richardson extrapolation.

Table C.5: Richardson extrapolation on the production grid pair (200×300 coarse, 400×600 fine): $\hat{V} = (4V_{\text{fine}} - V_{\text{coarse}})/3$. Applied to the baseline Scenario G result and to the $t = 0$ slice of Scenario H; the intermediate time slices plotted in the boundary-evolution figure are single fine-grid solves, for which the correction at $t = 0$ is 0.012 \$/100mi, but *not* applied to sensitivity analyses (ρ, κ, η), which use a single-grid solve at 400×600. For sensitivity analyses only threshold *differences* across parameter values are reported, so absolute level accuracy is not required, and isotonic regression and the smoothing spline suffice.

Quantity	Correction (fine vs. $\hat{V}$)
$d^*(C_0)$ (\$/100 mi)	$\approx$0.008
$V(C_0, D_0)$ (\$)	\$25 (0.012%)

The 400×600 grid is well into the asymptotic regime.

The refinement study of Table C.5 varies the mesh at a fixed computational domain and therefore bounds discretisation error alone. Truncation error is assessed separately by re-solving on an extended domain, $D \in [-60, 30]$ with $N_D = 900$, holding both step sizes ($\Delta C = 50$, $\Delta D = 0.1$) and all solver settings fixed. At the reference state this moves $d^*(C_0)$ by 0.0007 \$/100 mi, an order of magnitude below the Richardson correction, and $V(C_0, D_0)$ by \$91, which is 0.7% of the option premium and 0.04% of the value function but is larger than the \$25 Richardson correction. Domain placement is therefore the binding source of error on the value function, while the reported threshold is insensitive to it. The domain placement of Table C.1 is therefore adequate.

C.4 PSOR Iteration Parameters

The tolerance is stated as an absolute change but is best read relative to the scale of the value function: at $V(C_0, D_0) \approx \$211{,}000$, a maximum nodal change of 10^{-3} is a relative change of about 5×10^{-9}. The iteration cap is set high enough that every run reported here terminates on the tolerance rather than on the cap. The threshold is not

Table C.6: PSOR parameters for the infinite-horizon solver.

Parameter	Symbol	Value	Notes
Initial relaxation factor	ω_0	1.45	Adapted downward when residual stalls; active on the warm-started fine solve only
Adaptive decrement	$\delta\omega$	0.005	$\omega \leftarrow \max(1.0,\ \omega - 0.005)$ per stall
Minimum relaxation factor	$\omega_{\min}$	1.0	—
Convergence tolerance	ε	10^{-3}	$\|V^{(m+1)} - V^{(m)}\|_\infty$; $\approx 5 \times 10^{-9}$ relative
Maximum iterations	—	150,000	Not binding; runs stop on ε

fully settled at this tolerance: tightening it to 10^{-4} moves $d^*(C_0)$ by 0.07 \$/100mi, which is larger than the Richardson correction of 0.008 and of the same order as the extraction corridor, and is the reason the corridor rather than the extrapolation residual is quoted as the error bar on a reported threshold.

Table C.7: PSOR parameters for the finite-horizon solver (Scenario H).

Parameter	Symbol	Value	Notes
Relaxation factor	ω	1.2	Fixed, no adaptive decrement
Convergence tolerance (per step)	ε	10^{-9}	$\|V^{(m+1)} - V^{(m)}\|_\infty$
Max iterations per time step	—	5,000	—
Time steps ($T = 5$ yr, 50 steps/yr)	N_T	250	$\Delta t = 0.02$ yr

Homotopy schedule. For infinite-horizon solves started from scratch, the mixed cross-derivative term (proportional to $\rho\eta\sigma V_{CD}$) is omitted during the first 99 PSOR iterations (`it < 100`); the full operator is restored at iteration 100 and iteration continues to convergence. Warm-started fine-grid solves (multigrid) retain the full operator throughout, since the interpolated coarse solution already reflects the mixed term. For finite-horizon solves, east/west boundary conditions are blended linearly over the first $N_{\text{ramp}} = 10$ time steps (weight $\min(1,\ k/N_{\text{ramp}})$), allowing the boundary layer to stabilise before the full free-boundary formulation is enforced.

C.5 Boundary Extraction, Validation, and Cost

Boundary extraction. The optimal-adoption boundary is not read off the value surface node by node. Near the free boundary the continuation premium $V - \hat{g}$ decays like the square of the distance to the boundary, so $\sqrt{V - \hat{g}}$ is fitted linearly in d within each cost column and the root of that line is taken as the boundary estimate; the column-wise estimates are then made monotone and smoothed. Table C.8 collects the settings of each step, together with the corridor half-width quoted as the error bar on a reported threshold.

Table C.8: Parameters for the optimal-adoption boundary extraction procedure described in Subsection 3.4.2.

Step	Parameter	Value
Fit window	Band of $V - \hat{g}$ used, relative to the column maximum	$[10^{-3}, 0.25]$
Fit	Least squares of $\sqrt{V - \hat{g}}$ on d	Last 25 nodes in the band
Boundary estimate	Root of the fitted line	$d_{\text{raw}}(c)$
Fit diagnostic	Relative residual of the fit	≈1% at production resolution
Monotonicity enforcement	Isotonic regression ($b'(c) \geq 0$)	—
Smoothing	UnivariateSpline	$s = n$ (n = # valid points, $1.5n$ finite)
Uncertainty corridor (half-width)	Grid ⊕ smoothing spread	0.05 \$/100 mi

Corridor half-width 0.05 \$/100 mi across the central cost range, at most 0.09 at the domain endpoints. A gradient-based term is excluded from the corridor: it is evaluated where smooth pasting makes the derivative vanish. Columns where the fit window holds fewer than five nodes, or where the fitted slope is non-negative, fall back to the first node at which the premium is no longer positive; columns with fewer than four positive-premium nodes are skipped. On the production surface none of the three branches is taken.

Monte Carlo validation. The finite-horizon boundary is cross-checked against a least-squares Monte Carlo solve that shares no discretisation with the finite-difference scheme, so the two methods fail in different ways if either is wrong. Table C.9 records the parameters of that run.

Table C.9: Finite-horizon Longstaff–Schwartz MC parameters used for the plotted boundary validation (Figure 30). Paths are seeded uniformly over the plotting rectangle rather than from the initial state, which is what makes a boundary rather than a single value recoverable, and the regression is fitted on in-the-money paths only. The reported boundary is the zero crossing of $\hat{g}^L - X\hat{\beta}$ at the time-zero coefficient vector.

Parameter	Value	Notes
Paths M	100,000	Figure run, see caption for 500k cross-check
Time step Δt	0.01 yr	—
Algorithm	Longstaff–Schwartz	American option regression
Regression basis	Degree-2 polynomial	$\{1, C_s, D_s, C_s^2, C_s D_s, D_s^2\}$
Feature scaling	$C_s = C/1000$, $D_s = D/10$	Improves conditioning
Min. in-the-money paths	11	Regression skipped at 10 or fewer
Regression method	OLS	`np.linalg.lstsq`

Computational cost. Table C.10 summarises the size of every solve reported in the thesis, so that the resolutions and iteration caps quoted in Chapters 3 and 4 can be placed against the effort each one costs.

Table C.10: Approximate computational cost. All finite-difference computations use Python with Numba JIT compilation for the inner PSOR sweep.

Run type	Grid $N_C \times N_D$	Iter./paths	Steps	Notes
Convergence study (each)	50–200×50–200	≤150,000	—	5 resolutions
Baseline	400×600	≤150,000	—	Richardson 200 × 300→400×600
Sensitivity (ρ, κ, η)	400×600	≤150,000	—	Single-grid, no Richardson; η runs on the extended domain
Scenario H	400×600	≤5,000/step	250	Richardson 200 × 300→400×600
MC validation (finite, LSM)	—	100,000	500	Figure run, 500k cubic cross-check separate

BIBLIOGRAPHY

Andrew B. Abel and Janice C. Eberly. Optimal investment with costly reversibility. *The Review of Economic Studies*, 63(4):581–593, 1996. doi: 10.2307/2297794.

Ilyes Abid, Stéphane Goutte, Khaled Guesmi, and Ibrahim Jamali. Transmission of shocks and contagion from us to mena equity markets: The role of oil and gas markets. *Energy Policy*, 134:110953, 2019. doi: 10.1016/j.enpol.2019.110953.

ACEA. The automobile industry pocket guide 2025/2026. Technical report, European Automobile Manufacturers' Association, September 2025. URL `https://www.acea.auto/files/ACEA-Pocket-Guide-2025-2026.pdf`.

Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American Economic Review*, 102(1): 131–166, 2012. doi: 10.1257/aer.102.1.131.

Roger Adkins and Dean Paxson. Subsidies for renewable energy facilities under uncertainty. *The Manchester School*, 84(2):222–250, 2016. doi: 10.1111/manc.12093. URL `https://onlinelibrary.wiley.com/doi/pdf/10.1111/manc.12093`.

Zack Aemmer, Daniel Malarkey, and Don MacKenzie. Emissions reductions from electrifying high-mileage vehicles. *Findings*, 2023. doi: 10.32866/001c.75133.

Philippe Aghion, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51, 2016. doi: 10.1086/684581.

Vishal V Agrawal and Ioannis Bellos. The potential of servicizing as a green business model. *Management Science*, 63(5):1545–1562, 2017. doi: 10.1287/mnsc.2015.2399.

Fernando Aguilar Lopez, Dirk Lauinger, François Vuille, and Daniel B Müller. On the potential of vehicle-to-grid and second-life batteries to provide energy and material security. *Nature communications*, 15(1):4179, 2024. doi: 10.1038/s41467-024-48554-0.

Larbi Alili and Andreas E. Kyprianou. Some remarks on first passage of Lévy processes, the American put and pasting principles. *The Annals of Applied Probability*, 15(3):2062–2080, 2005. doi: 10.1214/105051605000000377.

Hunt Allcott and Michael Greenstone. Is there an energy efficiency gap? *Journal of Economic perspectives*, 26(1):3–28, 2012. doi: 10.1257/jep.26.1.3.

Hunt Allcott and Judd B Kessler. The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics*, 11(1):236–76, 2019. doi: 10.1257/app.20170328.

Eugene L. Allgower and Kurt Georg. Numerical continuation methods: An introduction. *Springer Series in Computational Mathematics*, 13, 2012. doi: 10.1007/978-3-642-61257-2.

Osman Alp, Tarkan Tan, and Maximiliano Udenio. Transitioning to sustainable freight transportation by integrating fleet replacement and charging infrastructure decisions. *Omega*, 109:102595, 2022. doi: 10.1016/j.omega.2022.102595.

Amir H Ansaripoor, Fernando S Oliveira, and Anne Liret. Recursive expected conditional value at risk in the fleet renewal problem with alternative fuel vehicles. *Transportation Research Part C: Emerging Technologies*, 65:156–171, 2016. doi: 10.1016/j.trc.2015.12.010.

Volodymyr Babich, Ruben Lobel, and Şafak Yücel. Promoting solar panel investments: Feed-in-tariff vs. tax-rebate policies. *Manufacturing & Service Operations Management*, 22(6):1148–1164, 2020. doi: 10.1287/msom.2019.0860.

Guy Barles and Panagiotis E Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4(3):271–283, 1991. doi: 10.3233/asy-1991-4305.

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`.

Stacie E Beck. Cointegration and market efficiency in commodities futures markets. *Applied economics*, 26(3):249–257, 1994. doi: 10.1080/00036849400000006.

BloombergNEF. Battery price survey 2024. `https://about.bnef.com/blog/battery-pack-prices-fall-to-an-average-of-139-per-kilowatt-hour-in-2024/`, 2024. Accessed January 2025.

BloombergNEF. Electric vehicle outlook 2025: Executive summary. Industry report, Bloomberg Finance L.P., June 2025. URL `https://about.bnef.com/electric-vehicle-outlook/`.

Severin Borenstein, Meredith Fowlie, and James Sallee. Paying for electricity in california: How residential rate design impacts equity and electrification. Technical report, Energy Institute at Haas, University of California, Berkeley, September 2022. URL `https://haas.berkeley.edu/wp-content/uploads/archive/WP330.pdf`. Energy Institute at Haas Working Paper 330.

Andrei N Borodin and Paavo Salminen. *Handbook of Brownian motion-facts and formulae*. Springer Science & Business Media, 2015.

Michael J Brennan and Eduardo S Schwartz. Finite difference methods and jump processes arising in the pricing of contingent claims: A synthesis. *Journal of Financial and Quantitative Analysis*, 13(3):461–474, 1978. doi: 10.2307/2330152.

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

Ricardo J. Caballero. Aggregate investment. In *Handbook of Macroeconomics*, volume 1, pages 813–862. Elsevier, 1999.

Min-Je Cho and Juneseuk Shin. Identifying a combination of key resources to overcome the entry barriers in the electric vehicle market. *IEEE Access*, 10: 60373–60386, 2022. doi: 10.1109/access.2022.3180729.

Robert M Corless, Gaston H Gonnet, David E G Hare, David J Jeffrey, and Donald E Knuth. On the Lambert W function. *Advances in Computational Mathematics*, 5 (1):329–359, 1996. doi: 10.1007/BF02124750.

Richard W Cottle, Jong-Shi Pang, and Richard E Stone. *The Linear Complementarity Problem*. Classics in Applied Mathematics. SIAM, Philadelphia, 2009.

Robin Cowan and Staffan Hultén. Escaping lock-in: the case of the electric vehicle. *Technological Forecasting and Social Change*, 53(1):61–79, 1996. doi: 10.1016/0040-1625(96)00059-5.

C. W. Cryer. The solution of a quadratic programming problem using systematic overrelaxation. *SIAM Journal on Control*, 9(3):385–392, 1971. doi: 10.1137/0309027.

Leila Dagher and Fakhri J Hasanov. Oil market shocks and financial instability in asian countries. *International Review of Economics & Finance*, 84:182–195, 2023. doi: 10.1016/j.iref.2022.11.008.

Leila Dagher, Ibrahim Jamali, and Oussama Abi Younes. Extreme energy poverty: The aftermath of lebanon‘s economic collapse. *Energy Policy*, 183:113783, 2023. doi: 10.1016/j.enpol.2023.113783.

You Daming, Yang Xiaohui, Desheng Dash Wu, and Chen Guofan. Option game with poisson jump process in company radical technological innovation. *Technological Forecasting and Social Change*, 81:341–350, 2014. doi: 10.1016/j.techfore.2013.04.012.

Donald A Darling, Thomas Liggett, and Howard M Taylor. Optimal stopping for partial sums. *The Annals of Mathematical Statistics*, 43(4):1363–1368, 1972. doi: 10.1214/aoms/1177692491.

Lucas W Davis. How much are electric vehicles driven? *Applied economics letters*, 26(18):1497–1502, 2019. doi: 10.1080/13504851.2019.1582847.

Savas Dayanik and Ioannis Karatzas. On the optimal stopping problem for one-dimensional diffusions. *Stochastic Processes and Their Applications*, 107(2): 173–212, 2003. doi: 10.1016/s0304-4149(03)00076-0.

Carl de Boor. *A Practical Guide to Splines*. Springer-Verlag, New York, 1978.

Melissa Demartini, Marco Ferrari, Kannan Govindan, and Flavio Tonelli. The transition to electric vehicles and a net zero economy: A model based on circular economy, stakeholder theory, and system thinking approach. *Journal of Cleaner Production*, 410:137031, 2023. doi: 10.1016/j.jclepro.2023.137031.

Jérôme Detemple. *American-Style Derivatives: Valuation and Computation*. Chapman and Hall/CRC, Boca Raton, FL, 2006. ISBN 9781584885672.

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

Daniel J Duffy. *Finite difference methods in financial engineering: a partial differential equation approach*. John Wiley & Sons, 2013.

Anupam Dutta, Elie Bouri, Tareq Saeed, and Xuan Vinh Vo. Impact of energy sector volatility on clean energy assets. *Energy*, 212:118657, 2020. doi: 10.1016/j.energy.2020.118657.

Olga Efimova and Apostolos Serletis. Energy markets volatility modelling using garch. *Energy Economics*, 43:264–273, 2014. doi: 10.1016/j.eneco.2014.02.018.

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.

Lei Fang, Yanlin Li, and Kannan Govindan. Entry mode selection for a new entrant of the electric vehicle automaker. *European Journal of Operational Research*, 313(1):270–280, 2024. doi: 10.1016/j.ejor.2023.08.014.

Wei Feng and Miguel Figliozzi. An economic and technological analysis of the key factors affecting the competitiveness of electric commercial vehicles: A case study from the usa market. *Transportation Research Part C: Emerging Technologies*, 26:135–145, 2013. doi: 10.1016/j.trc.2012.06.007.

Paul A. Geroski. Models of technology diffusion. *Research Policy*, 29(4-5):603–625, 2000. doi: 10.1016/S0048-7333(99)00092-X.

Mike B. Giles and Richard Carter. Convergence analysis of crank-nicolson and rannacher time-marching. *The Journal of Computational Finance*, 9(4):89–112, 2006. doi: 10.21314/jcf.2006.152.

Paul Glasserman. *Monte Carlo Methods in Financial Engineering*, volume 53. Springer Science & Business Media, 2003.

Nico Goetzel and M. Hasanuzzaman. An empirical analysis of electric vehicle cost trends: A case study in germany. *Research in Transportation Business & Management*, 43:100825, 2022. doi: 10.1016/j.rtbm.2022.100825. URL `https://linkinghub.elsevier.com/retrieve/pii/S2210539522000463`.

Gary Haq and Martin Weiss. Time preference and consumer discount rates-insights for accelerating the adoption of efficient energy and transport technologies. *Technological Forecasting and Social Change*, 137:76–88, 2018. doi: 10.1016/j.techfore.2018.06.045.

Hassan Harajli, Vahakn Kabakian, Jad El-Baba, Ahmad Diab, and Carla Nassab. Commercial-scale hybrid solar photovoltaic-diesel systems in select arab countries with weak grids: an integrated appraisal. *Energy Policy*, 137:111190, 2020. doi: 10.1016/j.enpol.2019.111190.

Michael K Hidrue, George R Parsons, Willett Kempton, and Meryl P Gardner. Willingness to pay for electric vehicles and their attributes. *Resource and Energy Economics*, 33(3):686–705, 2011. doi: 10.1016/j.reseneeco.2011.02.002.

Winfred B Hirschmann. Profit from the learning curve. *Harvard Business Review*, 42:125, 1964.

Willem Hundsdorfer and Jan G Verwer. *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*, volume 33. Springer Science & Business Media, 2013.

Samuli Ikonen and Jari Toivanen. Operator splitting methods for pricing american options under stochastic volatility. *Numerische Mathematik*, 113:299–324, 2009. doi: 10.1007/s00211-009-0227-5.

ING Research. Us electric vehicle transition delayed but not derailed by subsidy exit. Technical report, ING Bank, 2025.

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.

International Energy Agency. Global ev outlook 2023. Technical report, IEA, Paris, 2023.

International Energy Agency. World energy outlook 2024. Technical report, IEA, Paris, 2024. URL `https://www.iea.org/reports/world-energy-outlook-2024`. OECD/IEA Publication.

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`.

Henrik Klinge Jacobsen and Erika Zvingilaite. Reducing the market impact of large shares of intermittent energy in denmark. *Energy Policy*, 38(7):3403–3413, 2010. doi: 10.1016/j.enpol.2010.02.014.

Adam B Jaffe and Robert N Stavins. The energy-efficiency gap what does it mean? *Energy policy*, 22(10):804–810, 1994. doi: 10.1016/0301-4215(94)90138-4.

Boyan Jovanovic. Selection and the evolution of industry. *Econometrica*, 50(3): 649–670, May 1982. doi: 10.2307/1912606.

Raul Kangro and Roy Nicolaides. Far field boundary conditions for black–scholes equations. *SIAM Journal on Numerical Analysis*, 38(4):1357–1368, 2000. doi: 10.1137/s0036142999355921.

Selin Karagoz, Nesrin Aydin, and Vladimir Simic. End-of-life vehicle management: a comprehensive review. *Journal of Material Cycles and Waste Management*, 22 (2):416–442, 2020. doi: 10.1007/s10163-019-00945-y.

Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*. springer, 2014.

Janina C Ketterer. The impact of wind power generation on the electricity price in germany. *Energy Economics*, 44:270–280, 2014. doi: 10.1016/j.eneco.2014.04.003.

Yerkin Kitapbayev, John Moriarty, and Pierluigi Mancarella. Stochastic control and real options valuation of thermal storage-enabled demand response from flexible district energy systems. *Applied Energy*, 137:823–831, 2015. doi: 10.1016/j.apenergy.2014.07.019.

Noah Kittner, Felix Lill, and Daniel M Kammen. Energy storage deployment and innovation for the clean energy transition. *Nature Energy*, 2(9):1–6, 2017. doi: 10.1038/nenergy.2017.125.

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.

Florian Knobloch, Steef V Hanssen, Aileen Lam, Hector Pollitt, Pablo Salas, Unnada Chewpreecha, Mark AJ Huijbregts, and Jean-Francois Mercure. Net emission reductions from electric cars and heat pumps in 59 world regions over time. *Nature Sustainability*, 3(6):437–447, 2020. doi: 10.1038/s41893-020-0488-7.

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.

Steven G Kou and Hui Wang. First passage times of a jump diffusion process. *Advances in Applied Probability*, 35(2):504–531, 2003. doi: 10.1239/aap/1051201658.

Aileen Lam and Jean-Francois Mercure. Which policy mixes are best for decarbonising passenger cars? simulating interactions among taxes, subsidies and regulations for the united kingdom, the united states, japan, china, and india. *Energy Research & Social Science*, 75:101951, 2021. doi: 10.1016/j.erss.2021.101951.

Randall J. LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1st edition, 2007. ISBN 978-0898716299.

Feng Liu, Yingshuang Tan, Sudipto Sarkar, Xueqing Zhang, and Xingjun Huang. When to invest in electric vehicles under dual credit policy: A real options approach. *Managerial and Decision Economics*, 44(4):2186–2198, 2023. doi: 10.1002/mde.3811.

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.

Nigel Meade. Oil prices-brownian motion or mean reversion? a study using a one year ahead density forecast criterion. *Energy Economics*, 32(6):1485–1498, 2010. doi: 10.1016/j.eneco.2010.07.010.

J-F Mercure, Aileen Lam, Sophie Billington, and Hector Pollitt. Integrated assessment modelling as a positive science: private passenger road transport policies to meet a climate target well below 2°c. *Climatic Change*, 151:109–129, 2018. doi: 10.1007/s10584-018-2262-7.

Matthew Metz and Janelle London. Gasoline superusers. Technical report, Coltura, July 2021. URL `https://coltura.org/gasoline-superusers/`.

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

Ernesto Mordecki. Optimal stopping and perpetual options for Lévy processes. *Finance and Stochastics*, 6(4):473–493, 2002b. doi: 10.1007/s007800200070.

M Nijhuis, M Gibescu, and JFG Cobben. Assessment of the impacts of the renewable energy and ict driven energy transition on distribution networks. *Renewable and Sustainable Energy Reviews*, 52:1003–1014, 2015. doi: 10.1016/j.rser.2015.07.124.

Alexander A Novikov and Albert N Shiryaev. On an effective solution of the optimal stopping problem for random walks. *Theory of Probability & Its Applications*, 49(2):344–354, 2004. doi: 10.1137/S0040585X97981093.

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.

Bernt Øksendal and Agnes Sulem. *Applied stochastic control of jump diffusions*, volume 3. Springer, 2007.

Ugur Korkut Pata, Selin Karlilar, and Mustafa Tevfik Kartal. On the road to sustainable development: the role of ict and r&d investments in renewable and nuclear energy on energy transition in germany. *Clean Technologies and Environmental Policy*, 26(7):2323–2335, 2024. doi: 10.1007/s10098-023-02677-y.

Samuel Pelletier, Ola Jabali, Jorge E Mendoza, and Gilbert Laporte. The electric bus fleet transition problem. *Transportation Research Part C: Emerging Technologies*, 109:174–193, 2019. doi: 10.1016/j.trc.2019.10.012.

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.

Robert S. Pindyck. Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, 29(3):1110–1148, 1991.

David M Pooley, Peter A Forsyth, and Kenneth R Vetzal. Numerical convergence properties of option pricing pdes with uncertain volatility. *IMA Journal of Numerical Analysis*, 23(2):241–267, 2003. doi: 10.1093/imanum/23.2.241.

Vinay Ramani, Debabrata Ghosh, and ManMohan S Sodhi. Understanding systemic disruption from the covid-19-induced semiconductor shortage for the auto industry. *Omega*, 113:102720, 2022. doi: 10.1016/j.omega.2022.102720.

Eva Regnier. Oil and energy price volatility. *Energy Economics*, 29(3):405–427, 2007. doi: 10.1016/j.eneco.2005.11.003.

Lewis Fry Richardson and John Arthur Gaunt. The deferred approach to the limit. *Philosophical Transactions of the Royal Society of London. Series A*, 226(636-646):299–342, 1927. doi: 10.1098/rsta.1927.0008.

Mohammadhosein Safari. Battery electric vehicles: Looking behind to move forward. *Energy Policy*, 115:54–65, 2018. doi: 10.1016/j.enpol.2017.12.053.

Michael Saidani, Alissa Kendall, Bernard Yannou, Yann Leroy, and François Cluzel. Management of the end-of-life of light and heavy vehicles in the us: comparison with the european union in a circular economy perspective. *Journal of Material Cycles and Waste Management*, 21:1449–1461, 2019. doi: 10.1007/s10163-019-00897-3.

Eduardo Schwartz and James E Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000. doi: 10.1287/mnsc.46.7.893.12034.

Rüdiger Seydel. *Tools for Computational Finance*. Springer, 2017.

Zuo-Jun Max Shen, Bo Feng, Chao Mao, and Lun Ran. Optimization models for electric vehicle service operations: A literature review. *Transportation Research Part B: Methodological*, 128:462–477, 2019. doi: 10.1016/j.trb.2019.08.006.

Lingling Shi and Bin Hu. Battery as a service: Flexible electric vehicle battery leasing. *Forthcoming at Manufacturing and Service Operations Management*, 2022. doi: 10.1287/msom.2022.0587.

Aljoša Slameršak, Giorgos Kallis, and Daniel W O'Neill. Energy requirements and carbon emissions for a low-carbon energy transition. *Nature Communications*, 13 (1):6932, 2022. doi: 10.1038/s41467-022-33976-5.

Vi Kie Soo, Jef Peeters, Paul Compston, Matthew Doolan, and Joost R Duflou. Comparative study of end-of-life vehicle recycling in australia and belgium. *Procedia CIRP*, 61:269–274, 2017. doi: 10.1016/j.procir.2016.11.222.

Vi Kie Soo, Matthew Doolan, Paul Compston, Joost R Duflou, Jef Peeters, and Yasushi Umeda. The influence of end-of-life regulation on vehicle material circularity: A comparison of europe, japan, australia and the us. *Resources, Conservation and Recycling*, 168:105294, 2021. doi: 10.1016/j.resconrec.2020.105294.

Lars Stentoft. Assessing the least squares monte-carlo approach to american option valuation. *Review of Derivatives research*, 7(2):129–168, 2004. doi: 10.1023/B:REDR.0000031176.24759.e6.

John C. Strikwerda. *Finite Difference Schemes and Partial Differential Equations*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2nd edition, 2004. ISBN 978-0898716294.

Dávid Zoltán Szabó and Randall Martyr. Real option valuation of a decremental regulation service provided by electricity storage. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375 (2100):20160300, 2017. doi: 10.1098/rsta.2016.0300.

Dávid Zoltán Szabó, Péter Csóka, and Réka Janosik. The optimal timing of clean technology adoption: A stochastic cost–benefit analysis. *Technological Forecasting & Social Change*, 219:124276, 2025. doi: 10.1016/j.techfore.2025.124276.

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.

Domingo Tavella and Curt Randall. *Pricing Financial Instruments: The Finite Difference Method*. John Wiley & Sons, 2000.

VJ Thomas and Elicia Maine. Market entry strategies for electric vehicle start-ups in the automotive industry–lessons from tesla motors. *Journal of Cleaner Production*, 235:653–663, 2019. doi: 10.1016/j.jclepro.2019.06.284.

Lenos Trigeorgis and Andrianos E Tsekrekos. Real options in operations research: A review. *European Journal of Operational Research*, 270(1):1–24, 2018. doi: 10.1016/j.ejor.2017.11.055.

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.

Christian Urom, Khaled Guesmi, Ilyes Abid, and Leila Dagher. Dynamic integration and transmission channels among interest rates and oil price shocks. *The Quarterly Review of Economics and Finance*, 87:296–317, 2023. doi: 10.1016/j.qref.2021.04.008.

U.S. Geological Survey. Mineral commodity summaries 2025. `https://pubs.usgs.gov/publication/mcs2025`, January 2025. Lithium chapter prepared by Brian W. Jaskula.

Mengyu Wang, Jeffrey Wurgler, and Hong Zhang. Policy uncertainty reduces green investment. Technical report, National Bureau of Economic Research, 2023.

Claire Weiller. Electric vehicle rental services: Project in Okinawa, Japan. Technical report, Cambridge Service Alliance, Institute for Manufacturing (IfM), University of Cambridge, Cambridge, UK, 2014. URL `https://www.ifm.eng.cam.ac.uk/uploads/Resources/Okinawa_HQP.pdf`.

M Weiss, A Zerfass, and E Helmers. Fully electric and plug-in hybrid cars - an analysis of learning rates, user costs, and costs for mitigating co2 and air pollutant emissions. *Journal of Cleaner Production*, 212:1478–1489, 2019. doi: 10.1016/j.jclepro.2018.12.019.

Martin Weiss, Martin K Patel, Martin Junginger, Adolfo Perujo, Pierre Bonnel, and Geert Van Grootveld. On the electrification of road transport-learning rates and price forecasts for hybrid-electric and battery-electric vehicles. *Energy Policy*, 48:374–393, 2012. doi: 10.1016/j.enpol.2012.05.038.

Martin Weiss, Peter Dekker, Alberto Moro, Harald Scholz, and Martin K Patel. On the electrification of road transportation–a review of the environmental, economic, and social performance of electric two-wheelers. *Transportation research part D: transport and environment*, 41:348–366, 2015. doi: 10.1016/j.trd.2015.09.007.

JH Wesseling, J Faber, and MP Hekkert. How competitive forces sustain electric vehicle development. *Technological Forecasting and Social Change*, 81:154–164, 2014. doi: 10.1016/j.techfore.2013.02.005.

Paul Wilmott. *Paul Wilmott on Quantitative Finance*. John Wiley & Sons, 2nd edition, 2006.

Paul Wilmott, Sam Howison, and Jeff Dewynne. *The Mathematics of Financial Derivatives: A Student Introduction*. Cambridge University Press, 1995.

Jonas Winkelmann, Stefan Spinler, and Thomas Neukirchen. Green transport fleet renewal using approximate dynamic programming: A case study in german heavy-duty road transportation. *Transportation Research Part E: Logistics and Transportation Review*, 186:103547, 2024. doi: 10.1016/j.tre.2024.103547.

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.

THE AUTHOR'S 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